

CHARACTERISING THE ORBITS OF LONG PERIOD EXOPLANETS

Diana Dragomir

Master of Science

Department of Physics

McGill University

Montreal, Quebec

August 22, 2008

A thesis submitted to McGill University in partial fulfilment of the
requirements of the degree of Master of Science

©Diana Dragomir 2008

ABSTRACT

To date nearly 300 exoplanets have been discovered, most of them through measurements of the wobble induced by the planet in the motion of its host star. We have developed a program, based on Bayesian inference, to fit eccentric Keplerian orbits to exoplanet radial velocity data. The data consist of optical spectra obtained using the HIRES echelle spectrometer on the Keck I telescope. We have applied the program to 58 sets of measurements. We have obtained probability distributions for the orbital period and eccentricity for each set. We have found that clear upper and lower limits can be placed on the period while the eccentricity proves more difficult to constrain. From the average period probability distribution we prepared, we preliminarily concluded that there is a much higher probability to find exoplanets with periods below 10000 days than with periods above 10000 days. We also suspect the existence of a correlation between the period and the eccentricity, as well as that of a possible trend of decreasing period with increasing stellar metallicity.

ABRÉGÉ

À date presque 300 exoplanètes ont été découvertes, la plupart à travers des mesures du vacillement provoqué par la planète dans le mouvement de son étoile mère. Nous avons mis au point un programme, basé sur l'inférence bayésienne, pour modéliser des données de vitesse radiale pour la recherche d'exoplanètes par des orbites excentriques képlériennes. Les données sont des spectres optiques et ont été recueillies avec le spectromètre échelle HIRES au télescope Keck I. Nous avons utilisé le programme pour analyser 58 ensembles de données. Nous avons obtenu des distributions de probabilité pour la période orbitale ainsi que pour l'excentricité pour chaque ensemble. Nous avons trouvé qu'il est possible de placer des limites supérieures et inférieures sur la période, mais que l'excentricité est plus difficile à contraindre. À partir de la moyenne des distributions de probabilité que nous avons obtenu pour la période, nous

avons préliminairement conclu qu'il y a une beaucoup plus grande probabilité de trouver des exoplanètes avec des périodes inférieures à 10000 jours qu'avec des périodes supérieures à 10000 jours. Nous suspectons la présence d'une corrélation entre la période et l'excentricité, ainsi qu'une possible tendance de la période à décroître à mesure que la métallicité stellaire augmente.

ACKNOWLEDGEMENTS

I want to thank my supervisor, Andrew Cumming, for his constant support and encouragement, as well as his availability and patience. He has helped me through both programming difficulties and endless questioning of basic principles, some of which he has explained to me countless times with the same indulgence. I also wish to thank Paul Mercure for assistance with everything from (computer) networking to printing. Ian Dobbs-Dixon has provided me with many expert tips and advice on IDL programming, without which I would have spent weeks searching in the dark. I am grateful to Hugo Olivares, Michael Zamfir and Eric Ting-kuei Chou (fellow research group members) for both useful and entertaining conversations as needed.

Contents

List of Figures	6
1 Introduction	11
1.1 Definitions	13
1.2 A Note on Size Conventions	14
1.3 Celestial Mechanics	14
1.3.1 Model of the Planetary System	15
1.4 Exoplanet Detection Methods	17
1.5 Observed Properties of Exoplanets	21
1.6 Planet Formation and Evolution	25
1.6.1 Formation	25
1.6.2 Migration	27
1.6.3 Orbit Alterations	28
2 Bayesian Inference at Work in Exoplanet Orbit Analysis	30
2.1 Introduction to Bayesian Inference	30
2.2 Markov Chain Monte Carlo	33
2.3 Parallel Tempering	35
2.4 Non-Bayesian Statistical Methods	36
2.5 Motivation for this Thesis Work	37
3 Method	40
3.1 Circular Orbits	40
3.1.1 Choice of Priors	41

3.1.2	Calculating the Probabilities	42
3.1.3	Marginalization	42
3.1.4	The Odds Ratio	49
3.1.5	Improving the Efficiency of the Calculations	50
3.2	Eccentric Orbits	51
3.2.1	Choice of Priors	52
3.2.2	Kepler's Equation	53
3.2.3	Analytic Solution for C_{bf}	53
3.2.4	Calculating the Probabilities	54
3.2.5	Marginalization	55
3.2.6	Running Time Factors and Limitations	57
4	Results and Discussion	59
4.1	Comparison with Previous Work	60
4.2	Probability Distributions	61
4.3	A Closer Look at the Period-Eccentricity Probability Distributions	67
4.4	Average Period Probability Distribution	70
4.5	Period-Metallicity Trend	72
5	Conclusion and Future Work	79
	Bibliography	81

List of Figures

1.1	Keplerian orbital parameters (i , ω and Ω). In the context of exoplanet orbits, the reference plane is the plane of the sky. Image is courtesy of <i>Eric Weisstein's World of Physics</i>	16
1.2	Velocity components of a star.	19
1.3	Minimum mass distribution for the 292 known nearby exoplanets with $m \sin i < 15M_J$. We produced this figure interactively at <i>exoplanet.eu</i>	22
1.4	Distribution of orbital periods for the 103 known nearby hot Jupiters. We produced this figure interactively at <i>exoplanet.eu</i>	23
1.5	Minimum mass distribution of the 276 known nearby exoplanets with a < 6.5 AU. We produced this figure interactively at <i>exoplanet.eu</i>	23
1.6	Distribution of orbital eccentricities as a function of minimum mass for the 253 known nearby exoplanets with $m \sin i < 13M_J$ for which the eccentricity is known. We produced this figure interactively at <i>exoplanet.eu</i>	24
2.1	All confirmed planets (48) and unconfirmed candidates (76) flagged among the 585 stars analysed. The open circles represent the planets and the solid triangles represent the candidates. The vertical dashed line corresponds to the duration of the survey (8 years). The dotted lines correspond to exoplanet masses of $m \sin i = 0.1, 1$ and $10M_J$, for a solar mass star. This figure is used with the permission of A. Cumming.	38

3.1	Two overlapping probability distributions for K for the same radial velocity data set. The top curve was obtained with ∞ as the upper limit for the marginalization over σ_i ; the bottom curve was obtained with 6 m/s as the upper limit.	44
3.2	Probability distribution for K : for the top curve, we used direct numerical integration over ϕ ; for the bottom curve, we used a Gaussian approximation to the integrand.	47
3.3	The solid curve is a typical example of the probability distribution for ϕ at low values of K . The dashed curve is the fitted Gaussian, which includes more probability than the actual distribution.	48
3.4	The solid curve is a typical example of the probability distribution for ϕ at various combinations of values of K , P and C . The dashed curve is the fitted Gaussian, which only includes one of the peaks.	49
3.5	Probability distribution for ϕ , with $e = \mathbf{0.5}$, $K = 31.7$ m/s, $P = 59.6$ days and $\omega = \pi/3$	56
3.6	Probability distribution for ϕ , with $e = \mathbf{0.7}$, $K = 31.7$ m/s, $P = 59.6$ days and $\omega = \pi/3$	56
3.7	Probability distribution for ϕ , with $e = \mathbf{0.9}$, $K = 31.7$ m/s, $P = 59.6$ days and $\omega = \pi/3$	56
4.1	Probability distribution marginalized over all parameters except the P and e , for HD 72659. The solid contours indicate the 1, 2 and 3 σ confidence intervals, corresponding to 68.3%, 95.4% and 99.73% of the probability distribution. a) Our contour plot. b) The contour plot taken from [27] (the dotted contours represent 1, 2 and 3 σ confidence intervals obtained through a different method, and the dot represents the best-fit orbital solution based on later observations made after the rest of the plot had been composed). The plot in subfigure b) is used with the permission of E.B. Ford.	61

4.2	Radial velocity measurements and $p_{10/200}$ values (see text for details) for six unconfirmed candidates.	62
4.3	Period probability distributions for six unconfirmed candidates.	63
4.4	Eccentricity probability distributions for six unconfirmed candidates.	64
4.5	Period-eccentricity probability distributions for six unconfirmed candidates. The contours indicate the 1, 2 and 3 σ confidence intervals, corresponding to 68.3%, 95.4% and 99.73% of the probability distribution.	65
4.6	Period-eccentricity probability distributions, showing a) two branches and b) a single branch.	68
4.7	Period-eccentricity probability distributions. a) The upper limit on K corresponds to 10 M_J . b) The upper limit on K corresponds to 200 M_J	69
4.8	Average probability distribution for the period for the 58 long period confirmed planets and candidates.	71
4.9	P_{bf} -stellar mass distribution for the 58 long period confirmed planets and candidates ($P > 300$ days). The error bars correspond to the 1 σ confidence intervals of the associated probability distributions.	72

- 4.10 P_{bf} –stellar metallicity distribution for the 37 long period unconfirmed candidates. The diamonds correspond to candidates with $p_{10/200} < 0.1$ while the circular points correspond to candidates with $p_{10/200} > 0.1$. The error bars correspond to the 1σ confidence intervals of the associated probability distributions. The metallicity is expressed as $[\text{Fe}/\text{H}]$, where 0.0 corresponds to the metallicity of the Sun. The circled points constitute the outer envelope of a possible trend (see text for details). Not plotted here for visual ease are the error bars for the metallicity values: the magnitude of these error bars corresponding to the 1σ confidence interval is 0.019 [76], which is roughly the diameter of the data points on this plot. 74
- 4.11 P_{bf} –stellar metallicity distribution for the 37 unconfirmed candidates and all 257 confirmed planets (for which stellar metallicity information was available) combined. The confirmed planets are represented by dots and the candidates are represented by asterisks. The metallicity is expressed as $[\text{Fe}/\text{H}]$, where 0.0 corresponds to the metallicity of the Sun. 75
- 4.12 Duration of observations versus metallicity for the 37 unconfirmed candidates. The diamonds correspond to candidates with $p_{10/200} < 0.1$ while the circular points correspond to candidates with $p_{10/200} > 0.1$. The circled candidates are the same as the circled envelope candidates in Figure 4.10. The metallicity is expressed as $[\text{Fe}/\text{H}]$, where 0.0 corresponds to the metallicity of the Sun. The magnitude of the error bars corresponding to the 1σ confidence interval for the metallicity values is 0.019 [76], which is roughly the diameter of the data points on this plot. . . 76

4.13 Mean observational uncertainty versus metallicity for the 37 unconfirmed candidates. The diamonds correspond to candidates with $p_{10/200} < 0.1$ while the circular points correspond to candidates with $p_{10/200} > 0.1$. The circled candidates are the same as the circled envelope candidates in Figure 4.10. The metallicity is expressed as $[\text{Fe}/\text{H}]$, where 0.0 corresponds to the metallicity of the Sun. The magnitude of the error bars corresponding to the 1σ confidence interval for the metallicity values is 0.019 [76], which is roughly the diameter of the data points on this plot. 77

Chapter 1

Introduction

In 1992, the first exoplanet was discovered orbiting a pulsar by astronomers Aleksander Wolszczan and Dale Frail [81]. Three years later, the first exoplanet around a solar-type was detected by Michel Mayor and Didier Queloz of the University of Geneva [57]. Since 1995, nearly 300 exoplanets have been discovered at an increasing rate. Over 20 multiple-planet systems are presently known. The most productive search method is the radial velocity technique, but several other methods are also producing results (see section 1.4). A large number of ground- and space-based missions, using a variety of search methods, are currently in operation or being planned.

The exoplanets found until now are extremely diverse, with masses ranging from a few times that of the Earth to several times the mass of Jupiter. They have been revealed to exist around pulsars, solar-type stars and brown dwarfs. Their orbital period can be as short as 1.3 days [83] and as long as 14 years [24]. Their orbital properties in general have baffled astronomers and forced them to reconsider previously proposed planetary formation models based uniquely on the solar system.

The statistical properties of these exoplanet populations can reveal useful information regarding their formation and evolution, as well as the frequency of these companions. In the cases of exoplanets with periods of the order of or longer than the duration of the observations, it is difficult to adequately

characterise their orbits as multiple solutions can fit the data equally well. As such, a good understanding of the uncertainties in orbital parameters will become increasingly important. Bayesian inference can be a convenient and effective way towards this goal. The Bayesian method produces a probability distribution for a given parameter, thus providing a means to estimate the uncertainty associated with the best-fit value of that parameter.

The focus of this thesis on long period exoplanets is motivated by a lack of information regarding their occurrence rate and formation processes. The longer the orbital period of a planet, the farther away from its star the planet is located. It is not clear whether a planet massive enough to be detected from Earth can form in the far outer regions of the protoplanetary disk, where the disk material is less dense and orbits the star slower [52]. These conditions may inhibit the coagulation of dust to form rocky bodies and prevent such bodies from colliding and sticking together, two processes essential to planet formation. If a significant number of exoplanets do exist at such distances from their host star, could they have formed there or were they scattered outwards by another planet in the system? The second possibility may in turn explain the existence of a large number of exoplanets on eccentric orbits. Through a statistical analysis of the radial velocity data associated with 58 long period exoplanets and planetary candidates, this thesis aims to place additional constraints on the period and eccentricity probability distributions of these companions, and attempts to shed more light on the questions described above. Furthermore, direct imaging surveys which search for exoplanets with long orbits, such as the Gemini Deep Planet Survey [49], can use estimates for the constraints on the orbital properties of such planets to better plan the surveys. Our results could therefore serve as input for such surveys.

The outline of the thesis is as follows. In this chapter we give an introduction to extrasolar planets, describe the means by which they are detected and summarise the current knowledge regarding their formation and evolution. In chapter 2, we provide an introduction to Bayesian statistics and review previous work pertaining to the object of this thesis. In chapter 3 we describe our

analysis method and in chapter 4 we present and discuss the results obtained. In chapter 5 we conclude and suggest a few avenues for future work.

1.1 Definitions

To begin, we should define an exoplanet. An exoplanet, or extra-solar planet, is a planet beyond the solar system. According to the International Astronomical Union (IAU), a planet is an object with true mass below the limiting mass for thermonuclear fusion of deuterium that orbits a star; this limiting mass is currently calculated to be 13 Jupiter masses for objects of solar metallicity [11]. Objects above this mass are generally classified as brown dwarfs, but how the boundary between these two classes of objects should be defined is still being debated.

The IAU has also placed a lower mass limit on objects that can be called planets: a planet is a body that has sufficient mass to assume hydrostatic equilibrium (this can be as low as 0.01% of the Earth’s mass) and has cleared the neighborhood around its orbit. The second part of this definition is difficult to verify observationally for exoplanets, but at present this is not a concern because all known exoplanets around solar-type stars have masses significantly higher than the IAU-defined lower limit.

Most exoplanets discovered so far have masses close to or exceeding that of Jupiter (see next section) and orbit very close to their host star (at about 5% of the Earth-Sun distance). As a consequence of the second property, they are believed to be extremely hot. For this reason, they are referred to as “hot Jupiters”.

Finally, we include a note on naming conventions for exoplanets. A lower-case letter is placed after the star name, starting with “b” for the first planet found in the system (for example, 51 Peg b). For the next planet in the system the letter “c” is assigned, and so on. The letter “a” was intended to refer specifically to the host star, but it is rarely used. Letters are assigned based on the order in which the exoplanets are discovered and not on their

position. This is also the standard nomenclature for binary stars or multiple star systems.

1.2 A Note on Size Conventions

Many of the exoplanets discovered thus far are similar in size to Jupiter. Consequently, throughout this thesis we will refer to the radius or the mass of an exoplanet in relation to that of Jupiter (M_J and R_J). For reference, Jupiter has a radius of 7.1492×10^6 m and a mass of 1.8986×10^{27} kg. We also refer to R_\odot and M_\odot , the radius and the mass of the Sun, respectively. These values are 6.9599×10^8 m and 1.9891×10^{30} kg. We sometimes use M_\oplus which is the mass of the Earth and is equal to 5.9736×10^{24} kg. Finally, we measure distances in AU (astronomical unit) which corresponds to the distance between the Sun and the Earth and has a value of 1.496×10^{11} m, or in ly (light-year) which corresponds to the distance light travels in one year and is equal to 9.4607×10^{15} m.

1.3 Celestial Mechanics

Kepler's Laws of Motion revolutionized our understanding of celestial mechanics. The inverse square nature of gravity gives rise to these laws, from which the principles and equations that govern planetary motion around a star have been derived.

Kepler's three laws are:

1. The planets move in ellipses with the Sun at one focus.
2. A line joining a planet and the Sun sweeps out equal areas during equal intervals of time.
3. The square of the orbital period of a planet is proportional to the cube of its semi-major axis:

$$\left(\frac{P}{2\pi}\right)^2 = \frac{a^3}{G(M+m)} \quad (1.1)$$

where P is the period, a is the semi-major axis, G is the gravitational constant, M is the mass of the Sun and m is the mass of the planet.

1.3.1 Model of the Planetary System

An orbit involving two spherical objects and governed by gravitational forces is called a Keplerian orbit. It can be completely described by six parameters (the orbital period, P ; the orbital eccentricity, e ; the inclination relative to the plane of the sky, i ; the argument of periastron measured from the plane of the sky, ω ; the longitude of ascending node, Ω ; and the mean anomaly, M). Some of those parameters are shown in figure 1.1¹. M is not shown in the figure because it varies with time; it is the angle (as seen from the host star) between the periastron of the orbit (the point where the planet is closest to its star) and the position at which the planet would be if it was travelling at constant velocity along a *circular* orbit. Another way to describe it is as the measure of the phase of the planet's motion around the orbit.

Using Newton's Laws of Motion and the fact that a Keplerian orbit is circular in velocity space, it is possible to obtain the velocity equation for an object moving in a Keplerian orbit [33], which depends on five of the six parameters listed above (all except Ω), on K (the velocity semi-amplitude) and on C (the velocity constant, equivalent to the mean line-of-sight velocity of the center of motion of the planetary system):

$$f = K(\cos(\omega + \theta) + e \cos \omega) + C \quad (1.2)$$

where θ is the true anomaly and depends on P , e , the orbital phase ϕ and time as follows

¹The image in figure 1.1 comes from the following web address: <http://scienceworld.wolfram.com/physics/OrbitalElements.html>

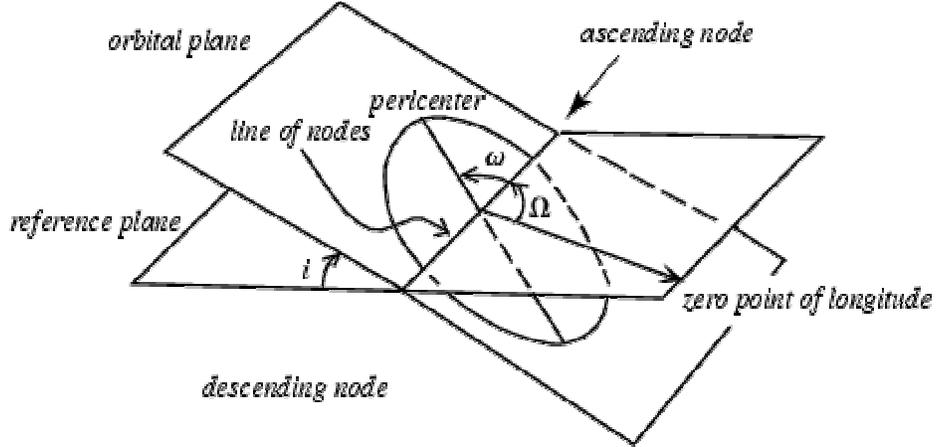


Figure 1.1 Keplerian orbital parameters (i , ω and Ω). In the context of exoplanet orbits, the reference plane is the plane of the sky. Image is courtesy of *Eric Weisstein's World of Physics*.

$$\tan \frac{\theta}{2} = \sqrt{\frac{1+e}{1-e}} \tan \frac{E}{2}. \quad (1.3)$$

E is called the eccentric anomaly and it is determined by solving Kepler's equation

$$E - e \sin E = M. \quad (1.4)$$

M is the mean anomaly

$$M = \frac{2\pi t}{P} + \phi, \quad (1.5)$$

where t is time.

By setting e and ω to 0, we obtain the velocity equation for an object moving in a circular Keplerian orbit

$$f = K \sin \left(\frac{2\pi t}{P} + \phi \right) + C. \quad (1.6)$$

We chose f for the velocity “predicted” by the model in order to distinguish it from the measured velocity v , which will appear in Chapter 3.

For multiple-planet systems, a multi-Keplerian orbit is fitted to the data. It can be expressed as follows:

$$f_N = C + \sum_{j=1}^N (K_j (\cos(\omega_j + \theta_j) + e_j \cos \omega_j)). \quad (1.7)$$

However, for the purpose of this thesis only single-Keplerian orbits have been considered. Furthermore, the six independent parameters we will refer to from here on are K , P , e , ω , ϕ and C . (M , θ and E are dependent on one or more of these six parameters.)

K can also be related to the planet mass, M_P , by [19]

$$K = \frac{28.4 \text{ m/s}}{\sqrt{1-e^2}} \left(\frac{M_P \sin i}{M_J} \right) \left(\frac{P}{1 \text{ yr}} \right)^{-1/3} \left(\frac{M_\star}{M_\odot} \right)^{-2/3}, \quad (1.8)$$

where M_J is the mass of Jupiter, M_\star is the mass of the host star and M_\odot is the mass of the Sun.

As described in section 1.4, the inclination (i) cannot be obtained from radial velocity measurements alone; the longitude of the ascending node (Ω) and the mass of the planet (M_P) cannot be determined either. Only $m_P \sin i$, which corresponds to the minimum mass of the planet, can be calculated. The stellar mass can either be obtained from stellar properties catalogs ([73], [76]), or by using the $B - V$ stellar mass relation [1].

1.4 Exoplanet Detection Methods

Currently, several techniques can be used for the detection of exoplanets. In this section we briefly describe the most commonly used of these techniques.

The most productive exoplanet detection method so far is the radial velocity (RV) method, which was used to obtain the data analysed in this thesis.

About 200 exoplanets have been discovered with this method. A planet in orbit around a star induces an orbital motion in the latter. This motion translates into redshifted or blueshifted lines in the star’s spectra, from which RV values can be obtained. As an example, Jupiter induces a velocity variation in the Sun’s motion of 12 m/s. Such variations are now relatively easy to detect in nearby stars with high-precision spectrometers, such as HIRES (High Resolution Echelle Spectrometer) [12] at the Keck optical telescope in Hawaii and the HARPS (High Accuracy Radial Velocity Planet Searcher) [62] spectrometer at the La Silla Observatory. Indeed, these instruments can pick up signals with measurement errors down to 1 m/s [14].

A major limitation is the signal-to-noise ratio (SNR), so such high precision can only be achieved for nearby stars (within a few hundred light-years) [15]. Furthermore, as described in more detail in the next section, only the minimum mass of an exoplanet can be obtained from RV measurements. RV measurements do not provide any information on the inclination (i) of the orbit, so only the component of the RV which lies along our line of sight ($v \sin i$) can be measured (see figure 1.2). Since the true mass (m) depends on the true velocity of the star, without the inclination only $m \sin i$ can be estimated using this detection method.

The second most efficient exoplanet detection technique is transit photometry, which has been used to observe over 30 exoplanets. When the orbital plane of an exoplanet lies within a few degrees of the observer’s line of sight (about 10% of all known exoplanets have such alignment), a dip in the host star’s luminosity can be detected as the planet passes in front of it if the latter has a small enough orbit. The radius of the planet and the inclination of the orbital plane can then be calculated from the duration and depth of this drop in luminosity. The transit method often gives rise to false detections, because several phenomena not related to a planet can also cause variations in its luminosity. These include grazing binaries, small-radius stellar companions and stellar activity patterns [59]. Therefore, RV follow-ups are typically carried out in order to confirm the presence of a planet. These follow-ups are also

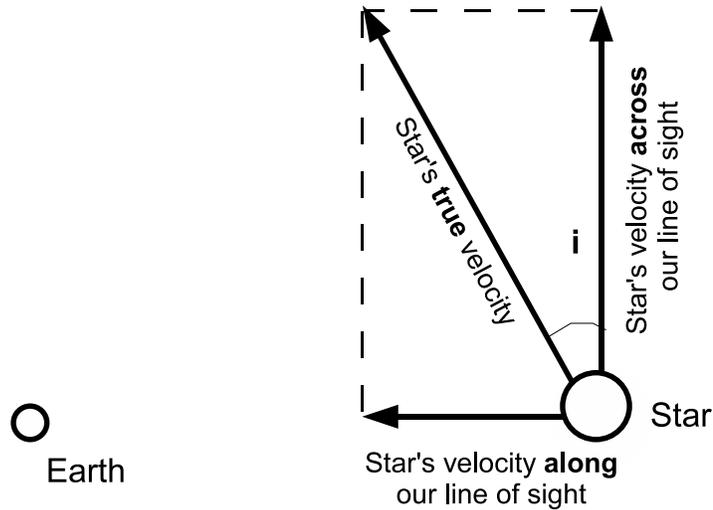


Figure 1.2 Velocity components of a star.

used in conjunction with the transit measurements to determine both the true mass (since i can be deduced from the transit data) and the radius of the planet, and thus its density. The transit method also allows the study of an exoplanet's atmosphere. It is possible to detect elements present in the upper atmosphere of the planet by closely examining light from the host star that passed through it ([67], [38]).

Furthermore, by subtracting the luminosity of the star when the planet is behind it (giving rise to a secondary eclipse) from the luminosity observed just before and after the secondary eclipse, the radiation arising from the planet alone can be obtained. This information can in turn be used to estimate the planet's temperature through the Stefan-Boltzmann law. Charbonneau et al. [16] and Deming et al. [21] have carried out this type of measurements on the exoplanets TrES-1 and HD 209458 b, respectively.

The first three planets outside our solar system were discovered by pulsar timing around PSR 1257+12 [81]. The rotation of a pulsar is extremely regular, so slight anomalies in the timing due to its motion in an orbit stand out

and can indicate the presence of planets. This method is so sensitive that it can detect planets with a mass as low as 0.04% of that of the Earth. Although another exoplanet has been found in orbit around PSR B1620-26 [75], pulsars themselves are quite rare and hence it is not likely that many planets will be found in this way. However, in 2007 the timing method has also been used to discover a planet around V 391 Pegasi, a B-type subdwarf star which exhibits pulsations with a very stable frequency [70].

Gravitational microlensing occurs when the light of a distant star is magnified by the gravitational field of a foreground star acting as a lens. A lensing event requires the two stars to be almost precisely aligned with each other and is usually brief (of the order of a few days to a few weeks). If the foreground star hosts a planet, the latter can make a detectable contribution to the lensing effect. Such events are rare, but microlensing searches such as OGLE (Optical Gravitational Lensing Experiment) monitoring over 100 million stars have already achieved six exoplanet detections using this method, the first of which was announced in 2004 [8]. Most of these stars lie between the Earth and the Galactic Bulge which provides a large number of background stars. The fact that a given lensing measurement cannot be repeated represents another disadvantage of this technique. In addition, follow-ups with other methods are usually impossible because the lensing stars are most often located several thousand light-years away.

There are two other methods currently employed in exoplanet searches and studies which are rising in popularity due to some recent successes. Astrometry involves measuring the change in position of a star due to a planetary companion. This method is expected to work best for systems with low-inclination orbital planes. No planetary companions have yet been discovered in this way, but high-precision astrometric observations have helped determine the mass of a few already known exoplanets ([4], [5], [58]), and have shown that HD 33636 b, which was previously classified as a planet [63], is in fact a low-mass star

[3]. Although exoplanets are generally extremely faint relative to their host stars, the VLT (Very Large Telescope), Gemini and Subaru telescopes are in the process of being equipped with instruments capable of directly imaging exoplanets. Even so, in 2004 the first image of an exoplanet around the brown dwarf 2M1207 has been obtained [18].

1.5 Observed Properties of Exoplanets

Because the amount of information that can be obtained for a single exoplanet is in most cases limited to its orbital properties and to a certain extent its physical properties, we can learn more from studying these properties statistically for all planets than for each planet in particular. In this section we briefly review the statistical distributions for some of these properties and the conclusions that can be drawn from them.

Figures 1.3 to 1.6 show properties pertaining to a fraction of or all of 292 planets known to date. The data used for the figures was obtained from the *Interactive Extra-solar Planets Catalog*².

Figure 1.3 shows a considerable decrease in the number of planets at high masses. The low end of the distribution is affected by a selection effect common to radial velocity surveys: low-mass planets induce small velocity variations in the motion of their host star; they are thus more difficult to detect and are under-represented in this distribution. However, massive planets are easier to detect so the scarcity of planets with minimum masses larger than $12 M_J$ is real. This also confirms the existence of the brown dwarf desert, a region in the mass function of stellar companions between the planetary and the brown dwarf mass regimes where very few objects have been found; the brown dwarf desert indicates that there is a fairly clear distinction between planets and brown dwarfs.

Figure 1.4 shows the distribution of periods for nearby hot Jupiters. There

²The *Interactive Extra-solar Planets Catalog* can be found at <http://exoplanet.eu/catalog.php>

is a “pile-up” of such planets at periods of around 3 days, suggesting that the migration mechanism which brought them close to their host stars ceases once they reach this period. Another explanation is the existence of a braking mechanism which stops the giant planets at this period, or loses strength at periods shorter than 3 days and sends the planets into the star. It should be noted that there is no significant selection effect contributing to this pile-up since radial velocity surveys have the same sensitivity to hot Jupiters for all the periods in figure 1.4.

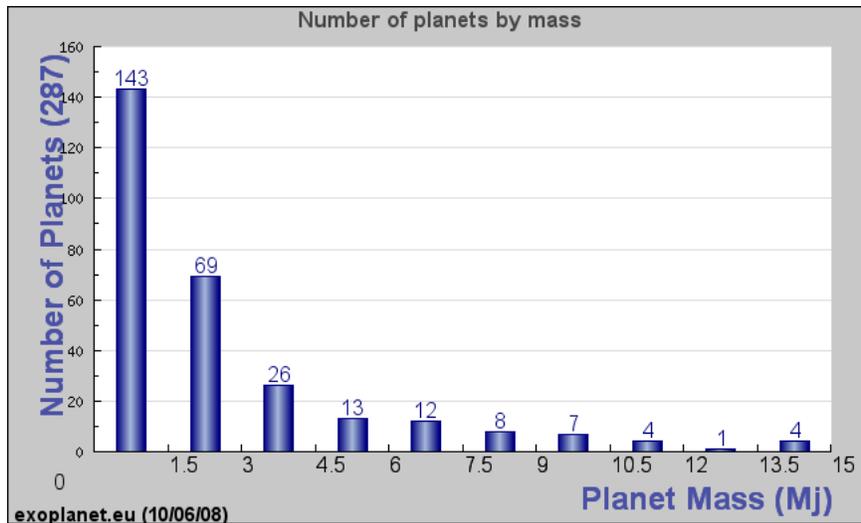


Figure 1.3 Minimum mass distribution for the 292 known nearby exoplanets with $m \sin i < 15M_J$. We produced this figure interactively at *exoplanet.eu*.

There seems to be a lack of close-in planets with high masses, as indicated in figure 1.5. However, selection effects make low-mass planets beyond 1 AU difficult to detect because such planets induce small velocity variations in their host stars and require longer observation times to cover a full period, so it is not clear whether the mass distribution for such planets is different from that of close-in planets.

Figure 1.6 shows orbital eccentricity as a function of minimum mass for nearby exoplanets. Many of the planets at $e \approx 0$ are located very close to their host star ($a < 0.1$ AU) and their orbits are likely to have been tidally circularised. No strong correlation between eccentricity and mass is apparent,

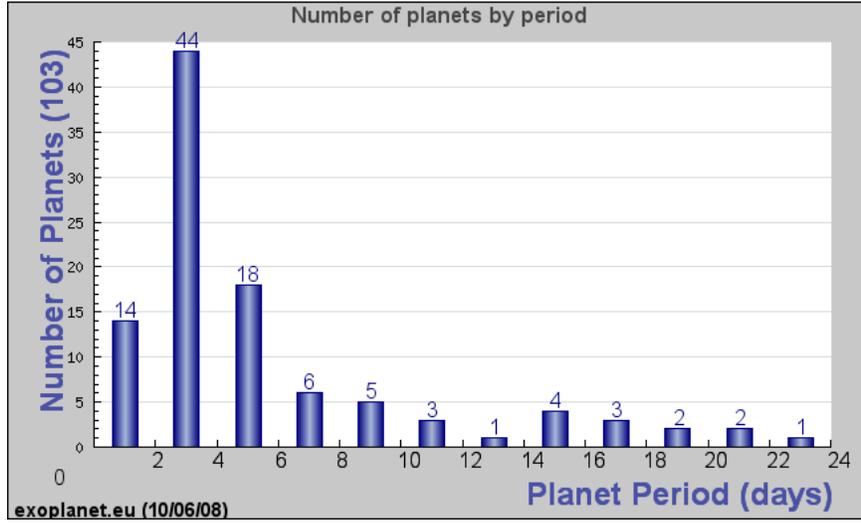


Figure 1.4 Distribution of orbital periods for the 103 known nearby hot Jupiters. We produced this figure interactively at *exoplanet.eu*.

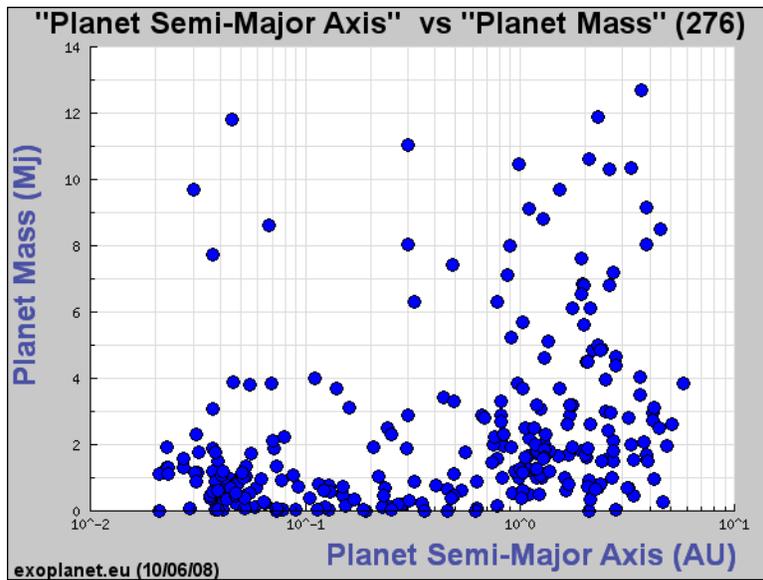


Figure 1.5 Minimum mass distribution of the 276 known nearby exoplanets with a < 6.5 AU. We produced this figure interactively at *exoplanet.eu*.

but close inspection shows that high-mass exoplanets ($m \sin i > 5M_J$) have a higher median eccentricity than lower-mass exoplanets. Radial velocity surveys have uniform sensitivity for eccentricities below 0.7 [19].

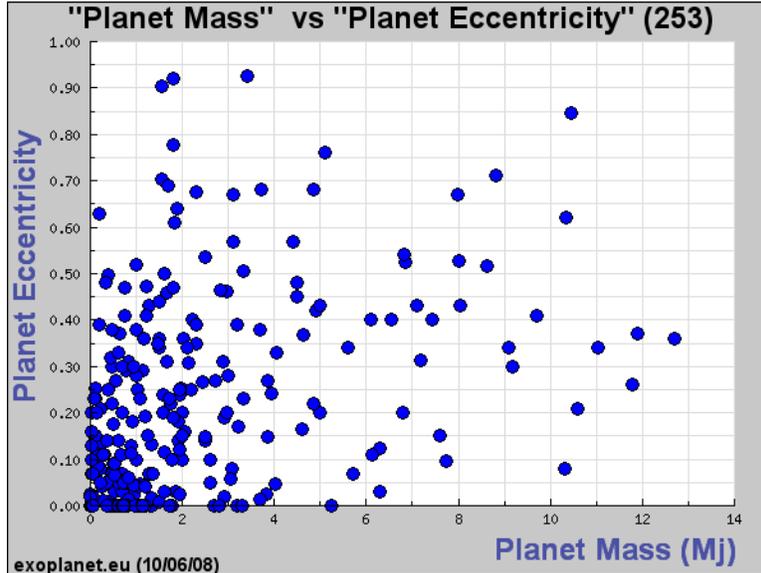


Figure 1.6 Distribution of orbital eccentricities as a function of minimum mass for the 253 known nearby exoplanets with $m \sin i < 13M_J$ for which the eccentricity is known. We produced this figure interactively at *exoplanet.eu*.

In addition, Johnson et al. (2007) [45] have found that the fraction of stars with planets increases with the stellar mass. However, since low-mass ($< 0.7M_\odot$) and high-mass ($> 1.3M_\odot$) stars have only recently started to be searched for planets, conclusions with respect to this trend should be drawn with caution. Fischer & Valenti (2005) [25] have also shown a correlation between the fraction of stars with planets and the stellar metallicity. They found that planet frequency rises rapidly with stellar metallicity according to a power law in terms of the fraction of metal atoms: $p(\text{planet}) = 0.03[(N_{Fe}/N_H)/(N_{Fe}/N_H)_\odot]^{2.0}$.

Furthermore, radial velocity observations to date have also shown that multiple planet systems are common, with over 25 found to date. Transit searches have also produced some interesting results, such as revealing a population of planets with extremely short periods (between 1.2 and 3 days). In addition, the radii obtained for transiting (and hence close-in) exoplanets confirm their gas giant nature.

1.6 Planet Formation and Evolution

1.6.1 Formation

The complete planet formation process is not known with certainty. The currently accepted theory is that as a star forms, ever-larger bodies form in the surrounding protoplanetary disk through accretion of dust particles. Aerodynamic forces and the gravity from the star are dominant until these bodies attain sizes of 1-100 km, becoming planetesimals. However, the growth mechanism of centimeter and meter sized rocky objects to planetesimal size is still unknown [22]. These rocky objects have high velocities and can migrate towards the star very quickly (hundreds of years), due to a “headwind” from the slower rotating gas [80]. This implies a rapid growth to kilometer sized objects to avoid being absorbed into the star. One possibility is that planetesimals form via continuous collisions, possibly aided by the concentration of particles into small regions due to local turbulence, but these meter scale boulders are expected to stick together poorly [6]. Alternatively, the formation of planetesimals may take place through gravitational collapse of regions with high concentrations of solid particles [82], but turbulence will prevent sedimentation of solids to a sufficiently high density, and we know turbulence is required to explain gas accretion in protostellar disks [40]. However, Johansen et al. (2007) may have found a way to overcome this challenge [44].

Planetesimals are massive enough to continue accreting at an accelerated rate and their increasingly spherical shape means they are only weakly affected by aerodynamic forces. Planetesimals can collide and form larger bodies or undergo close encounters with one another, which alter their orbits. Larger bodies typically grow more rapidly than smaller ones at the expense of the latter, as part of the runaway and oligarchic growth processes ([48], [47]). At this stage, the newly formed large objects are known as protoplanets. Further collisions and radioactive decay heat the protoplanets until they melt at least partially, causing dense elements such as iron to sink to the center and form a core surrounded by a rocky mantle [42]. The evolution from planetesimal

to protoplanet is expected to last one million years or less. The resulting population is made up of lunar-sized protoplanets which continue to collide and then accumulate atmospheres, which greatly increase the capture rate of remaining planetesimals [43], thus forming terrestrial planets on a time scale of up to 100 million years. Planetesimals that have avoided collision can become natural satellites or asteroids.

There are two models for gas giant formation. The standard theory is core accretion [64], whereby a core of the order of 10 Earth masses forms through a process similar to the formation of terrestrial planets, and subsequently accretes gas from the protoplanetary disk. Gas giants form beyond the snow line, at a few AU from the host star, since the inner region of the disk does not contain enough solid material to form a massive enough core. The snow line defines the radius from the star beyond which water can condense. Initially, the core is surrounded by a low mass atmosphere which grows steadily more massive as the gas cools and contracts onto the core. As the latter exceeds the critical core mass, beyond which a hydrostatic envelope can no longer be maintained, gas begins to flow onto the core at an increasing rate as the planet gains mass. The time scale for gas giant formation through this model is difficult to estimate, but it is clear that the planet must form before the star ignites and the disk evaporates, terminating the supply of gas.

The alternative theory is that of gravitational disk instability [23]. A sufficiently massive and/or cold gas disk becomes gravitationally unstable. If this instability can lead to fragmentation, formation of massive planets can occur. An important advantage of this mechanism is that it works extremely rapidly.

Nevertheless, core accretion is the preferred model. One reason for this is even though protoplanetary disks may be massive enough to be unstable, calculations show that they are unlikely to cool fast enough to fragment (except perhaps at large radii) [7]. Another reason is that the core accretion model naturally supports the existence of ice giants like Neptune and Uranus, although such planets have not yet been confirmed outside our solar system.

Finally, the observed correlation between the frequency of exoplanets and the metallicity of their host stars [25] can be explained as a consequence of core accretion: if the disk is rich in solids, a critical mass core can form more easily. It is not clear whether this correlation can be explained by the gravitational instability model. Additional observations may either validate one of these models, or suggest that a combination of the two is most appropriate.

1.6.2 Migration

Interest in planetary migration mechanisms increased with the discovery of 51 Peg b [57], an exoplanet with an orbital period of 4.2 days, which places it extremely close to its host star. The planet could not have formed at that distance since protoplanetary disk material so close to the star would have been absorbed before the planet could have been created. Two gas disk and one planetesimal-driven migration mechanisms have been proposed.

Low-mass planets undergo type I migration [79], where the surface density profile of the gas disk is only mildly affected by the planet. The migration rate is proportional to the planet's mass, and the planet remains completely embedded in the gas. Interaction with the gas disk interior to its orbit adds angular momentum to the planet, while interaction with the exterior disk removes angular momentum. Consequently, the planet can either migrate inward or outward. Simulations suggest that it will almost always migrate inward, potentially on a time scale as short as one million years [74].

Massive planets perturb the gas disk much more strongly. During type II migration [79], the exchange of angular momentum between the planet and the disk repels gas from the vicinity of the former, creating a surrounding gap with low gas surface density. The direction and speed of the migration depends on how quickly the gas disk can flow back toward the gap. The motion of the planet is locked to the viscous evolution of the disk. If the gas is flowing inward, the planet also moves inward, while if the gas flows outward, the planet also moves outward. Type II migration is typically slower than Type I migration.

Planets can also exchange angular momentum with smaller bodies through

collisions or close encounters. During such an event, if a planet ejects a planetesimal from the planetary system, it must give up energy and thus move inward towards the star. A planet that scatters planetesimals into shorter period orbits gains energy and migrates outward. Ida et al. (2000) [41] have shown through numerical simulations that planetesimal-driven migration can indeed be significant.

So far, there is no direct observational evidence for gas disk or planetesimal migration, but type I and type II migration are widely accepted as explanations for the existence of hot Jupiters. Furthermore, no satisfactory braking mechanism has yet been developed to explain how the migration of gas giants is stopped.

1.6.3 Orbit Alterations

Planet-planet scattering can occur after the protoplanetary disk has been depleted or evaporated. An initially unstable planetary system could evolve via ejection of one or more of the lighter planets (arising from collisions or close encounters) or an increase in the orbital separation of the planets toward a more stable configuration (through migration for example). Numerical calculations have shown that planet-planet scattering can reproduce the observed eccentricity distribution of massive exoplanets, which makes this mechanism a leading candidate for explaining why a significant fraction of exoplanets have non-circular orbits [26].

Planets that form on or acquire eccentric orbits can undergo tidal circularization to reach circular, short-period orbits observed in many extrasolar systems [66]. The periastron of a planet on an eccentric orbit is in the neighborhood of the star, where tidal effects become important. The gravitational potential gradient of the star creates a differential attraction between the inner side (facing the star) and the outer side of the planet. As a result, a tire-shaped structure forms at the surface of the planet, which gives rise to the energy dissipation by internal friction within the planet that tends to circularize the orbit. A side effect of this process is that the rotation and revolution

of the planet will become synchronized; consequently, the planet will become tidally locked, always showing the same side to the star. This circularization process takes about one billion years for the planet to attain a final period of a few days [39].

Chapter 2

Bayesian Inference at Work in Exoplanet Orbit Analysis

2.1 Introduction to Bayesian Inference

In this section we present a brief overview of Bayesian inference. For a more complete resource, see Sivia (1996) [71].

Bayesian inference interprets the concept of probability as a measure of belief in or a state of knowledge concerning a given hypothesis. As evidence accumulates, the degree of belief in a hypothesis should either increase or decrease. Bayesian inference uses a numerical estimate of the degree of belief in a hypothesis *before* evidence has been observed (the prior probability) and calculates a numerical estimate of the degree of belief in the hypothesis after evidence has been observed (the posterior probability), sometimes repeatedly as new evidence becomes available.

The essential equations of Bayesian inference can be derived from the basic algebra of probability theory. Starting with the joint probability distribution $p(x, y)$ (where x and y can be scalars or arrays of several variables), we can produce the marginalized probability distribution for x by integrating over y :

$$p(x) = \int p(x, y) dy. \quad (2.1)$$

Using the fact that we can write the joint probability distribution as a product of equation (2.1) and a conditional probability distribution as follows

$$p(x, y) = p(x)p(y|x), \quad (2.2)$$

Bayes' theorem can be expressed as

$$p(x|y) = \frac{p(x, y)}{p(y)} = \frac{p(y)p(x|y)}{\int p(y)p(x|y)dy}. \quad (2.3)$$

If we now replace x with a set of observational data d , and y with a set of model parameters H (the orbital parameters in the context of this thesis), equation (2.3) becomes

$$p(H|d, \mathcal{M}) = \frac{p(d, H|\mathcal{M})}{p(d|\mathcal{M})} = \frac{p(H|\mathcal{M})p(d|H, \mathcal{M})}{\int p(H|\mathcal{M})p(d|H, \mathcal{M})dH} \quad (2.4)$$

where \mathcal{M} is the model (equation (1.2) in this thesis) which includes the meaning of the model parameters H and their relationship to the data d . In equation (2.4) $p(H|\mathcal{M})$ represents the prior probability distribution that is based on previous knowledge of the model parameters; $p(d|H, \mathcal{M})$ is the probability of the data given the model parameters (also known as the likelihood function); and $p(H|d, \mathcal{M})$ is the posterior probability distribution.

The power of Bayes' theorem lies in the fact that it relates the quantity of interest, the probability that a given set of parameters is true given the data and the model ($p(H|d, \mathcal{M})$), to the term that we have a better chance of being able to assign, the probability that we would have observed the measured data if a given set of parameters and the model were true ($p(d|H, \mathcal{M})$).

For a Keplerian orbit there are six parameters to consider. If we want to obtain the probability distribution for a single one of them (for example K), we can integrate equation (2.4) over the remaining parameters as follows:

$$p(K) = p(H(K)) = \int p(H(K, P, e, \omega, \phi, C))dPded\omega d\phi dC. \quad (2.5)$$

In the same way, we can obtain probability distributions for two or more parameters as well.

Bayesian inference has several advantages over frequentist statistical methods. First, since all Bayesian inferences are based on the posterior probability distributions, there is a rigorous basis for quantifying uncertainties in the parameters. Frequentist methods generally result in point estimates for these parameters so in order to estimate the associated uncertainties, they are often combined with techniques such as bootstrap which rely on fictitious observations simulated by the experimenter. The rigorousness of Bayesian methods arises from the fact that posterior distributions depend only on the observations that actually took place. Another advantage consists of the ability to incorporate a variety of information and observations in the prior and the likelihood function, thus updating the posterior distributions as new knowledge or data becomes available. Furthermore, the Bayesian framework provides a natural basis for making predictions about future observations [29].

The necessity of specifying a prior probability distribution before performing any Bayesian analysis is the main inconvenience of Bayesian inference. It is important to choose a prior which reflects the state of knowledge (or ignorance) of the data as accurately as possible; when in doubt, one should select the more conservative prior. Additionally, one should always check the dependence of the results and conclusions on the choice of prior.

A major difficulty in Bayesian inference is the computation of the lower integral in equation (2.4), especially when H has a large number of dimensions. We chose simple, direct numerical integration for this thesis, partly because we were not concerned with multiple planetary systems (which would have dramatically increased the number of parameters on which H depends), but especially because we wanted a simple and low maintenance program which would allow us to focus on the features of the probability distributions and possible correlations between parameters. However, when the main goal is to search for exoplanets and characterise their orbits using a large number of data sets, more efficient methods become necessary such as Markov Chain Monte

Carlo ([27], [28]) and parallel tempering ([34], [35], [36]) algorithms. Work with these methods also encompasses most of the Bayesian statistics applications in the field of exoplanets so far.

2.2 Markov Chain Monte Carlo

During the last decade, the Markov Chain Monte Carlo (MCMC) method has been applied successfully to several astronomical data and problems, including the analysis of UV and X-ray spectra [46], star formation history [61], cluster weak lensing and the Sunyaev-Zeldovich effect [56] and the cosmic microwave background [77], to name a few.

Ford has developed a MCMC algorithm based on Bayesian inference to characterise the uncertainties in the orbits of exoplanets. As described in Ford (2005) [27], the goal of the MCMC method is to generate a sequence of sets of parameter values that are sampled from a desired probability distribution, which can be calculated from an initial set of parameter values (H_0) and a transition probability ($p(H_{n+1}|H_n)$). The Monte Carlo aspect of MCMC refers to the randomness in the generation of each set of parameters (state). The Markov property indicates that the probability distribution for H_{n+1} can depend on H_n but not on previous states. The Markov chain will converge to the posterior probability distribution if it is aperiodic, irreducible and reversible [31]. The chain being irreducible means that it should be able to reach any state with nonzero probability from any other state with nonzero probability.

The Metropolis-Hastings (M-H) algorithm satisfies these conditions. It can draw samples from the probability distribution $p(H)$ and generates a Markov chain in which each state H_{n+1} depends only on H_n . The algorithm uses a proposal distribution $q(H'|H, \mathcal{M})$ to generate a new proposed sample H' . The proposal is accepted as the next value if the acceptance probability (α), drawn from $U(0,1)$ satisfies

$$\alpha < \frac{p(d|H', \mathcal{M})q(H|H', \mathcal{M})}{p(d|H, \mathcal{M})q(H'|H, \mathcal{M})}. \quad (2.6)$$

If the proposal is not accepted, the current value of H is conserved: $H_{n+1} = H_n$.

The M-H algorithm does not specify when the chain will converge. The choice of $q(H'_\mu|H_\mu, \mathcal{M})$ is very important as it can have an extremely significant effect on the rate of convergence. Based on tests which involved monitoring the fraction of states that are accepted, Ford opted for a Gaussian proposal distribution,

$$q(H'_\mu|H_\mu, \mathcal{M}) = \frac{1}{\sqrt{2\pi}\psi_\mu} \exp\left[-\frac{(H'_\mu - H_\mu)^2}{w\psi_\mu^2}\right], \quad (2.7)$$

where each ψ_μ is a parameter that controls the scale for the steps taken [27]. The values of ψ_μ are automatically selected by the MCMC algorithm, which saves a lot of time that would otherwise be spent on trial runs in order to manually determine the most appropriate set of ψ_μ values. In short, MCMC explores probability space by taking short steps in regions of high probability, and longer steps in regions of low probability.

In this paper [27], Ford also showed that Bayesian analyses work particularly well for systems where the orbital period is comparable to the duration of observations. Ford et al. (2005) have used this algorithm to analyse the dynamical interactions of the three planets in the ν Andromedae system [30].

In Ford (2006) [28], he introduced several improvements to the MCMC method. He tested several alternative proposal probability distributions which improve the efficiency of MCMC. For example, it was shown that the rate of convergence can be increased significantly when specific proposal distributions are used for planets with a large range of possible orbital periods, high eccentricity or low eccentricity. The optimised algorithms were applied to several exoplanet systems. In one case (HD 117207) it was demonstrated that the uncertainties in the orbital parameters are larger than estimated using frequentist resampling methods. In another (HD 37124), a good choice of proposal distributions allowed the MCMC algorithm to identify small peaks in strongly multimodal (comprising more than one significant solution) posterior probability distributions.

Although the MCMC method was also used on a few 2- and 3- planet systems [28], improvements are needed in order to efficiently simultaneously consider models with zero, one, two or more planets. One possible alternative is parallel tempering [35].

2.3 Parallel Tempering

Gregory (2005) [34] used an approach similar to that of Ford (2005) [27] to implement the MCMC method, involving a M-H algorithm. In addition, he introduced parallel tempering to better handle multimodal probability distributions, which most often arise for multiple planet systems and for planets with periods longer than or comparable to the duration of observations [28].

In parallel tempering, several copies of a MCMC simulation are run in parallel. They each have a different temperature T . (Here, temperature refers to the step size in the Monte Carlo algorithm: high temperature corresponds to large steps, and low temperature corresponds to small steps.) Instead of varying T from 0 to infinity, $\beta = 1/T$ is used, which varies between 0 and 1. The desired posterior probability distribution corresponds to the simulation with $\beta = 1$. The remaining simulations have increasingly higher T (or increasingly lower β). At a given MCMC step, a pair of adjacent simulations is randomly chosen and a proposal is made to swap their parameter configurations. The proposal is accepted or rejected based on a M-H criterion. If the swap takes place, it allows for an exchange of information across the ensemble of simulations. Simulations with a higher temperature can experience radically different parameter configurations after a swap, while lower temperature simulations can refine their configurations. Adjacent simulations must overlap somewhat in order for a swap to take place. In short, by a careful choice of temperatures and number of simulations, it is possible to improve the mixing properties of a set of Monte Carlo simulations.

Using parallel tempering with MCMC, Gregory found two alternative orbital solutions for the HD 73526 system [34]. In addition, he claimed that the

data for HD 208487 system shows a 95% probability for a second planet [35], and the data for HD 11964 indicates the presence of three planets [36].

2.4 Non-Bayesian Statistical Methods

Several non-Bayesian methods are also used to find best-fit orbital solutions for exoplanets. In this section we very briefly describe the main techniques.

A periodogram is essentially a plot of the minimum χ^2 versus the period, in which the velocity semi-amplitude and the phase can vary. Because it is comparable to fitting data with a simple sinusoid, it is not an accurate method for estimating the parameters of eccentric orbits. However, periodograms allow for quick identification of potential orbital periods in the observational data without requiring a simultaneous fit for all the other parameters, which can then be used as input (initial guesses) for other algorithms.

The Levenberg-Marquardt algorithm [65] searches the parameter space to minimise χ^2 . The disadvantage of this technique is the fact that it can only find one minimum for a given initial guess of parameters. It is most useful for refining the model parameters locally once the rough topography of χ^2 space is obtained through global search algorithms.

Simulated annealing [65] is one such algorithm. Initially it takes large steps (high temperature) and then gradually smaller steps (lower temperature). Two disadvantages are the large number of iterations and the need to specify a cooling curve (the rate at which the steps should be reduced).

A genetic algorithm is another kind of global search algorithm, which is generally based on biological evolution [17]. However, it can be extremely time-consuming, especially when applied to a large number of data sets. It has been used successfully on the *v* Andromedae [72] and the GJ 876 [51] systems.

2.5 Motivation for this Thesis Work

In 2004, Cumming statistically analysed the sensitivity of radial velocity surveys at large periods and low velocity semi-amplitudes [19]. In order to estimate the significance of a Keplerian fit to radial velocity data, he calculates the false alarm probability (FAP) associated with a particular fit using the Lomb-Scargle (LS) periodogram ([53], [69]). The periodogram uses power (z) as a measure of the goodness of the fit. Consider χ_{Kep}^2 , the χ^2 of a Keplerian fit, and χ_{mean}^2 , the χ^2 from the fit of a constant to the data. The power is

$$z = \frac{(\chi_{mean}^2 - \chi_{Kep}^2)/4}{\chi_{Kep}^2/\nu} \quad (2.8)$$

where ν is the number of degrees of freedom. A simple estimate of the FAP is

$$FAP \approx Mprob(z > z_0) \quad (2.9)$$

where M is the number of independent frequencies and $prob(z > z_0)$ is the probability that z is larger than z_0 (z_0 is obtained using fake data sets which contained only noise). For details on how $prob(z > z_0)$ is obtained and for a discussion of a Bayesian approach to the LS periodogram, see [19]. Among the other results of this paper are analytical expressions for the velocity thresholds for short and long periods. Of particular interest for this thesis is that for circular orbits at long periods, Cumming showed that a linear or quadratic relationship ($K \propto P$ or $K \propto P^2$) can fit the data as well as a Keplerian orbit (see section 4.3). He also showed that there are important selection effects to be considered for observations of orbits with $e \gtrsim 0.6$.

A sample of 585 stars were searched for planets using the false alarm probability by Cumming et al. (2008) [20]. They showed that all exoplanets with $P > 2000$ days, $K > 20$ m/s and $e \lesssim 0.6$ had already been announced, and they found evidence for 76 unconfirmed candidates. Of these, 27 candidates correspond to companions with masses $\gtrsim 20M_J$ (thus likely not planets). The remaining candidates have not yet been confirmed due to either low K ($\lesssim 20$

m/s) or high P ($\gtrsim 2000$ days), which make their orbital parameters difficult to constrain. Figure 2.1 summarises these results. A more in-depth characterisation of the orbits of the candidates to the right of the vertical line and below the $10M_J$ line in figure 2.1, including uncertainty estimates of their orbital parameters, constitutes a major motivation for this thesis.

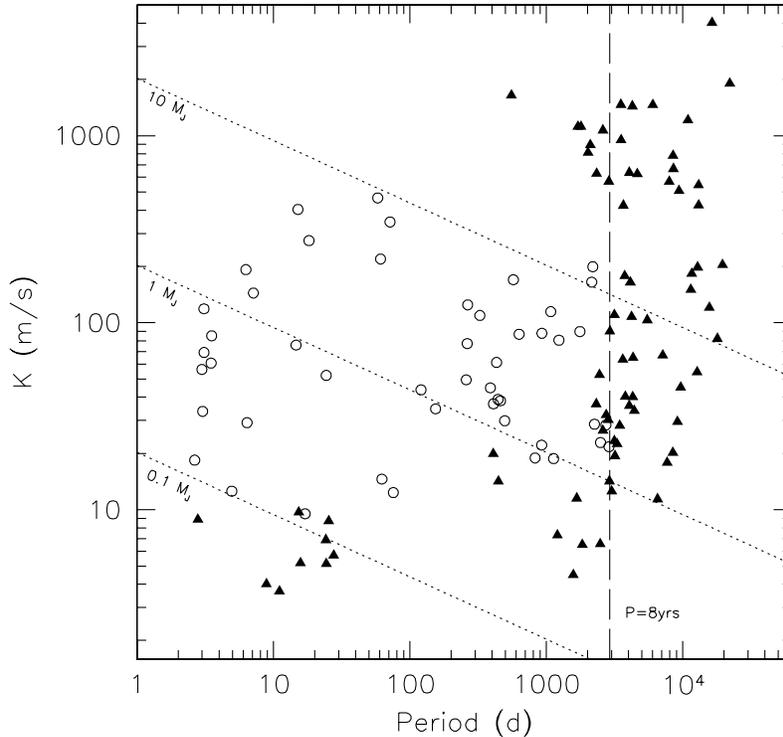


Figure 2.1 All confirmed planets (48) and unconfirmed candidates (76) flagged among the 585 stars analysed. The open circles represent the planets and the solid triangles represent the candidates. The vertical dashed line corresponds to the duration of the survey (8 years). The dotted lines correspond to exoplanet masses of $m \sin i = 0.1, 1$ and $10M_J$, for a solar mass star. This figure is used with the permission of A. Cumming.

Cumming et al. (2008) also found that 17–20% of stars should have planets with periods shorter than 90 years (3.2×10^4 days), but data show that 18% of stars host planets with periods shorter than 32 years (1.2×10^4 days). This may indicate that there is a lack of planets at periods larger than 32 years, or that there are more stars with planets than expected based on the assumptions

made in this paper. We compare our results with these expectations in section 4.4.

Finally, in Cumming et al. (2008) the eccentricity (especially if $e \gtrsim 0.6$) and associated selection effects were neglected in some parts of their analysis. This thesis aims to study the eccentricity distributions and consider the prospective results when placing constraints on e and P .

Chapter 3

Method

In this chapter we describe our method for calculating probability distributions for the exoplanet orbital parameters. We chose to base our approach on Bayesian inference because of its capacity to quantify uncertainty in the orbits of exoplanets and because it can easily incorporate new information in the form of priors or additional data. We used the IDL programming language to perform the necessary calculations.

3.1 Circular Orbits

It is simpler to constrain the parameters for circular orbits than for eccentric orbits. Therefore, developing a program which only characterizes orbits using circular orbit models as a first step allowed us to arrive at the most time efficient and straightforward set of computation and integration techniques.

The model used for fitting radial velocity data with circular orbits is

$$f_i = K \sin\left(\frac{2\pi t_i}{P} + \phi\right) + C \quad (3.1)$$

where K is the velocity semi-amplitude, P is the orbital period, ϕ is the orbital phase and C is the velocity constant. These are the parameters of the reflex motion of the star, although P , ϕ and C will be the same for the planet. f_i

is the set of velocities corresponding to a set of observation times t_i , given a particular set of the four parameters.

Our method involves choosing priors and limits for K , P and ϕ and solving analytically for C , scanning through a grid of values of the first three parameters, calculating χ^2 and then the probability for every combination of parameters and finally integrating (marginalizing) over all parameters except the one for which we wish to obtain a probability distribution.

3.1.1 Choice of Priors

It is important to choose a set of priors that accurately reflects our state of uncertainty about the truth of the model before analysing the data. For the period (P), we set a lower limit of 300 days and an upper limit of 30000 days. We use a Jeffreys prior, which corresponds to a uniform probability density in $\ln P$. Since the prior period range spreads over 3 log scale decades, a Jeffreys prior is a good choice as it says that the true value of P is just as likely to lie in any of the decades. It can be written as

$$p(P)dP = \frac{dP}{P \ln(P_H/P_L)} \quad (3.2)$$

where P_H and P_L are the upper and lower limits on P , respectively. The lower limit for the velocity semi-amplitude (K) is 1 m/s and the upper limit is $2(v_{max} - v_{min})$. The upper limit allows the velocity semi-amplitude to be up to twice as large as the largest difference between two velocities in the data set. We opted for a uniform prior for K with these limits. For the phase (ϕ) we used a uniform prior ranging from 0 to 2π radians. The prior for C is uniform and its range is from $-\infty$ to ∞ . The limits on C may seem extreme but they enable us to marginalize over C analytically, as described in more detail in section 3.1.3.

3.1.2 Calculating the Probabilities

Next, we set up a three-dimensional grid of K , ϕ and P values. The first two parameters were sampled uniformly while P was sampled uniformly in frequency (ν). We did not need to sample C because it is possible to find analytically the best-fit value of C for a given combination of K , P and ϕ , as shown in section 3.1.3. For a given set of measurements we calculated f_i for each combination of K , P , ϕ and C values using equation (3.1). We then determined χ^2 for each f_i as follows:

$$\chi^2 = \sum_{i=1}^N \left(\frac{v_i - f_i}{\sigma_i} \right)^2. \quad (3.3)$$

Finally, the probability was calculated using

$$p(d|K, P, \phi, C) = \frac{1}{(\sqrt{2\pi}\sigma_i)^N} e^{\frac{-\chi^2(K,P,\phi,C)}{2}} \quad (3.4)$$

where σ_i is the uncertainty associated with the i_{th} velocity measurement. Equation (3.4) comes from drawing each observed velocity from a Gaussian distribution [19], and multiplying together the probabilities for obtaining these individual observed velocities.

3.1.3 Marginalization

Unless otherwise specified, the method we used to marginalize over parameters is direct numerical integration. This consists of adding up the value of the integrand evaluated at each sampled value of the parameter over which we marginalized. We compared this to the Trapezoidal Rule and Simpson’s Rule (“Numerical Recipes”, section 4.1 [65]) and obtained only slightly differing results (i.e. the values obtained for the area under a given curve as calculated with direct numerical integration, and with the Trapezoidal Rule or Simpson’s rule, differed by only 0.1%). This means that we could use fewer grid points (sample the parameters less finely) and use Simpson’s rule to attain the same

accuracy as that obtained using direct numerical integration on more grid points. However, the former method would require the use of loops, while the latter method consists of simply summing the values of the integrand, thus requiring less operations and less computer time. Indeed, the time gained by the reduction in the number of grid points does not compensate the time lost through the use of loops. Therefore, we opted for direct numerical integration because it was the most efficient method.

Marginalizing over the observational uncertainties

Since each radial velocity measurement is calculated by averaging thousands of spectral lines over a large number of spectrum sections, the associated observational uncertainties can be accurately estimated and are almost Gaussian [13]. However, there can be additional causes of radial velocity variations that are not included in these uncertainties, such as stellar jitter or additional unseen companions around the same star [28]. We can produce the most conservative results by considering a fourth parameter s which we multiply by the measured observational uncertainty (σ_i) and over which we integrate equation (3.3) from 0 to ∞ . Essentially, this means that we assume the actual observational uncertainty can be infinitely large. This assumption may seem unrealistic, but we have verified that it is sufficiently accurate. Figure 3.1 shows two overlapping probability distributions for K for the same radial velocity data set. The difference between them is only a difference in the upper limits chosen for the marginalization over σ_i : for the top curve the upper limit was infinity while for the bottom curve it was 6 m/s (the mean value of the observational uncertainty for any data set was never larger than 4.5 m/s, so 6 m/s seemed a safe choice for a finite upper limit). The two distributions do not differ by much. Having verified that our choice did not affect the probability distributions, we opted for an upper limit of infinity. If anything, this choice produces more conservative probability distributions and, not least, it allowed us to solve the integral analytically. We used integration by parts repeatedly and solved a standard Gaussian integral to obtain

$$\int_0^{\infty} \frac{1}{(\sqrt{2\pi s\sigma_i})^N} e^{-\frac{\chi^2(K,P,\phi,C,s)}{2}} ds \propto [\chi^2(K,P,\phi,C)]^{(-N/2)}. \quad (3.5)$$

where the constant prefactor was dropped since only the odds ratio is used for the final result.

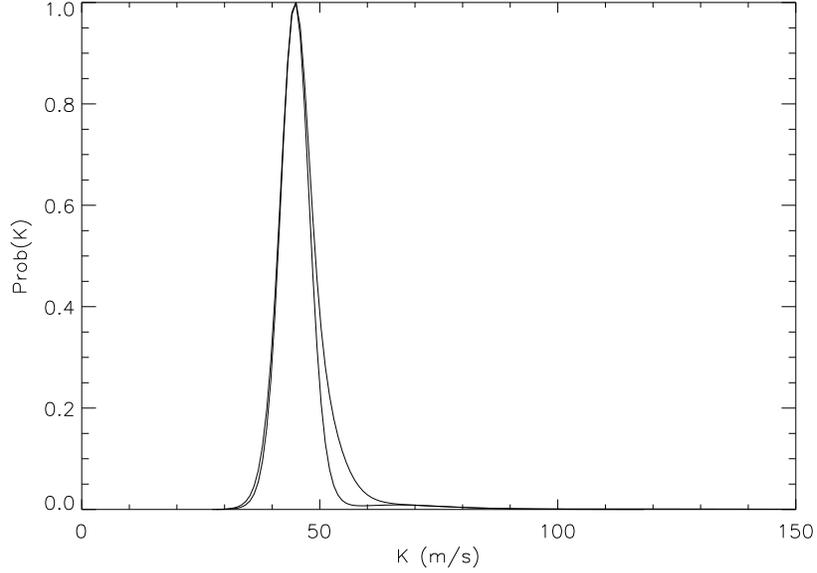


Figure 3.1 Two overlapping probability distributions for K for the same radial velocity data set. The top curve was obtained with ∞ as the upper limit for the marginalization over σ_i ; the bottom curve was obtained with 6 m/s as the upper limit.

Marginalizing over C

First, we define h_i as follows:

$$h_i = K \sin\left(\frac{2\pi t_i}{P} + \phi\right) \quad (3.6)$$

so we have

$$\chi^2 = \sum_{i=1}^N \left(\frac{v_i - C - h_i}{\sigma_i} \right)^2. \quad (3.7)$$

We can write equation (3.7) as a quadratic expression in C ,

$$\frac{\chi^2}{\sum_{i=1}^N \frac{1}{\sigma_i^2}} = C^2 - 2C \left(\frac{\sum_{i=1}^N \frac{(v_i - h_i)}{\sigma_i^2}}{\sum_{i=1}^N \frac{1}{\sigma_i^2}} \right) + \frac{\sum_{i=1}^N \frac{(v_i - h_i)^2}{\sigma_i^2}}{\sum_{i=1}^N \frac{1}{\sigma_i^2}}. \quad (3.8)$$

$(\chi^2)^{(-N/2)}$ can now be integrated analytically over C as follows:

$$\begin{aligned} \int_{-\infty}^{\infty} (\chi^2(K, P, \phi, C))^{(-N/2)} dC &\propto \left(\sum_{i=1}^N \frac{1}{\sigma_i^2} \right)^{(-1/2)} \left(\sum_{i=1}^N \frac{(v_i - h_i)^2}{\sigma_i^2} \right)^{..} \\ &\quad - \left(\sum_{i=1}^N \frac{1}{\sigma_i^2} \right)^{-1} \sum_{i=1}^N \left(\frac{v_i - h_i}{\sigma_i} \right)^2 \right)^{-\left(\frac{N-1}{2}\right)} \end{aligned} \quad (3.9)$$

Once again, we dropped the constant prefactor for the same reason as in the marginalization over σ_i .

In fact, equation (3.9) can be written in terms of the χ^2 corresponding to the best-fit C for each K , P and ϕ combination. As mentioned above, we can find the best-fit C analytically by taking the first derivative of equation (3.3) with respect to C and setting it equal to 0 to find the best-fit value of C , which corresponds to the minimum of χ^2 at a given combination of K , P and ϕ :

$$\left. \frac{\partial \chi^2}{\partial C} \right|_{k,p,\phi} = 2 \sum_{i=1}^N \left(\frac{v_i - f_i}{\sigma_i} \right) \left(-\frac{1}{\sigma_i} \frac{\partial f_i}{\partial C} \right) = 0. \quad (3.10)$$

We thus obtain

$$C_{bf} = \frac{\sum_{i=1}^N \left(\frac{v_i - h_i}{\sigma_i^2} \right)}{\sum_{i=1}^N \frac{1}{\sigma_i^2}}. \quad (3.11)$$

What equation (3.11) says is that to find C_{bf} all we need to do is subtract from each velocity measurement the velocity variation resulting from the orbital motion of the planet (which is a function of K , P and ϕ), and then find the weighted average of these residuals.

This means that equation (3.9) can be written very simply as

$$\int_{-\infty}^{\infty} (\chi^2(K, P, \phi, C))^{(-N/2)} dC \propto \left(\sum_{i=1}^N \frac{1}{\sigma_i^2} \right)^{(-1/2)} (\chi_o^2(K, P, \phi))^{-(\frac{N-1}{2})} \quad (3.12)$$

where χ_o^2 is just χ^2 evaluated at C_{bf} for each set of K , P and ϕ .

Marginalizing over ϕ

To simplify the integration over ϕ , we initially adopted a Gaussian approximation to the integrand. However, upon comparison with direct numerical integration, this approximation proved to be inaccurate. We noticed this important detail in the probability distribution for K . Using both methods, we ran the program on the radial velocity data for several known exoplanets and consistently obtained a discrepancy between the two methods. The top curve in figure 3.2 was obtained through direct numerical integration over ϕ while the bottom curve was obtained by using a Gaussian approximation to the integrand and solving the resulting Gaussian integral. There is less probability under the bottom curve because the Gaussian approximation method seems to have “missed” some of the ϕ probability distribution. The bottom curve also diverges at low values of K (less than ~ 7 m/s).

Next, we proceeded to understand the exact cause of this discrepancy, that is, why the Gaussian approximation fails. We looked at the probability

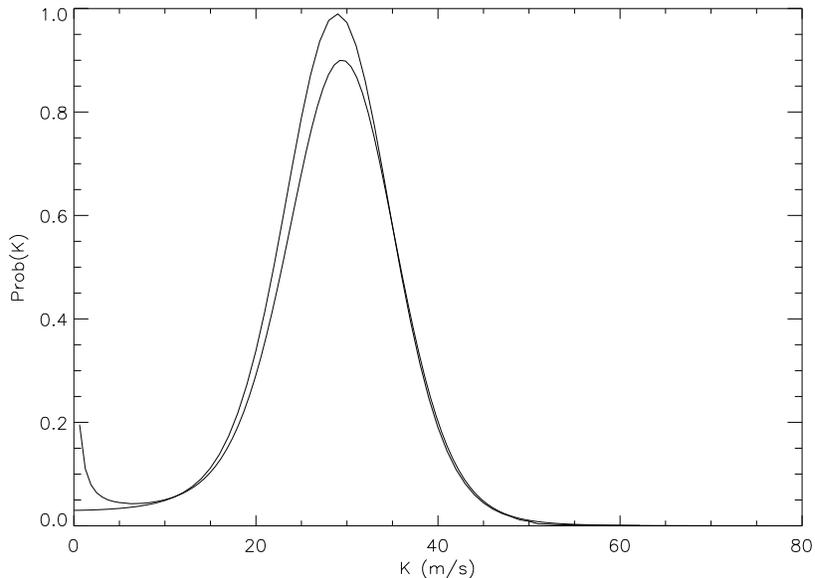


Figure 3.2 Probability distribution for K : for the top curve, we used direct numerical integration over ϕ ; for the bottom curve, we used a Gaussian approximation to the integrand.

distribution for ϕ computed for each of several combinations of K , P and C values. In several instances, the probability distribution was either multi-peaked or otherwise non-Gaussian. Figure 3.3 illustrates a typical example at low values of K . When fitting a Gaussian to such a distribution, the program will choose the largest value in the ϕ distribution to correspond to the peak of the Gaussian. The fitted Gaussian in this case includes more probability than the actual distribution, especially when considering the area under its tails, which extend to the left beyond 0 rad and to the right beyond 2π rad. Although some actual probability will still be “missed” (roughly between $3\pi/2$ and 2π on the figure), the effect just described will more than compensate for it, resulting in a calculated probability which is larger than is actually the case. This produces the divergence shown in figure 3.2 at low K .

Figure 3.4 shows a multi-peaked probability distribution for ϕ , which occurred for various combinations of values of the other three parameters. Fitting a Gaussian to the taller peak (the global maximum) while unaware of the pres-

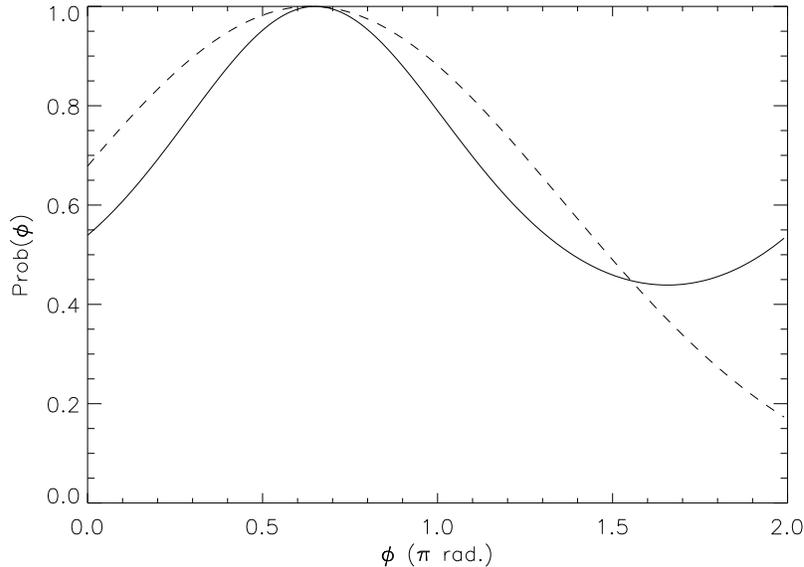


Figure 3.3 The solid curve is a typical example of the probability distribution for ϕ at low values of K . The dashed curve is the fitted Gaussian, which includes more probability than the actual distribution.

ence of the shorter peak will inevitably yield a lower probability than is the case, as is represented by the dashed curve. Such cases account for the “missing” probability represented by the bottom curve in figure 3.2. Although at values larger than ~ 7 m/s, the two curves agree within 10 to 15% and they have roughly the same overall shape, we could not be sure that this will be the case for all the radial velocity data sets we will analyse. Furthermore, since it originates from the integration over ϕ , this discrepancy could potentially affect the period, and for eccentric orbits, the eccentricity probability distributions. Finally, it is to be noted that the peaks of the two curves correspond to slightly different values of K . We could not predict whether this difference in the best-fit value of K would remain small for any of the candidates we analysed, nor any difference that might arise in the best-fit values of the period or eccentricity. To avoid these problems we employed direct numerical integration for marginalizing over ϕ .

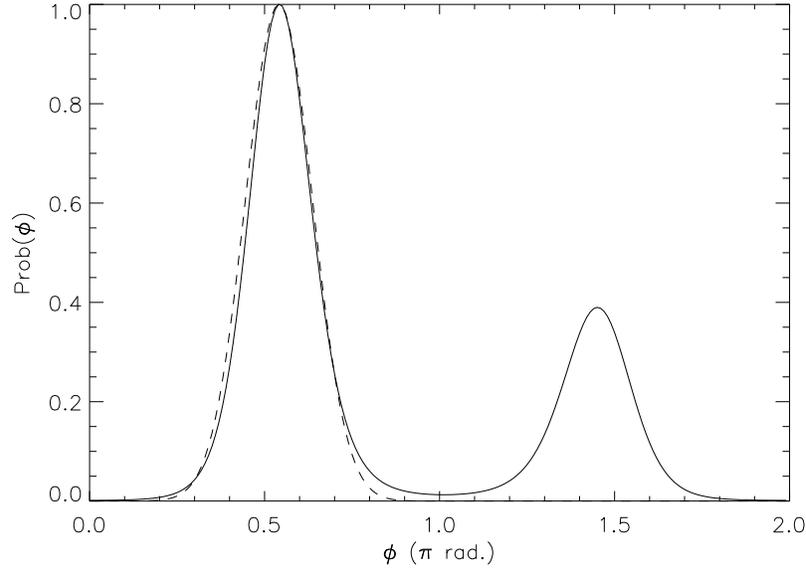


Figure 3.4 The solid curve is a typical example of the probability distribution for ϕ at various combinations of values of K , P and C . The dashed curve is the fitted Gaussian, which only includes one of the peaks.

Marginalizing over K and P

To obtain the probability distribution for K , we finally marginalized over P through direct numerical integration. To obtain the probability distribution for P , we marginalized over K instead.

3.1.4 The Odds Ratio

One other useful quantity which we can calculate is the odds ratio. This is basically the probability of the presence of a planet in the data vs. the probability that no planet is present. By marginalizing over all parameters for a given data set, we obtain a single value which we called p , the probability that a planet is present in the data. For each data set, we also find

$$q = \int_{-\infty}^{\infty} (\chi^2(C))^{-\frac{(N-1)}{2}} dC \quad (3.13)$$

which is the probability that no planet is present.

The odds ratio,

$$\Lambda = \frac{p}{q}, \quad (3.14)$$

tells us which model is preferred, planet or no planet.

3.1.5 Improving the Efficiency of the Calculations

In preparation for the eccentric orbit program, we proceeded to simplify and render our calculations more efficient. The improvements consisted of deriving an explicit formula for χ^2 in terms of the circular orbit parameters and the radial velocity measurements for a given exoplanet. In addition to the substitutions defined below, we will also use \sum instead of $\sum_{i=1}^N$.

$$\begin{aligned} \Omega &= \frac{2\pi}{P}, & w_i &= \frac{1}{\sigma_i^2}, & \langle v \rangle &= \frac{\sum w_i(v_i)}{\sum w_i}, & \langle v^2 \rangle &= \frac{\sum w_i(v_i^2)}{\sum w_i}, \\ \langle B_c \rangle &= \frac{\sum w_i(\sin \Omega t_i)}{\sum w_i}, & \langle A_c \rangle &= \frac{\sum w_i(\cos \Omega t_i)}{\sum w_i}, \\ \langle B_c^2 \rangle &= \frac{\sum w_i(\sin \Omega t_i)^2}{\sum w_i}, & \langle A_c^2 \rangle &= \frac{\sum w_i(\cos \Omega t_i)^2}{\sum w_i}, \\ \langle v B_c \rangle &= \frac{\sum w_i(v_i \sin \Omega t_i)}{\sum w_i}, & \langle v A_c \rangle &= \frac{\sum w_i(v_i \cos \Omega t_i)}{\sum w_i}, \\ \langle B_c A_c \rangle &= \frac{\sum w_i(\cos \Omega t_i \sin \Omega t_i)}{\sum w_i}. \end{aligned}$$

Using these substitutions, C_{bf} becomes

$$C_{bf} = \langle v \rangle - K \langle B_c \rangle \cos \phi - K \langle A_c \rangle \sin \phi \quad (3.15)$$

and the equation for $\frac{\chi^2}{\sum w_i}$ is

$$\begin{aligned} \frac{\chi^2}{\sum w_i} &= K^2(\langle A_c^2 \rangle \sin^2 \phi + 2 \langle B_c A_c \rangle \sin \phi \cos \phi + \langle B_c^2 \rangle \cos^2 \phi) \dots \\ &+ \langle v^2 \rangle - C_{bf}^2 - 2K(\langle v A_c \rangle \sin \phi + \langle v B_c \rangle \cos \phi). \end{aligned} \quad (3.16)$$

There are two reasons for re-writing χ^2 in the manner described above. The first is efficiency. The IDL programming language allows the use of vector operations. This means that, using equation (3.16), we can calculate not a single value but an entire matrix of χ^2 values with one line of code. The alternative would have been the use of significantly less efficient multiple loops through all the values of all the parameters we sampled.

The second reason becomes clear when we consider that all the values inside $\langle \rangle$ in equation (3.16) only need to be computed when P changes, and not when K or ϕ change, because these values only depend on P and the velocity measurements (of course, $\langle v^2 \rangle$ was computed only once for each data set since it does not depend on any parameters). For example, if we were to sample 10 values for each of the three parameters, we would only need to calculate the values inside $\langle \rangle$ 10 times as opposed to 10^3 times. This minimizes the number of operations required to generate values of χ^2 and thus allows us to do so more quickly. This is also very useful for the eccentric orbits program, as detailed in section 3.2.4.

With these tools in hand and using what we have learned from our work on circular orbit models, we began developing a program capable of fitting eccentric orbits.

3.2 Eccentric Orbits

As mentioned in section 1.3.1, the model used for fitting radial velocity data with eccentric orbits is

$$f_i = K(\cos(\omega + \theta_i) + e \cos \omega) + C \quad (3.17)$$

where K is the velocity semi-amplitude, e is the eccentricity, ω is the argument of periastron (measured from the plane of the sky), θ is the true anomaly and C is the velocity constant. P and ϕ are included in θ_i . This means that altogether we have six independent parameters. Once again, these parameters correspond to the reflex motion of the star, although e , P , ϕ and C will be the same for the planet. f_i and t_i are as described in section 3.1.

We chose priors and limits for K , P , e , ω and ϕ , and solved analytically for C . Next, we scanned through a grid of the first five parameters, calculated χ^2 and the probability for every combination of parameters, and obtained probability distributions for the parameters of interest by marginalizing over the other parameters.

3.2.1 Choice of Priors

For eccentric orbits we made a more careful choice of priors, as the resulting program is the one actually used in the analysis of our data. It is to be noted that a program for eccentric orbits can also accommodate circular orbits with $e, \omega = 0$.

We define a long period exoplanet to have a mass equal to or less than $10M_J$ and a period between 300 and 3×10^5 days (≈ 820 years). Although there are currently no confirmed planets with such large periods, the upper cut-off corresponds approximately to a period where perturbations from nearby stars and the galactic tide would disturb the planetary candidate's orbit [37]. We mostly adopted the priors used in Gregory (2007) [35]. As for circular orbits, we used a Jeffreys prior for P , with a uniform probability density in $\ln P$, since the prior period range spreads over 4 log scale decades.

We also used a Jeffreys prior for K , with K ranging between 1 m/s and $284(1\text{yr}/P_i)^{1/3}(1/\sqrt{1-e_i^2})(M_\star/M_\odot)^{-2/3}$ m/s. The upper cut-off for K corresponds to a $10M_J$ planet and comes from [19]

$$K = \frac{28.4\text{m/s}}{\sqrt{1-e^2}} \left(\frac{M_P \sin i}{M_J} \right) \left(\frac{P}{1\text{yr}} \right)^{-1/3} \left(\frac{M_\star}{M_\odot} \right)^{-2/3} \quad (3.18)$$

where $M_P \sin i / M_J$ was set to 10. For a given set of measurements, as the value of P we can fit through the data increases, the value of K must decrease in order to have an object of not more than $10M_J$. Additionally, as the value of e we can fit through the data increases, the value of K must decrease to fulfill the same condition. Thus, the upper limit on K depends both on P and eccentricity (e). The masses of the stars were obtained from Takeda et al. 2007 [73] and from Valenti and Fischer 2005 [76].

For the eccentricity e we used a uniform prior with a range from 0 to 1. The priors for ω and ϕ were also uniform, with a range from 0 to 2π .

3.2.2 Kepler’s Equation

In order to obtain θ_i (used in equation (3.17)), we had to solve Kepler’s equation

$$E_i - e \sin E_i = M_i \quad (3.19)$$

where E and M are related to θ as defined in section 1.3.1.

We solved Kepler’s equation (3.19) using an algorithm which finds the root(s) of a function within a given range (here, M lies between 0 and 2π) using Brent’s method. The value is refined until it reaches the accuracy specified by the user. For more information on Brent’s method, refer to “Numerical Recipes”, section 9.3 [65].

3.2.3 Analytic Solution for C_{bf}

The equation for C_{bf} is the same as equation (3.11),

$$C_{bf} = \frac{\sum_{i=1}^N \left(\frac{v_i - h_i}{\sigma_i^2} \right)}{\sum_{i=1}^N \frac{1}{\sigma_i^2}} \quad (3.20)$$

except that here h_i corresponds to an eccentric orbit:

$$h_i = K(\cos(\omega + \theta_i) + e \cos \omega). \quad (3.21)$$

As for circular orbits, the equation for C_{bf} says that it is simply the weighted average of the residuals obtained by subtracting from each velocity measurement the velocity variation resulting from the orbital motion of the planet (which for eccentric orbits is a function of K , P , e , ω and ϕ).

Introducing two more substitutions,

$$\langle A_e \rangle = \frac{\sum w_i \cos \theta_i}{\sum w_i} \quad \text{and} \quad \langle B_e \rangle = \frac{\sum w_i \sin \theta_i}{\sum w_i},$$

we can write C_{bf} as

$$C_{bf} = \langle v \rangle - Ke \cos \omega - K \langle A_e \rangle \cos \omega + K \langle B_e \rangle \sin \omega. \quad (3.22)$$

3.2.4 Calculating the Probabilities

Before calculating the probabilities, we first write χ^2 in terms of the eccentric orbit parameters and the radial velocity measurements for a given planet,

$$\begin{aligned} \frac{\chi^2}{\sum w_i} &= \langle v^2 \rangle - (C_{bf} + Ke \cos \omega)^2 - 2K(\langle vA_e \rangle \cos \omega - \langle vB_e \rangle \sin \omega) \dots \\ &+ K^2(\langle A_e^2 \rangle \cos^2 \omega + \langle B_e^2 \rangle \sin^2 \omega - 2 \langle B_e A_e \rangle \sin \omega \cos \omega) \end{aligned} \quad (3.23)$$

where $\langle vB_e \rangle$, $\langle vA_e \rangle$, $\langle B_e A_e \rangle$, $\langle B_e^2 \rangle$ and $\langle A_e^2 \rangle$ are analogous to the expressions defined in section 3.1.5 for circular orbits, with B_c and A_c replaced by B_e and A_e , respectively.

In section 3.1.5 we considered two reasons for writing χ^2 using these substitutions. We elaborate on how the second justification applies to the eccentric orbits case. Equation (3.17) shows that we need to calculate θ_i every time P , e or ϕ change, but not when K and ω change. So if we were to sample 10 values

for each of the five parameters, we would only need to calculate θ_i 10^3 times as opposed to 10^5 times. This increases the efficiency and thus the speed of the program by minimizing the number of operations needed to calculate the grid of χ^2 values.

For the eccentric orbit program we set up a five-dimensional grid of K , P , e , ω and ϕ values. We sampled 200 values of K , 400 values of P , 30 values of e and 100 values of ω . The sampling in ϕ was determined by an adaptive integration algorithm, which selected as many values as needed to achieve the specified level of accuracy in the results of the integration. This procedure is described in more detail below.

We calculated the probability using equation (3.4).

3.2.5 Marginalization

We marginalized over the observational uncertainty and over C in the same way as in section 3.1.3, using equation (3.21) instead of equation (3.6). The marginalization over the five remaining parameters is described in the following subsections. Unless otherwise specified, the method used to marginalize over parameters is direct numerical integration (see section 3.1.3).

Marginalizing over ϕ

Given the various shapes the probability distribution for ϕ can exhibit and in order to ensure we did not miss any peaks in the distribution, we decided to use an adaptive integration algorithm for the marginalization over this parameter. This IDL-based algorithm is called QPINT1D. It locates regions of the integration interval which contain the highest error (difference between the values of the integrand computed at two adjacent points) and concentrates on those regions. It does this by successively bisecting the starting interval (in this case 0 to 2π), assigning an error estimate to each subinterval and dividing further until each subinterval carries the same amount of error, which was specified by the user.

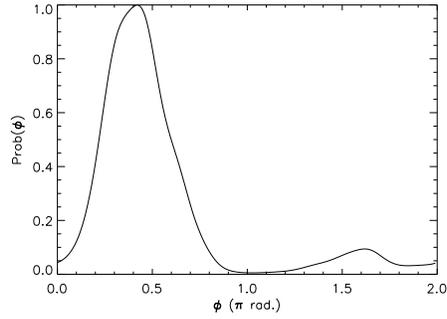


Figure 3.5 Probability distribution for ϕ , with $e = 0.5$, $K = 31.7$ m/s, $P = 59.6$ days and $\omega = \pi/3$.

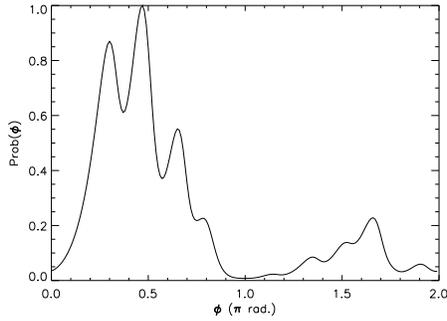


Figure 3.6 Probability distribution for ϕ , with $e = 0.7$, $K = 31.7$ m/s, $P = 59.6$ days and $\omega = \pi/3$.

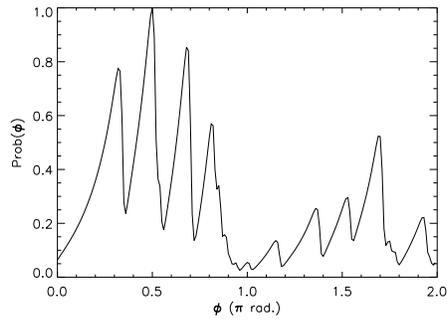


Figure 3.7 Probability distribution for ϕ , with $e = 0.9$, $K = 31.7$ m/s, $P = 59.6$ days and $\omega = \pi/3$.

For high values of e , the probability distribution for ϕ presents a large number of variations which are greater than the user-specified error. Thus, although this method proved indeed to be accurate, it is also slow because at high values of e the algorithm would compute the integrand for up to 6000 values of ϕ . Nonetheless, it was reasonably fast for our purpose and we used it throughout the entire analysis. Some examples of the probability distribution for ϕ at $e = 0.5$, $e = 0.7$ and $e = 0.9$ are shown in figures 3.5, 3.6 and 3.7, respectively. K , P and ω have the same values for each of the three figures.

Marginalizing over K , P , e and ω

Once again, we simply used direct numerical integration for the remaining four parameters. Ultimately we were interested in the probability distributions for P and e . Therefore we only stored the probability distributions marginalized over all parameters except those two.

We can also find the odds ratio as described in section 3.1.4.

3.2.6 Running Time Factors and Limitations

Although the duration of execution of the circular orbit program never exceeded 20 seconds, we nonetheless strove to decrease this time by implementing equation (3.16), in preparation for the eccentric orbit program. In the case of the latter, the addition of two more parameters (e and ω) increased the running duration by two or three orders of magnitude, depending on how finely the new parameters were sampled. In order to keep this effect under control, we tried running the program with several combinations of parameter sampling and observed the variations in the resulting probability distributions. The sampling ultimately chosen for each parameter was such that making it more fine no longer significantly affected the probability distributions for P and e (the ones that most interested us). The number of values sampled mentioned in section 3.2.4 are based on these tests.

We now take a closer look at how these choices affected the running time of the program. We sampled 200 values of K , 400 values of P , 30 values of e ,

100 values of ω and an average of the order of 500 values of ϕ . Multiplying these numbers gives 1.2×10^{11} calculations of χ^2 . The program takes on average 3 to 4 hours to run on a 2.1 GHz processor, depending on the number of measurements in a given data set. If it runs for 3 hours, this means each calculation of χ^2 for a combination of five parameters takes about 10^{-7} seconds.

One important limitation that this creates is the number of planets we can search for in a given data set. To fit a multi-Keplerian orbit to a set of radial velocity measurements as follows,

$$\mathcal{F}_i = \sum_{j=1}^M h_{i,j} \quad (3.24)$$

where M is the number of planets we wish to search for and h_i corresponds to equation (3.21), we would effectively need to sample $5 \times M$ parameters, or 5 additional parameters for each additional Keplerian orbit. If we were to keep the same sampling as for a single Keplerian orbit, searching for even two planets would then require $(1.2 \times 10^{11})^2 = 1.4 \times 10^{22}$ computations of χ^2 , which translates into an impractical duration of execution. The purpose of this thesis involves fitting single Keplerian orbits, but if we wished to accommodate multiple planet systems we would have to radically modify our method.

Chapter 4

Results and Discussion

The data to be analysed consisted of radial velocity measurements for 58 solar type stars, selected out of the 585 stars used in Cumming et al (2008) [20] (where single values for the best-fit period and exoplanet mass had been estimated using frequentist methods). The data were obtained with the HIRES echelle spectrometer at the Keck 1 optical telescope in Hawaii as part of the Keck Planet Search program [78]. The uncertainty on the radial velocity measurements is typically 3-5 m/s. Only data pertaining to stars with at least 10 observations over a period of two years or more were used in this analysis. The reader should refer to Cumming et al (2008) [20] for additional details on how the radial velocity data were obtained.

The selection was based on period (we chose the long period candidates, with $P > 300$ days) and estimated exoplanet mass (we chose candidates with $M < 10M_J$). In terms of mass, an exoplanet is loosely defined as having a mass lower than $12 - 13M_J$. Objects above this mass are often classified as brown dwarfs (see section 1.1). However, the boundary between the two is fairly blurred. Consequently, we conservatively chose $10M_J$ as an upper mass limit.

Of the 58 stars, 21 have confirmed planets (i.e. a published orbital solution) and the remaining 37 are unconfirmed candidates. We did not discriminate between those two categories during our analysis. It is to be noted that,

although the confirmed planets have a confirmed mass which corresponds to an exoplanet mass as described above, some or all of the unconfirmed candidates may turn out to be objects with masses above $10M_J$. This is because the values for the best-fit period obtained in [20] are based on the data currently available, which in most cases does not cover an entire period.

4.1 Comparison with Previous Work

Before producing probability distributions for all the data sets we used, we wanted to make sure our program produced accurate results. To accomplish this, we compared the period-eccentricity probability contour plot for HD72659 (Fig. 4.1(a)) with the one produced in Ford 2005 [27] (Fig. 4.1(b)). This is a plot of the probability distribution marginalized over all parameters except the period and eccentricity. We used the same data set as in [27], which consists of fewer measurements than we have available today. To obtain the contours, we calculated the height of the probability distribution above which 68.3% (1σ), 95.4% (2σ) or 99.73% (3σ) of the total volume encompassed by the distribution is contained. In other words, 68.3% of the distribution lies above the height defined by the 1σ contour, and so on.

Overall, the probability distributions match. Nevertheless, a few differences are present. In Figure 4.1(a), the 1σ contour does not extend to as high periods as in Figure 4.1(b). Additionally, it does not include $e = 0$ in the left-hand plot, while it does in the right-hand plot. We believe these discrepancies originate from the finite grid sampling used in our method, whereas the right-hand plot was produced using Markov Chain Monte Carlo (see section 2.2). The latter technique is adaptive and thus can sample more finely in some regions of $P - e$ space as needed, while our approach uses the same sampling throughout the entire $P - e$ space. Furthermore, the high e and P end of the contour plot tail reaches higher on the e axis in the left-hand plot than on the right-hand plot. We are not sure where this difference comes from, but we believe the difference in sampling may partly account for it as well. In conclusion, we attributed the

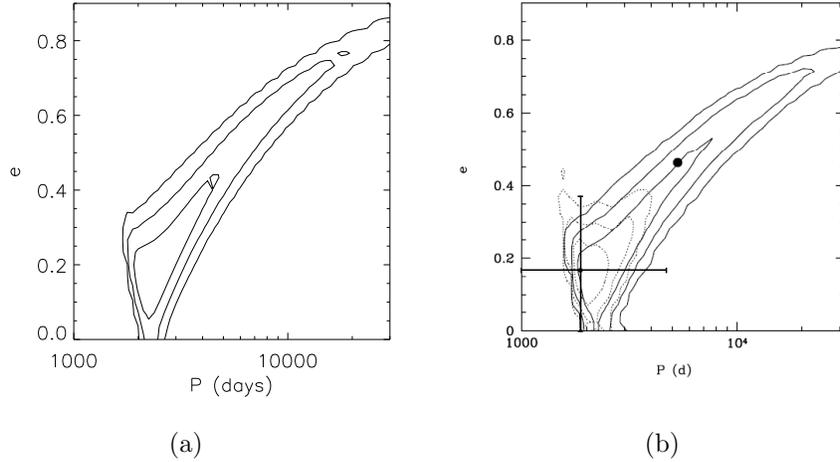


Figure 4.1 Probability distribution marginalized over all parameters except the P and e , for HD 72659. The solid contours indicate the 1, 2 and 3 σ confidence intervals, corresponding to 68.3%, 95.4% and 99.73% of the probability distribution. a) Our contour plot. b) The contour plot taken from [27] (the dotted contours represent 1, 2 and 3 σ confidence intervals obtained through a different method, and the dot represents the best-fit orbital solution based on later observations made after the rest of the plot had been composed). The plot in subfigure b) is used with the permission of E.B. Ford.

minor differences to the different methods used to produce each plot (simple grid sampling for Figure 4.1(a) and MCMC for Figure 4.1(b)).

We then proceeded to analyse the remainder of the data sets.

4.2 Probability Distributions

Having run our program on the 58 sets of measurements, we produced probability distributions for P and e for each set. We include plots of the associated radial velocity measurements (Figure 4.2). We also include six examples of distributions for each of the two parameters as well as six examples of the contour plots for both parameters (Figures 4.3, 4.4 and 4.5. These six examples correspond to six unconfirmed candidates.

Figure 4.3 shows a good variety of shapes and breadths of the period probability distributions. In general, the period is well constrained within approxi-

mately one order of magnitude. We can also observe this in the contour plots: the 1σ confidence interval most often encompasses less than one order of magnitude for the period. This tells us that even with radial velocity measurements covering only a fraction of an orbit, it is still possible to place reasonably good constraints on the periods of these candidates.

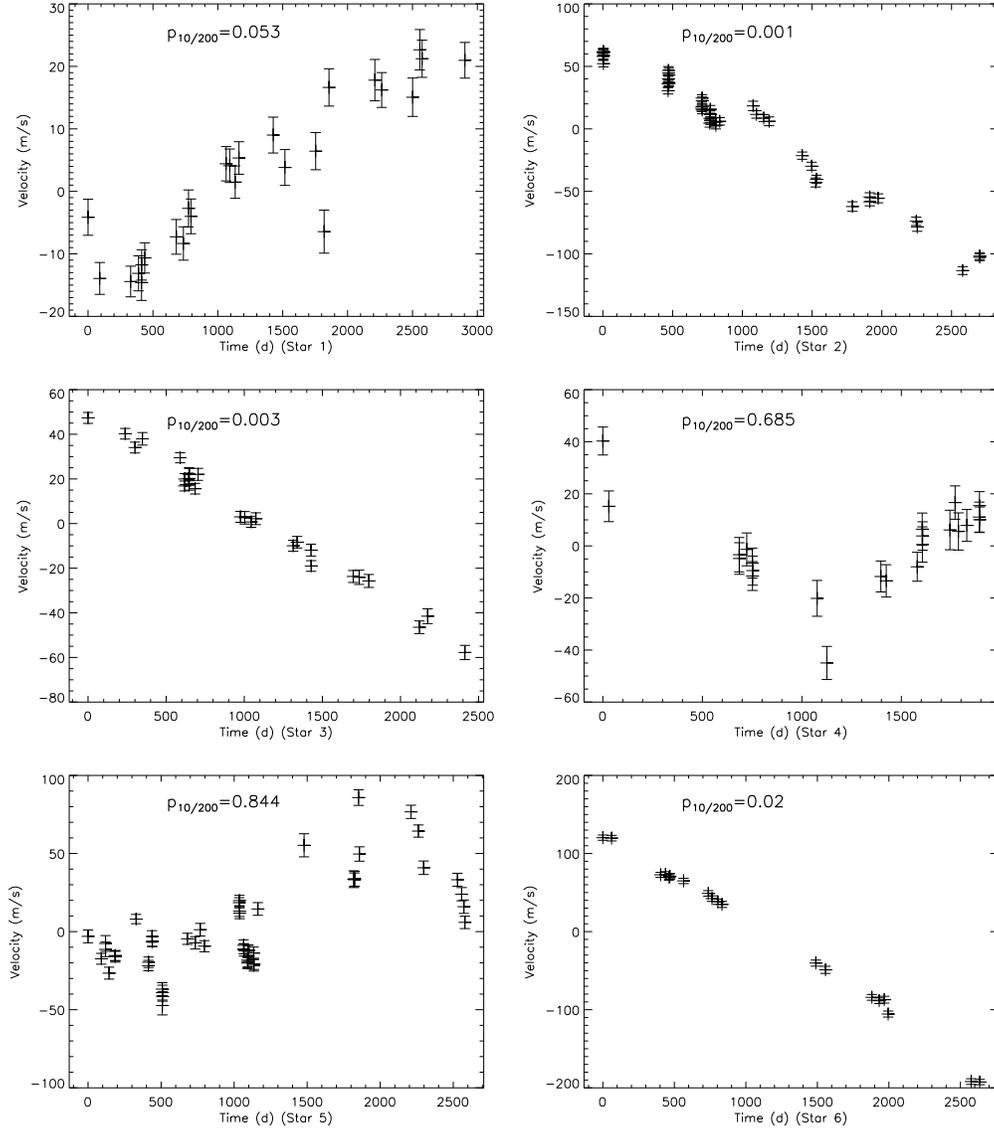


Figure 4.2 Radial velocity measurements and $p_{10/200}$ values (see text for details) for six unconfirmed candidates.

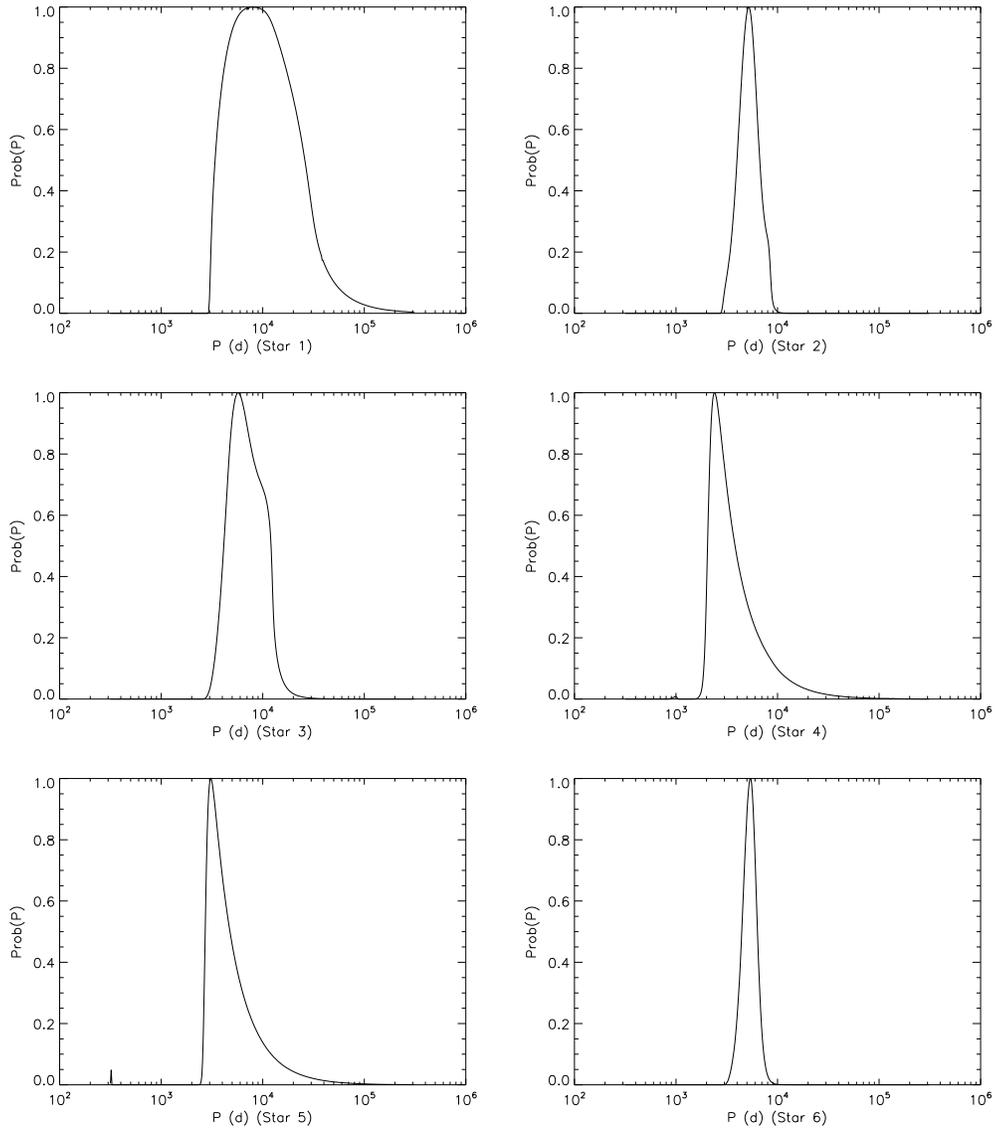


Figure 4.3 Period probability distributions for six unconfirmed candidates.

On the other hand, the eccentricity (Fig. 4.4) is much less constrained. Again, this can also be seen in the contour plots, where in many cases the 1σ confidence interval extends over the entire range of eccentricity. One way to explain this effect is as follows. For the period, we can always establish a fairly clear lower limit; any period below this limit simply cannot fit the data at all, no matter how we vary the other parameters. For example, in the plot for Star

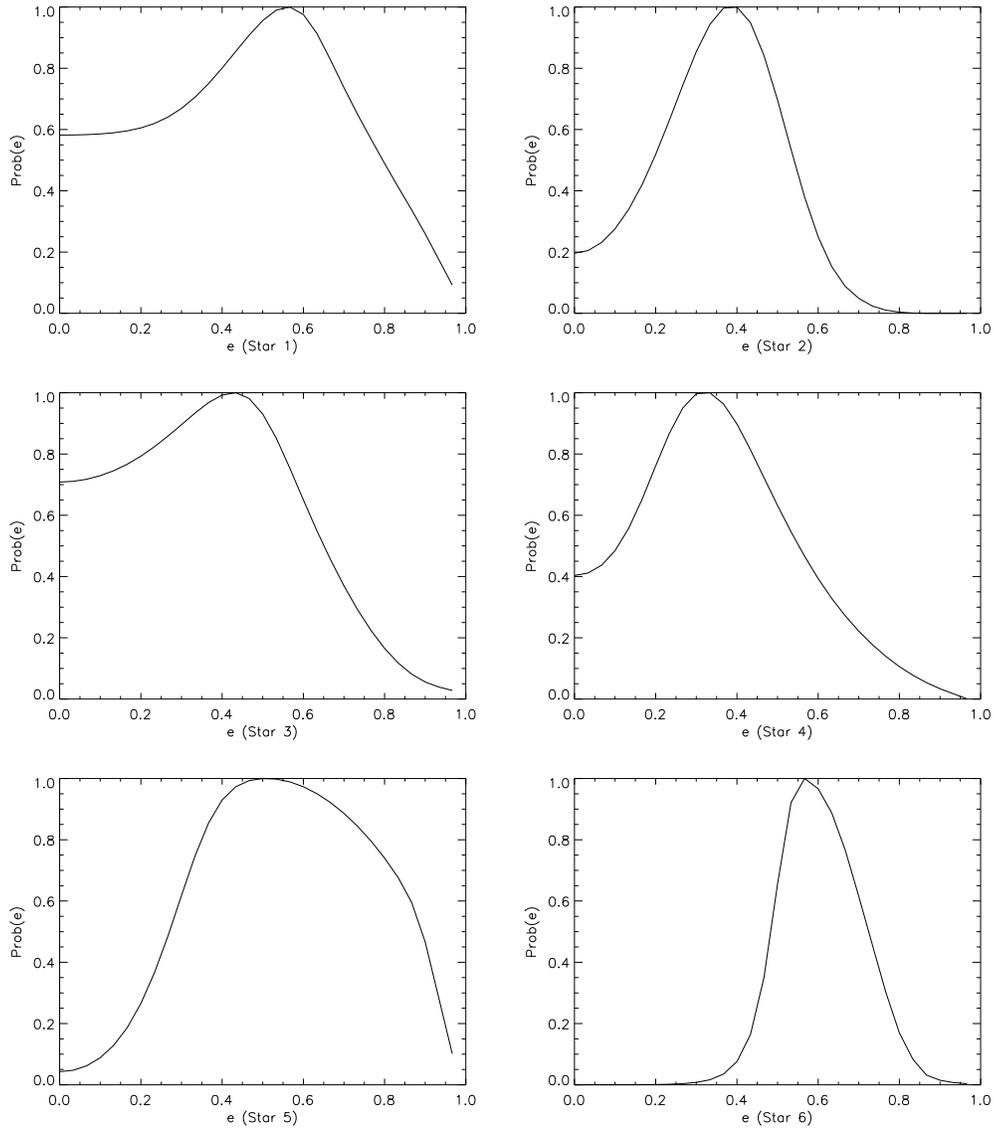


Figure 4.4 Eccentricity probability distributions for six unconfirmed candidates.

3 in Figure 4.2, we would be hard put to fit an orbit with a period of less than 2500 days. Although less abrupt than the lower limit, we can also recognise an upper limit on period since the tail of the distribution does not generally reach 3×10^5 days (the upper end of the P sampling). This upper limit is due to the upper limit we chose on exoplanet mass ($10 M_J$) and hence on K .

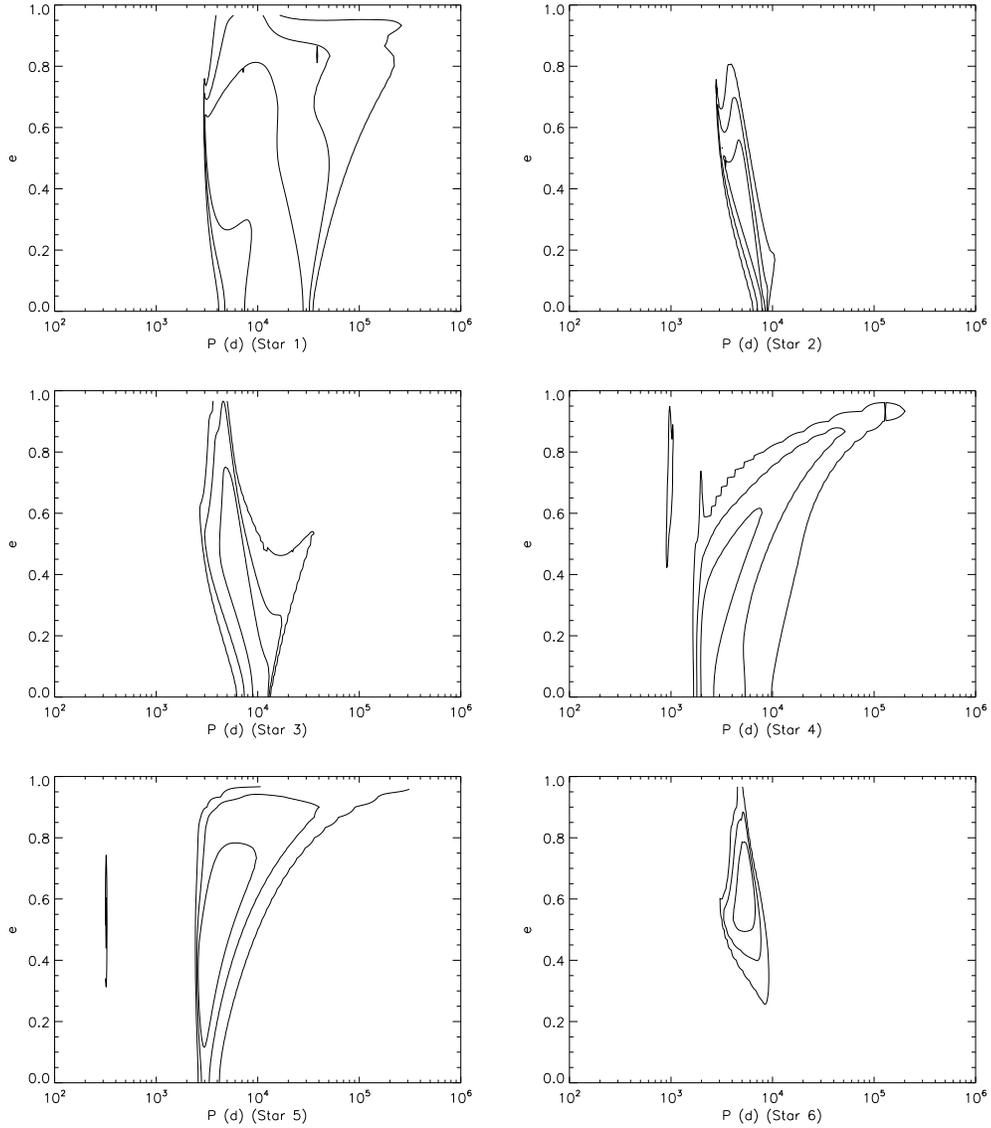


Figure 4.5 Period-eccentricity probability distributions for six unconfirmed candidates. The contours indicate the 1, 2 and 3 σ confidence intervals, corresponding to 68.3%, 95.4% and 99.73% of the probability distribution.

This says that for the orbit of a candidate to have a period larger than the upper limit seen in the corresponding distribution, the candidate must be more massive than $10 M_J$, which is beyond the K values sampled. For example, the candidate orbiting Star 3 cannot have a period longer than $\approx 3 \times 10^4$ days if

it has a mass lower than or equal to $10 M_J$. In the case of the eccentricity, however, we can often vary the other parameters in such a way as to obtain a decently fitting orbit of any eccentricity. Thus, a clear lower or upper limit for this parameter cannot be easily defined.

Most of the probability distributions as a function of both period and eccentricity (Fig. 4.5) exhibit a common feature: part of the distribution (in some cases, the entire distribution) tends to curve towards the upper right corner of the plot. In other words, the best-fit eccentricity increases with the best-fit period according to a specific function. Our attempt to predict this correlation is described in section 4.3.

As discussed in the previous chapter and at the beginning of this chapter, we selected the confirmed planets and unconfirmed candidates for the analysis based on the $10M_J$ criterion (as well as the previously estimated best-fit period) and chose an upper limit for the velocity semi-amplitude accordingly. However, some of these objects may prove to have higher masses. To account for this, we wished to find the probability that a given object is a planet (as defined above) versus a higher-mass object ($p_{10/hm}$). We ran our program on all 58 data sets once again, modifying the upper limit on K to correspond to $200M_J$. We found that beyond $200M_J$, the probability distributions for P and e as well as $p_{10/hm}$ did not change much, so we chose this value as a “maximum” mass for a higher-mass object. Thus, $p_{10/hm}$ became $p_{10/200}$, which is calculated as follows:

$$p_{10/200} = \frac{p(d|\mathcal{M}_{10})}{p(d|\mathcal{M}_{200})} \quad (4.1)$$

where d is the data, \mathcal{M}_{10} is the model for an object with mass less than $10M_J$, and \mathcal{M}_{200} is the model for an object with mass less than $200M_J$. We computed this value for each of the 58 data sets. It is noteworthy that generally $p_{10/200}$ was much higher for the confirmed planets than for the unconfirmed candidates, as expected.

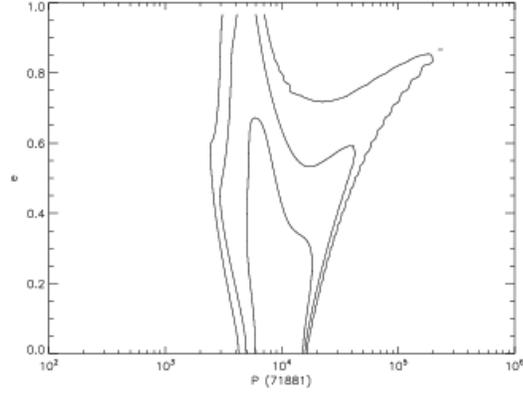
4.3 A Closer Look at the Period-Eccentricity Probability Distributions

In Figure 4.5, the period-eccentricity probability contour plots for Stars 1,4 and 5 exhibit a branch that curves to higher eccentricity at higher period, although it is less pronounced in the first plot. A second branch is often present as well, as in the plot for Star 3 in the same figure. In fact, over half of the contour plots we prepared display both branches. A representative example is shown in Figure 4.6(a). In some cases, only the right-most branch is present (Figure 4.6(b)).

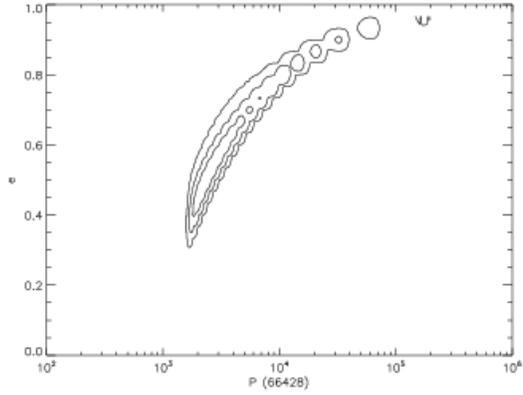
There are also a few distributions that look like HD 190406 (Figure 4.5) and for which P and e seem relatively well constrained. This is due to the limited range of K values we searched. As we discussed before, the upper limit we used for K corresponded to $10 M_J$. In cases where the previously estimated planet mass [20] lay between 4 and $10 M_J$ the contour plots looked like that for Star 2 in Figure 4.5. It seemed somewhat suspicious that the planets with the largest estimated masses (and likely the largest best-fit K) would be the ones best constrained in the P - e plane. When we raised the upper limit on K to correspond to $200 M_J$, keeping everything else unchanged, we found that the period-eccentricity distributions for these objects invariably broadened and assumed the two-branch shape present for other objects. A good illustration of this can be found in Figure 4.7. Upon increasing the sampling range for K , we obtained many more possible orbital solutions within a given confidence interval. By essentially eliminating the (planetary) upper limit on K , orbital solutions are found that fit the data much better (i.e. that have much higher probability) than the solutions found for an object between 0 and $10 M_J$. In other words, the probability distribution of Figure 4.7(a) hides in the probability distribution of Figure 4.7(b) which is much higher.

This observation does not affect other aspects of our analysis, but it was useful in interpreting the probability distributions.

Returning to the two branches described at the beginning of this section, we



(a)



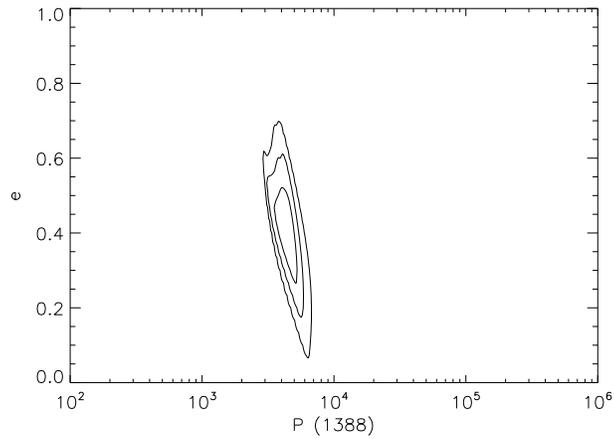
(b)

Figure 4.6 Period-eccentricity probability distributions, showing a) two branches and b) a single branch.

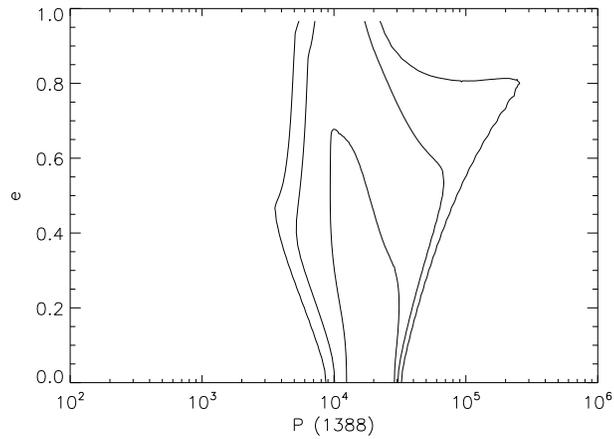
believe they can be predicted analytically. In Cumming 2004 [19], it was found that for some long period candidates, when the data cover only a fraction of the orbit, it is possible to fit a quadratic to the data $f_i = at_i^2 + bt_i + c$ and obtain as good a fit as with a circular orbit sinusoid (equation (3.1)). A quadratic fit then gives the following relation between K and P :

$$K = \frac{P}{2\pi} \left(b^2 + P^2 \frac{c^2}{\pi^2} \right)^{1/2}. \quad (4.2)$$

Essentially, equation (4.2) says that when a quadratic is a good fit to the data, these two parameters scale as $K \propto P^2$. We have decided to use the same



(a)



(b)

Figure 4.7 Period-eccentricity probability distributions. a) The upper limit on K corresponds to $10 M_J$. b) The upper limit on K corresponds to $200 M_J$.

approach for eccentric orbits. By fitting a quadratic to data that cover only part of a full period, we hope to find a relation between P and e (possibly dependent on some of the other parameters as well) which describes the two branches apparent in most of the period-eccentricity probability distributions. At the time of this writing, it is still a work in progress.

4.4 Average Period Probability Distribution

The next step in this statistical analysis consisted of producing a weighted average of the period distributions. The distributions were weighted by their corresponding probability of representing an object with mass $< 10M_J$, $p_{10/200}$ (obtained as described in section 4.2), and then averaged over the 58 long period exoplanets and candidates. The weighted average of the period distributions is shown in Figure 4.8. (The weighting step also allows us to use only the term “planet” when discussing this plot.)

One conclusion that can be drawn from this plot is that there is a much higher probability to have planets with periods between 300 and 10^4 days than planets with periods above 10^4 days. This is consistent with results obtained in Cumming et al. 2008 [20], where the period distribution for $P < 2000$ days was extrapolated to predict the occurrence rate of long period planets. It was found that 17% to 19% of solar type stars are expected to host a gas giant with period less than 3.2×10^4 days. However, the actual numbers of confirmed planets and candidates detected to date indicate that 18% of solar type stars host a planet with period less than only 1.2×10^4 days. These figures would suggest that extremely few stars are expected to have planets with periods between 1.2×10^4 and 3.2×10^4 days. Figure 4.8 supports this conclusion. Nevertheless, the uncertainties in the orbital parameters must be included in the statistical analyses before we can extract more convincing hypotheses.

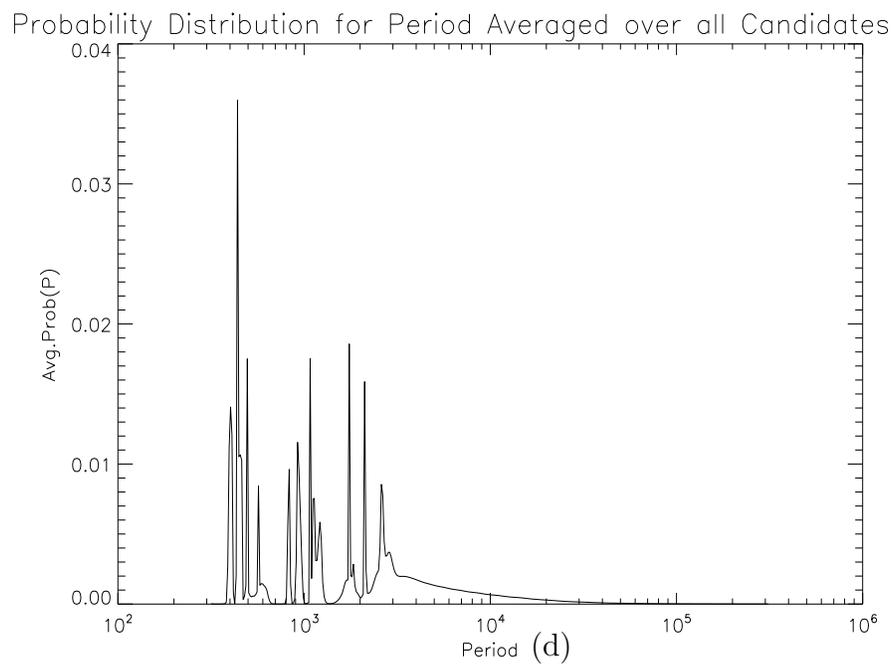


Figure 4.8 Average probability distribution for the period for the 58 long period confirmed planets and candidates.

4.5 Period-Metallicity Trend

Inspired by recent metallicity and stellar mass studies in the context of exoplanets ([25], [32]), we wanted to see if a correlation might exist between the period and the stellar metallicity. Since there already seems to be a connection between metallicity and stellar mass [50], we also plotted the period versus the stellar mass. We extracted the best-fit period for each planet and candidate from the corresponding period probability distribution. The stellar masses were obtained from [73] and [76].

The plot of the best-fit periods versus the stellar mass is shown in Figure 4.9. Although no trend is apparent, it should be noted that the masses of most of these stars are close to solar mass, and as such this distribution is not representative of a very wide range of stellar masses.

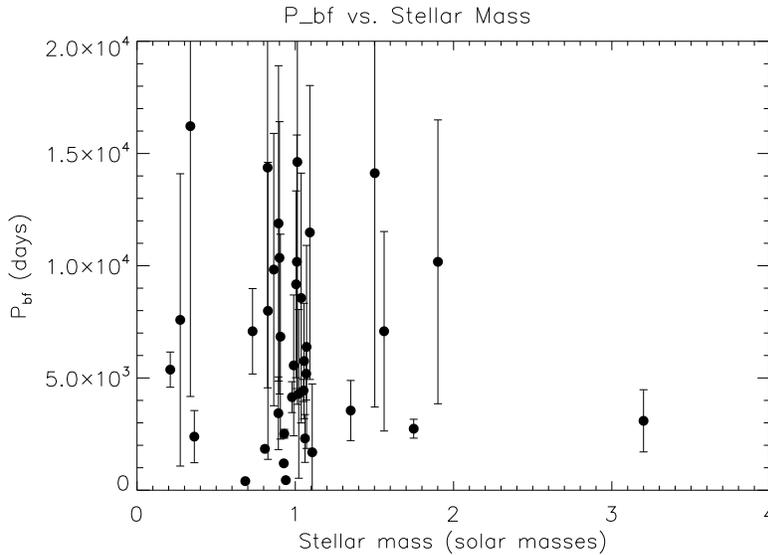


Figure 4.9 P_{bf} –stellar mass distribution for the 58 long period confirmed planets and candidates ($P > 300$ days). The error bars correspond to the 1σ confidence intervals of the associated probability distributions.

Stellar metallicity is generally expressed as “[Fe/H]”, which represents the logarithm of the ratio of a star’s iron abundance compared to that of the Sun. We obtained the stellar metallicity for each of the 37 stars with unconfirmed

candidates from [54], [9], [60], [73] and [76]. The plot of the best-fit periods versus the stellar metallicity is shown in Figure 4.10. We have then produced the same plot for 257 confirmed planets with available stellar metallicity data (regardless of period duration), for which we have obtained the best-fit periods and metallicities from *exoplanet.eu/catalog.php*, and for the unconfirmed candidates combined in Figure 4.11.

Figure 4.10 shows a trend of decreasing period with increasing metallicity. The circled points correspond define a possible outer envelope. When plotting the distributions for the known exoplanets and the unconfirmed candidates together in Figure 4.11, the trend is still apparent at periods longer than ≈ 5000 days; however, it is not noticeable among the confirmed planets at periods below 5000 days.

Before establishing whether what we observed is a real correlation between period and stellar metallicity, we must first check that this is not caused by any selection effects. We first looked at the duration of observations. Radial velocity measurements for many higher metallicity stars have only begun to be collected recently, motivated by studies of a possible correlation between stellar metallicity and planet frequency (see section 1.5). For data sets with shorter durations of observations, our program favors shorter best-fit periods. This can be seen in Fig. 4.3: the probability distributions have a more pronounced tail towards longer periods, with the best-fit period generally lying nearer the lower period cutoff than the higher period cutoff. Consequently, it was possible that we found lower best-fit periods for candidates around stars with higher metallicities, if the duration of observation was shorter than for stars with lower metallicities. Figure 4.12 shows our study of this possible selection effect.

There does seem to be a decrease in duration of observations at higher metallicities. However, only the right-most point corresponds to one of the circled envelope points which determine the trend in Figure 4.10. Although we cannot conclusively show it, we can tentatively assume that the duration of observations is not significantly responsible for the trend.

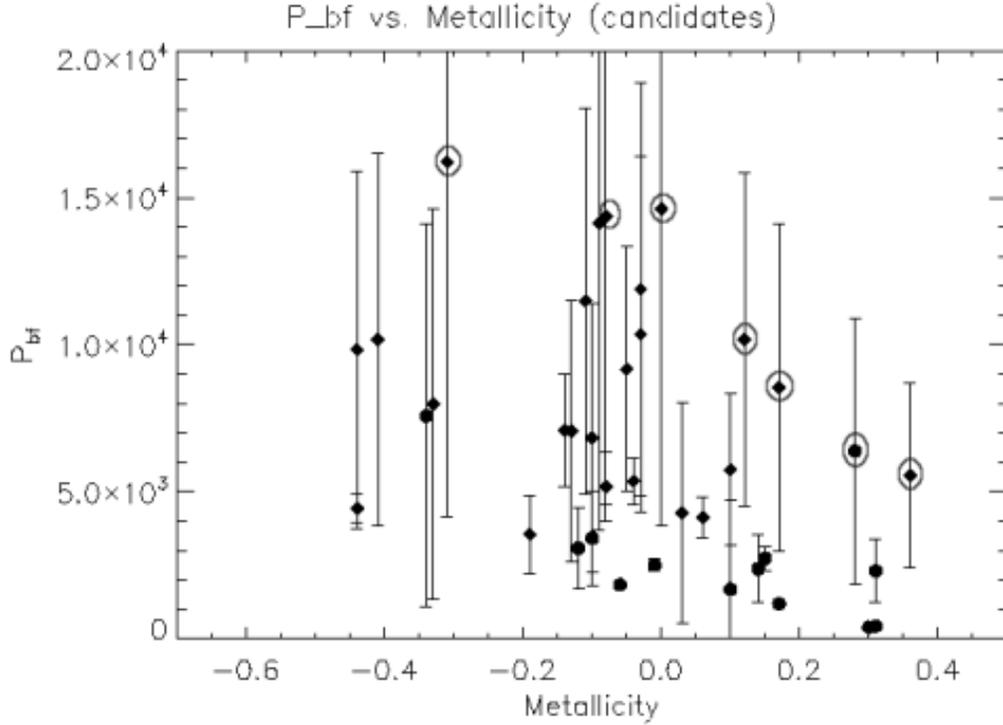


Figure 4.10 P_{bf} –stellar metallicity distribution for the 37 long period unconfirmed candidates. The diamonds correspond to candidates with $p_{10/200} < 0.1$ while the circular points correspond to candidates with $p_{10/200} > 0.1$. The error bars correspond to the 1σ confidence intervals of the associated probability distributions. The metallicity is expressed as $[\text{Fe}/\text{H}]$, where 0.0 corresponds to the metallicity of the Sun. The circled points constitute the outer envelope of a possible trend (see text for details). Not plotted here for visual ease are the error bars for the metallicity values: the magnitude of these error bars corresponding to the 1σ confidence interval is 0.019 [76], which is roughly the diameter of the data points on this plot.

Next, we checked whether the observational uncertainty in the radial velocity measurements might be of concern. We plotted the mean uncertainty for each candidate versus metallicity in Figure 4.13. There are a few outlying points, but no clear correlation. None of the outlying points correspond to any of the circled envelope points from Figure 4.10. We concluded that this effect is not likely to have caused the period-metallicity trend.

In Figures 4.10, 4.12 and 4.13 we have distinguished between candidates

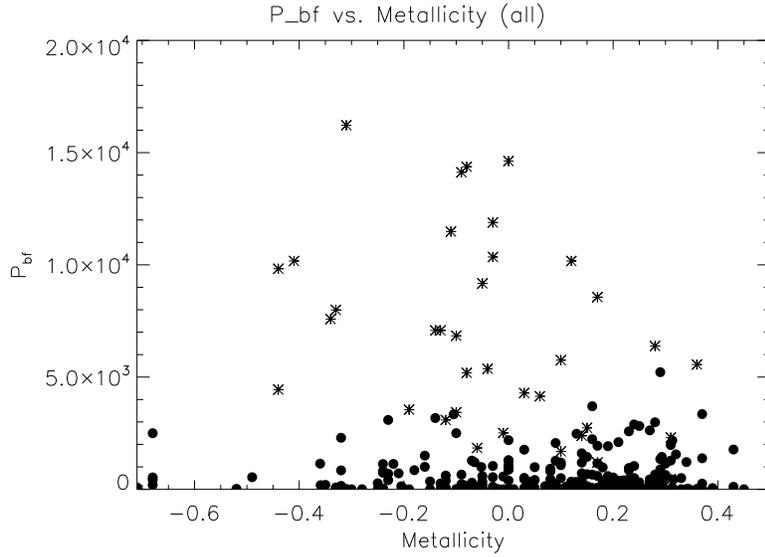


Figure 4.11 P_{bf} —stellar metallicity distribution for the 37 unconfirmed candidates and all 257 confirmed planets (for which stellar metallicity information was available) combined. The confirmed planets are represented by dots and the candidates are represented by asterisks. The metallicity is expressed as $[Fe/H]$, where 0.0 corresponds to the metallicity of the Sun.

with values of $p_{10/200}$ below and above 0.1. Indeed, since those candidates with $p_{10/200}$ less than 0.1 are unlikely to be planets (see section 4.2), they could justifiably be removed from these three plots. As more radial velocity measurements are added to the data sets of these candidates, their corresponding $p_{10/200}$ values may either increase or decrease, thus validating or invalidating the period metallicity trend. For now we plot them alongside candidates with $p_{10/200} > 0.1$ and distinguish them as described above.

Finally, we determined the correlation coefficient (r), which measures the strength and direction of a linear relationship between two variables. A value of r close to +1 or -1 indicates a high degree of correlation and a good fit to a positive or negative linear model respectively. A value of r close to 0 indicates a poor fit to a linear model. The correlation coefficient we find between the best-fit period and the stellar metallicity is -0.42. This indicates that there is some correlation, but the distribution is too scattered to be considered a good

fit to a linear model. It is also possible that the relationship between the two variables is accurately described by a non-linear model, which could explain why the value of r is closer to 0 than to -1.

We must also do a Kolmogorov-Smirnov (KS) test on the period-metallicity distribution in order to assess how significant the trend is statistically. Specifically, we want to quantify the difference between the distribution above 5000 days (candidates), and that below 5000 days (Figure 4.11). The KS test does this by computing the KS statistic, taking into account the number of points in each distribution. However, at the time of this writing we have not yet had time to perform this test.

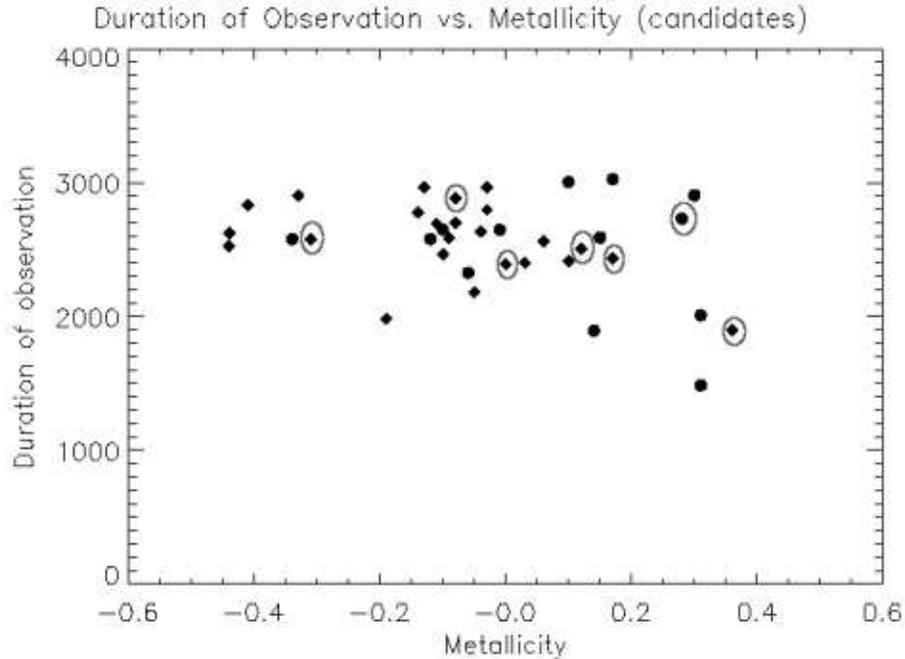


Figure 4.12 Duration of observations versus metallicity for the 37 unconfirmed candidates. The diamonds correspond to candidates with $p_{10/200} < 0.1$ while the circular points correspond to candidates with $p_{10/200} > 0.1$. The circled candidates are the same as the circled envelope candidates in Figure 4.10. The metallicity is expressed as $[\text{Fe}/\text{H}]$, where 0.0 corresponds to the metallicity of the Sun. The magnitude of the error bars corresponding to the 1σ confidence interval for the metallicity values is 0.019 [76], which is roughly the diameter of the data points on this plot.

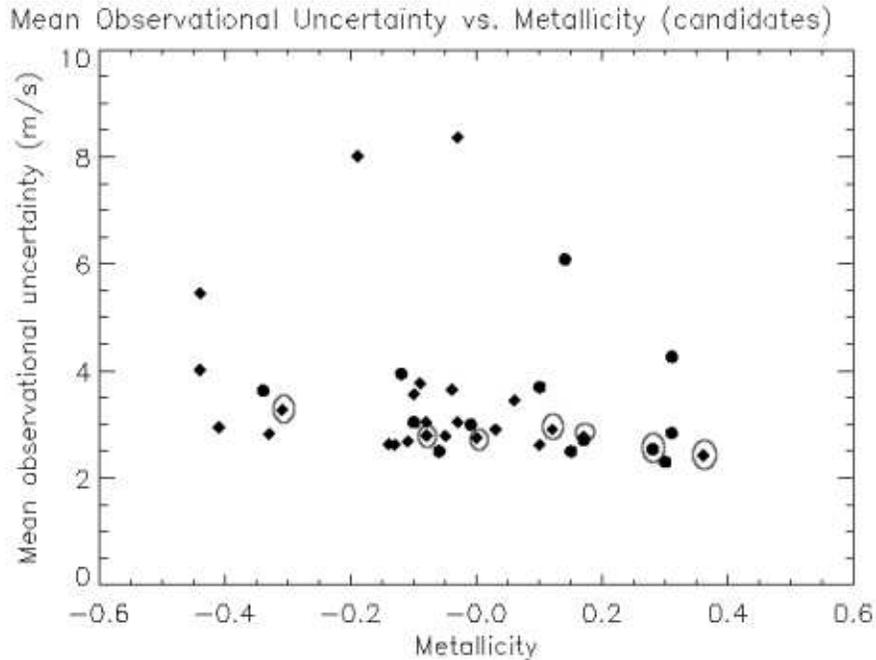


Figure 4.13 Mean observational uncertainty versus metallicity for the 37 unconfirmed candidates. The diamonds correspond to candidates with $p_{10/200} < 0.1$ while the circular points correspond to candidates with $p_{10/200} > 0.1$. The circled candidates are the same as the circled envelope candidates in Figure 4.10. The metallicity is expressed as $[\text{Fe}/\text{H}]$, where 0.0 corresponds to the metallicity of the Sun. The magnitude of the error bars corresponding to the 1σ confidence interval for the metallicity values is 0.019 [76], which is roughly the diameter of the data points on this plot.

If the decreasing period with increasing metallicity trend is indeed real, it could be explained by some current planet migration models. Sandquist et al. [68] have shown using simulations that it is possible for planets to migrate toward and become absorbed by their host star; depending on the composition of the planet’s interior, this process can result in an increase in the star’s metallicity. In this context, close-in planets (or planets with shorter periods, since the period is related to the semi-major axis through (1.1)) could constitute the tail of a stream of planets that have been consumed by the host star. Thus, we can speculate that planets which are further out (longer periods) are associated with lower metallicity stars which did not get a chance

to have their meal. In this frame of thought, we then expect such systems to host more than one planet. A follow-up project on this topic would be to search for other companions by fitting multi-Keplerian orbits to the data sets corresponding to the unconfirmed candidates.

Chapter 5

Conclusion and Future Work

The goal of this thesis was to place constraints on the period and eccentricity of long period exoplanets, and to shed some light on questions such as the incidence of such exoplanets. Motivated by the advantages of and the increasing use of Bayesian inference in the characterisation of exoplanet orbits ([27], [34]), we adopted a Bayesian approach for our own analysis. We first developed a program to fit circular Keplerian orbits to radial velocity data, and then upgraded to an eccentric orbit program. Along the way, we derived analytical solutions for some of the calculations while others we did numerically. We also found that a Gaussian is not always a good approximation to the probability distributions for the phase. From our results, we were able to place some qualitative constraints on the probability distributions for the period and eccentricity, including a possible correlation between those two parameters. Our average probability distribution for the period supported the results obtained by Cumming et al. (2008) [20] regarding the incidence of long period exoplanets. Finally, we revealed a possible trend suggesting that, at long periods, the orbital exoplanet period decreases with increasing host star metallicity.

Near future work involves obtaining the latest radial velocity data for the long period planets and candidates we analysed (the data used for this thesis

ends in 2004 (see [20]), and producing updated period and eccentricity distributions. Further work should be done on analytically predicting and understanding the period-eccentricity correlation. The possible period-metallicity trend should be verified more carefully, and a Kolmogorov-Smirnov test should be done to compare the candidates and confirmed planets distributions, thus assessing the significance of the trend. Nevertheless, it is important to keep in mind the large error bars on the best-fit period estimates. We hope that with more recent data we will be able to reduce these error bars for some of the candidates.

A far future goal would be to simultaneously search for multiple planets within a same system (same set of radial velocity data). However, as we've shown at the end of chapter 2, this would require either modifying our current computational approach, or using a different method (MCMC) or a combination of different methods.

Ongoing radial velocity surveys (HARPS [62], HIRES [12], ELODIE [2]) continue to acquire data while budding programs such as MARVELS [55] (due to go online in fall 2008) promise to search for planets around thousands more stars and to achieve unprecedented levels of sensitivity. Transit surveys such as the space based Kepler mission [10] (due for launch in 2008) will also provide further candidates for radial velocity follow-ups. With these new data and continuing progress on numerical simulations, we suspect the answers to many planet formation and occurrence questions are forthcoming.

Bibliography

- [1] C. W. Allen. *Allen's Astrophysical Quantities*. AIP Press, Springer, 4th, edition, 2000.
- [2] A. Baranne, D. Queloz, M. Mayor, G. Adrianzyk, G. Knispel, D. Kohler, D. Lacroix, J.-P. Meunier, G. Rimbaud, and A. Vin. ELODIE: A spectrograph for accurate radial velocity measurements. *A&AS*, 119:373–390, October 1996.
- [3] J. L. Bean, B. E. McArthur, G. F. Benedict, T. E. Harrison, D. Bizyaev, E. Nelan, and V. V. Smith. The Mass of the Candidate Exoplanet Companion to HD 33636 from Hubble Space Telescope Astrometry and High-Precision Radial Velocities. *AJ*, 134:749–758, August 2007. arXiv:0705.1861.
- [4] G. F. Benedict, B. E. McArthur, T. Forveille, X. Delfosse, E. Nelan, R. P. Butler, W. Spiesman, G. Marcy, B. Goldman, C. Perrier, W. H. Jefferys, and M. Mayor. A Mass for the Extrasolar Planet Gliese 876b Determined from Hubble Space Telescope Fine Guidance Sensor 3 Astrometry and High-Precision Radial Velocities. *ApJ*, 581:L115–L118, December 2002. arXiv:astro-ph/0212101.
- [5] G. F. Benedict, B. E. McArthur, G. Gatewood, E. Nelan, W. D. Cochran, A. Hatzes, M. Endl, R. Wittenmyer, S. L. Baliunas, G. A. H. Walker, S. Yang, M. Kürster, S. Els, and D. B. Paulson. The Extrasolar Planet ϵ Eridani b: Orbit and Mass. *AJ*, 132:2206–2218, November 2006. arXiv:astro-ph/0610247.

- [6] W. Benz. Low Velocity Collisions and the Growth of Planetesimals. *Space Science Reviews*, 92:279–294, April 2000.
- [7] A. C. Boley, R. H. Durisen, Å. Nordlund, and J. Lord. Three-Dimensional Radiative Hydrodynamics for Disk Stability Simulations: A Proposed Testing Standard and New Results. *ApJ*, 665:1254–1267, August 2007. arXiv:0704.2532.
- [8] I. A. Bond, A. Udalski, M. Jaroszyński, N. J. Rattenbury, B. Paczyński, I. Soszyński, L. Wyrzykowski, M. K. Szymański, M. Kubiak, O. Szewczyk, K. Żebruń, G. Pietrzyński, F. Abe, D. P. Bennett, S. Eguchi, Y. Furuta, J. B. Hearnshaw, K. Kamiya, P. M. Kilmartin, Y. Kurata, K. Masuda, Y. Matsubara, Y. Muraki, S. Noda, K. Okajima, T. Sako, T. Sekiguchi, D. J. Sullivan, T. Sumi, P. J. Tristram, T. Yanagisawa, and P. C. M. Yock. OGLE 2003-BLG-235/MOA 2003-BLG-53: A Planetary Microlensing Event. *ApJ*, 606:L155–L158, May 2004. arXiv:astro-ph/0404309.
- [9] X. Bonfils, X. Delfosse, S. Udry, N. C. Santos, T. Forveille, and D. Ségransan. Metallicity of M dwarfs. In F. Casoli, T. Contini, J. M. Hameury, and L. Pagani, editors, *SF2A-2005: Semaine de l’Astrophysique Française*, pages 327–+, December 2005.
- [10] W. J. Borucki, D. Koch, G. Basri, T. Brown, D. Caldwell, E. Devore, E. Dunham, T. Gautier, J. Geary, R. Gilliland, A. Gould, S. Howell, and J. Jenkins. Kepler Mission: a mission to find Earth-size planets in the habitable zone. In M. Fridlund, T. Henning, and H. Lacoste, editors, *Earths: DARWIN/TPF and the Search for Extrasolar Terrestrial Planets*, volume 539 of *ESA Special Publication*, pages 69–81, October 2003.
- [11] A. P. Boss, R. P. Butler, W. B. Hubbard, P. A. Ianna, M. Kürster, J. J. Lissauer, M. Mayor, K. J. Meech, F. Mignard, A. J. Penny, A. Quirrenbach, J. C. Tarter, and A. Vidal-Madjar. Working Group on Extrasolar Planets. *Transactions of the International Astronomical Union, Series A*, 26:183–186, March 2007.

- [12] R. P. Butler, G. W. Marcy, S. S. Vogt, and K. Apps. A Planet with a 3.1 Day Period around a Solar Twin. *PASP*, 110:1389–1393, December 1998.
- [13] R. P. Butler, G. W. Marcy, E. Williams, C. McCarthy, P. Dosanjh, and S. S. Vogt. Attaining Doppler Precision of 3 M s⁻¹. *PASP*, 108:500–+, June 1996.
- [14] R. P. Butler, S. S. Vogt, G. W. Marcy, D. A. Fischer, J. T. Wright, G. W. Henry, G. Laughlin, and J. J. Lissauer. A Neptune-Mass Planet Orbiting the Nearby M Dwarf GJ 436. *ApJ*, 617:580–588, December 2004. arXiv:astro-ph/0408587.
- [15] R. P. Butler, J. T. Wright, G. W. Marcy, D. A. Fischer, S. S. Vogt, C. G. Tinney, H. R. A. Jones, B. D. Carter, J. A. Johnson, C. McCarthy, and A. J. Penny. Catalog of Nearby Exoplanets. *ApJ*, 646:505–522, July 2006. arXiv:astro-ph/0607493.
- [16] D. Charbonneau, L. E. Allen, S. T. Megeath, G. Torres, R. Alonso, T. M. Brown, R. L. Gilliland, D. W. Latham, G. Mandushev, F. T. O’Donovan, and A. Sozzetti. Detection of Thermal Emission from an Extrasolar Planet. *ApJ*, 626:523–529, June 2005. arXiv:astro-ph/0503457.
- [17] P. Charbonneau. Genetic Algorithms in Astronomy and Astrophysics. *ApJS*, 101:309–+, December 1995.
- [18] G. Chauvin, A.-M. Lagrange, C. Dumas, B. Zuckerman, D. Mouillet, I. Song, J.-L. Beuzit, and P. Lowrance. A giant planet candidate near a young brown dwarf. Direct VLT/NACO observations using IR wavefront sensing. *A&A*, 425:L29–L32, October 2004. arXiv:astro-ph/0409323.
- [19] A. Cumming. Detectability of extrasolar planets in radial velocity surveys. *MNRAS*, 354:1165–1176, November 2004. arXiv:astro-ph/0408470.
- [20] A. Cumming, R. P. Butler, G. W. Marcy, S. S. Vogt, J. T. Wright, and D. A. Fischer. The Keck Planet Search: Detectability and the Mini-

- mum Mass and Orbital Period Distribution of Extrasolar Planets. *PASP*, 120:531–554, May 2008. arXiv:0803.3357.
- [21] D. Deming, S. Seager, L. J. Richardson, and J. Harrington. Infrared radiation from an extrasolar planet. *Nature*, 434:740–743, March 2005. arXiv:astro-ph/0503554.
- [22] C. Dominik, J. Blum, J. N. Cuzzi, and G. Wurm. Growth of Dust as the Initial Step Toward Planet Formation. In B. Reipurth, D. Jewitt, and K. Keil, editors, *Protostars and Planets V*, pages 783–800, 2007.
- [23] R. H. Durisen, A. P. Boss, L. Mayer, A. F. Nelson, T. Quinn, and W. K. M. Rice. Gravitational Instabilities in Gaseous Protoplanetary Disks and Implications for Giant Planet Formation. In B. Reipurth, D. Jewitt, and K. Keil, editors, *Protostars and Planets V*, pages 607–622, 2007.
- [24] D. A. Fischer, G. W. Marcy, R. P. Butler, S. S. Vogt, G. Laughlin, G. W. Henry, D. Abouav, K. M. G. Peek, J. T. Wright, J. A. Johnson, C. McCarthy, and H. Isaacson. Five Planets Orbiting 55 Cancri. *ApJ*, 675:790–801, March 2008. arXiv:0712.3917.
- [25] D. A. Fischer and J. Valenti. The Planet-Metallicity Correlation. *ApJ*, 622:1102–1117, April 2005.
- [26] E. B. Ford. Planet-Planet Scattering and Outer Planetesimal Disks. In *Bulletin of the American Astronomical Society*, volume 36 of *Bulletin of the American Astronomical Society*, pages 861–+, May 2004.
- [27] E. B. Ford. Quantifying the Uncertainty in the Orbits of Extrasolar Planets. *AJ*, 129:1706–1717, March 2005. arXiv:astro-ph/0305441.
- [28] E. B. Ford. Improving the Efficiency of Markov Chain Monte Carlo for Analyzing the Orbits of Extrasolar Planets. *ApJ*, 642:505–522, May 2006. arXiv:astro-ph/0512634.

- [29] E. B. Ford. Adaptive Scheduling Algorithms for Planet Searches. *AJ*, 135:1008–1020, March 2008. arXiv:astro-ph/0412703.
- [30] E. B. Ford, V. Lystad, and F. A. Rasio. Planet-planet scattering in the up-silon Andromedae system. *Nature*, 434:873–876, April 2005. arXiv:astro-ph/0502441.
- [31] W. R. Gilks. *Markov Chain Monte Carlo in Practice*. Chapman & Hall/CRC, December 1995.
- [32] G. Gonzalez. Parent stars of extrasolar planets - IX. Lithium abundances. *MNRAS*, 386:928–934, May 2008. arXiv:0802.0434.
- [33] A. Gonzalez-Villanueva, H. N. Nunez-Yepez, and A. L. Salas-Brito. In velocity space the Kepler orbits are circular. *European Journal of Physics*, 17:168–171, July 1996.
- [34] P. C. Gregory. A Bayesian Analysis of Extrasolar Planet Data for HD 73526. *ApJ*, 631:1198–1214, October 2005. arXiv:astro-ph/0509412.
- [35] P. C. Gregory. A Bayesian Kepler periodogram detects a second planet in HD208487. *MNRAS*, 374:1321–1333, February 2007. arXiv:astro-ph/0609229.
- [36] P. C. Gregory. A Bayesian periodogram finds evidence for three planets in HD 11964. *MNRAS*, 381:1607–1616, November 2007. arXiv:0709.0970.
- [37] P. C. Gregory and E. B. Ford. Bayesian model selection and extrasolar planet detection. In *Statistical Challenges in Modern Astronomy IV*, volume 371, 2007. arXiv:astro-ph/0608328v1.
- [38] C. J. Grillmair, D. Charbonneau, A. Burrows, L. Armus, J. Stauffer, V. Meadows, J. Van Cleve, and D. Levine. A Spitzer Spectrum of the Exoplanet HD 189733b. *ApJ*, 658:L115–L118, April 2007. arXiv:astro-ph/0702494.

- [39] J. L. Halbwachs, M. Mayor, and S. Udry. Statistical properties of exoplanets. IV. The period-eccentricity relations of exoplanets and of binary stars. *A&A*, 431:1129–1137, March 2005. arXiv:astro-ph/0410732.
- [40] L. Hartmann. *Accretion Processes in Star Formation*. Accretion processes in star formation / Lee Hartmann. Cambridge, UK ; New York : Cambridge University Press, 1998. (Cambridge astrophysics series ; 32) ISBN 0521435072., June 1998.
- [41] S. Ida, G. Bryden, D. N. C. Lin, and H. Tanaka. Orbital Migration of Neptune and Orbital Distribution of Trans-Neptunian Objects. , 534:428–445, May 2000.
- [42] S. Ida, Y. Nakagawa, and K. Nakazawa. The earth’s core formation due to the Rayleigh-Taylor instability. *Icarus*, 69:239–248, February 1987.
- [43] S. Inaba and M. Ikoma. Enhanced collisional growth of a protoplanet that has an atmosphere. *A&A*, 410:711–723, November 2003.
- [44] A. Johansen, J. S. Oishi, M.-M. M. Low, H. Klahr, T. Henning, and A. Youdin. Rapid planetesimal formation in turbulent circumstellar disks. *Nature*, 448:1022–1025, August 2007.
- [45] J. A. Johnson, R. P. Butler, G. W. Marcy, D. A. Fischer, S. S. Vogt, J. T. Wright, and K. M. G. Peek. A New Planet around an M Dwarf: Revealing a Correlation between Exoplanets and Stellar Mass. *ApJ*, 670:833–840, November 2007. arXiv:0707.2409.
- [46] V. Kashyap and J. J. Drake. Markov-Chain Monte Carlo Reconstruction of Emission Measure Distributions: Application to Solar Extreme-Ultraviolet Spectra. *ApJ*, 503:450–+, August 1998.
- [47] E. Kokubo and S. Ida. Oligarchic Growth of Protoplanets. *Icarus*, 131:171–178, January 1998.

- [48] S. J. Kortenkamp, E. Kokubo, and S. J. Weidenschilling. *Formation of Planetary Embryos*, pages 85–100. Origin of the earth and moon, edited by R.M. Canup and K. Righter and 69 collaborating authors. Tucson: University of Arizona Press., p.85-100, 2000.
- [49] D. Lafrenière, R. Doyon, C. Marois, D. Nadeau, B. R. Oppenheimer, P. F. Roche, F. Rigaut, J. R. Graham, R. Jayawardhana, D. Johnstone, P. G. Kalas, B. Macintosh, and R. Racine. The Gemini Deep Planet Survey. *ApJ*, 670:1367–1390, December 2007. arXiv:0705.4290.
- [50] G. Laughlin and F. C. Adams. Possible Stellar Metallicity Enhancements from the Accretion of Planets. *ApJ*, 491:L51+, December 1997. arXiv:astro-ph/9710110.
- [51] G. Laughlin and J. E. Chambers. Short-Term Dynamical Interactions among Extrasolar Planets. *ApJ*, 551:L109–L113, April 2001. arXiv:astro-ph/0101423.
- [52] J. J. Lissauer, J. B. Pollack, G. W. Wetherill, and D. J. Stevenson. Formation of the Neptune system. In *Neptune and Triton*, pages 37–108, 1995.
- [53] N. R. Lomb. Least-squares frequency analysis of unequally spaced data. *Ap&SS*, 39:447–462, February 1976.
- [54] R. E. Luck and U. Heiter. Dwarfs in the Local Region. *AJ*, 131:3069–3092, June 2006.
- [55] S. Mahadevan, J. Ge, J. van Eyken, B. Lee, S. Kane, P. Guo, S. W. Fleming, X. Wan, B. Zhao, J. R. Crepp, R. Cohen, S. Observers, and Staff. MARVELS: The Multi-Object APO Radial Velocity Survey. First Results from the Pilot Program. In *AAS/Division for Planetary Sciences Meeting Abstracts*, volume 39 of *AAS/Division for Planetary Sciences Meeting Abstracts*, pages 29.02–+, October 2007.

- [56] P. J. Marshall, M. P. Hobson, and A. Slosar. Bayesian joint analysis of cluster weak lensing and Sunyaev-Zel'dovich effect data. *MNRAS*, 346:489–500, December 2003. arXiv:astro-ph/0307098.
- [57] M. Mayor and D. Queloz. A Jupiter-Mass Companion to a Solar-Type Star. *Nature*, 378:355–+, November 1995.
- [58] B. E. McArthur, M. Endl, W. D. Cochran, G. F. Benedict, D. A. Fischer, G. W. Marcy, R. P. Butler, D. Naef, M. Mayor, D. Queloz, S. Udry, and T. E. Harrison. Detection of a Neptune-Mass Planet in the ρ^1 Cancri System Using the Hobby – Eberly Telescope. *ApJ*, 614 : L81 – –L84, October 2004. arXiv : astro – ph/0408585.
- [59] C. Moutou, F. Pont, and J.-L. Halbwachs. Detection and characterization of extrasolar planets: the transit method. *Formation planétaire et exoplanètes, Ecole thématique du CNRS, Goutelas (Loire), 23 - 27 mai 2005 Edited by J.-L. Halbwachs, D. Egret, and J.-M. Hameury. Strasbourg: Observatoire astronomique de Strasbourg et Société Française d'Astronomie et d'Astrophysique (SF2A), 2006, p. 55-79 <http://astro.u-strasbg.fr/goutelas/g2005/>, 28:55–79, March 2006.*
- [60] B. Nordström, M. Mayor, J. Andersen, J. Holmberg, F. Pont, B. R. Jørgensen, E. H. Olsen, S. Udry, and N. Mowlavi. The Geneva-Copenhagen survey of the Solar neighbourhood. Ages, metallicities, and kinematic properties of $\sim 14\,000$ F and G dwarfs. *A&A*, 418:989–1019, May 2004. arXiv:astro-ph/0405198.
- [61] B. Panter, A. F. Heavens, and R. Jimenez. Star formation and metallicity history of the SDSS galaxy survey: unlocking the fossil record. *MNRAS*, 343:1145–1154, August 2003. arXiv:astro-ph/0211546.
- [62] F. Pepe, M. Mayor, D. Queloz, W. Benz, X. Bonfils, F. Bouchy, G. L. Curto, C. Lovis, D. Mégevand, C. Moutou, D. Naef, G. Rupprecht, N. C. Santos, J.-P. Sivan, D. Sosnowska, and S. Udry. The HARPS search for

- southern extra-solar planets. I. HD 330075 b: A new “hot Jupiter”. *A&A*, 423:385–389, August 2004. arXiv:astro-ph/0405252.
- [63] C. Perrier, J.-P. Sivan, D. Naef, J. L. Beuzit, M. Mayor, D. Queloz, and S. Udry. The ELODIE survey for northern extra-solar planets. I. Six new extra-solar planet candidates. *A&A*, 410:1039–1049, November 2003. arXiv:astro-ph/0308281.
- [64] J. B. Pollack, O. Hubickyj, P. Bodenheimer, J. J. Lissauer, M. Podolak, and Y. Greenzweig. Formation of the Giant Planets by Concurrent Accretion of Solids and Gas. *Icarus*, 124:62–85, November 1996.
- [65] William Press, Saul Teukolsky, William Vetterling, and Brian Flannery. *Numerical Recipes in C*. Cambridge University Press, Cambridge, UK, 2nd edition, 1992.
- [66] F. A. Rasio and E. B. Ford. Dynamical instabilities and the formation of extrasolar planetary systems. *Science*, 274:954–956, November 1996.
- [67] L. J. Richardson, D. Deming, K. Horning, S. Seager, and J. Harrington. A spectrum of an extrasolar planet. *Nature*, 445:892–895, February 2007. arXiv:astro-ph/0702507.
- [68] E. Sandquist, R. E. Taam, D. N. C. Lin, and A. Burkert. Planet Consumption and Stellar Metallicity Enhancements. *ApJ*, 506:L65–L68, October 1998. arXiv:astro-ph/9808128.
- [69] J. D. Scargle. Studies in astronomical time series analysis. II - Statistical aspects of spectral analysis of unevenly spaced data. *ApJ*, 263:835–853, December 1982.
- [70] R. Silvotti, S. Schuh, R. Janulis, J.-E. Solheim, S. Bernabei, R. Østensen, T. D. Oswalt, I. Bruni, R. Gualandi, A. Bonanno, G. Vauclair, M. Reed, C.-W. Chen, E. Leibowitz, M. Paparo, A. Baran, S. Charpinet, N. Dolez, S. Kawaler, D. Kurtz, P. Moskalik, R. Riddle, and S. Zola. A giant

- planet orbiting the ‘extreme horizontal branch’ star V391 Pegasi. *Nature*, 449:189–191, September 2007.
- [71] D. S. Sivia. *Data Analysis: A Bayesian Tutorial (Oxford Science Publications)*. Oxford University Press, July 1996.
- [72] T. F. Stepinski, R. Malhotra, and D. C. Black. The Andromedae System: Models and Stability. *ApJ*, 545:1044–1057, December 2000. arXiv:astro-ph/0010279.
- [73] G. Takeda, E. B. Ford, A. Sills, F. A. Rasio, D. A. Fischer, and J. A. Valenti. Structure and Evolution of Nearby Stars with Planets. II. Physical Properties of ~ 1000 Cool Stars from the SPOCS Catalog. *ApJS*, 168:297–318, February 2007. arXiv:astro-ph/0607235.
- [74] H. Tanaka, T. Takeuchi, and W. R. Ward. Three-Dimensional Interaction between a Planet and an Isothermal Gaseous Disk. I. Corotation and Lindblad Torques and Planet Migration. *ApJ*, 565:1257–1274, February 2002.
- [75] S. E. Thorsett, Z. Arzoumanian, and J. H. Taylor. PSR B1620-26 - A binary radio pulsar with a planetary companion? *ApJ*, 412:L33–L36, July 1993.
- [76] J. A. Valenti and D. A. Fischer. Spectroscopic Properties of Cool Stars (SPOCS). I. 1040 F, G, and K Dwarfs from Keck, Lick, and AAT Planet Search Programs. *ApJS*, 159:141–166, July 2005.
- [77] L. Verde, H. V. Peiris, D. N. Spergel, M. R. Nolta, C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, G. S. Tucker, E. Wollack, and E. L. Wright. First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Parameter Estimation Methodology. *ApJS*, 148:195–211, September 2003. arXiv:astro-ph/0302218.

- [78] S. S. Vogt, G. W. Marcy, R. P. Butler, and K. Apps. Six New Planets from the Keck Precision Velocity Survey. *ApJ*, 536:902–914, June 2000. arXiv:astro-ph/9911506.
- [79] W. R. Ward. Protoplanet Migration by Nebula Tides. *Icarus*, 126:261–281, April 1997.
- [80] S. J. Weidenschilling. Aerodynamics of solid bodies in the solar nebula. *MNRAS*, 180:57–70, July 1977.
- [81] A. Wolszczan and D. A. Frail. A planetary system around the millisecond pulsar PSR1257 + 12. *Nature*, 355:145–147, January 1992.
- [82] A. N. Youdin and F. H. Shu. Planetesimal Formation by Gravitational Instability. *ApJ*, 580:494–505, November 2002. arXiv:astro-ph/0207536.
- [83] S. Zucker, T. Mazeh, N. C. Santos, S. Udry, and M. Mayor. Multi-order TODCOR: Application to observations taken with the CORALIE echelle spectrograph. II. A planet in the system HD 41004. *A&A*, 426:695–698, November 2004.