

An integrated analysis of radial velocities in planet searches

Andrew Cumming¹* and Diana Dragomir^{1,2}

¹*Department of Physics, McGill University, 3600 rue University, Montreal, QC H3A 2T8, Canada*

²*Department of Physics and Astronomy, University of British Columbia, 6224 Agricultural Road, Vancouver, BC V6T1Z1, Canada*

Accepted 2009 August 23. Received 2009 August 18; in original form 2009 January 18

ABSTRACT

We discuss a Bayesian approach to the analysis of radial velocities in planet searches. We use a combination of exact and approximate analytic and numerical techniques to efficiently evaluate χ^2 for multiple values of orbital parameters, and to carry out the marginalization integrals for a single planet including the possibility of a long-term trend. The result is a robust algorithm that is rapid enough for use in real-time analysis that outputs constraints on orbital parameters and false-alarm probabilities for the planet and long-term trend. The constraints on parameters and odds ratio that we derive compare well with previous calculations based on Markov Chain Monte Carlo methods, and we compare our results with other techniques for estimating false-alarm probabilities and errors in derived orbital parameters. False-alarm probabilities from the Bayesian analysis are systematically higher than frequentist false-alarm probabilities, due to the different accounting of the number of trials. We show that upper limits on the velocity amplitude derived for circular orbits are a good estimate of the upper limit on the amplitude of eccentric orbits for $e \lesssim 0.5$.

Key words: methods: statistical – binaries: spectroscopic – planetary systems.

1 INTRODUCTION

The analysis of a set of radial velocities in planet searches typically involves a number of different steps. First, the best-fitting Keplerian orbital parameters are found by minimizing χ^2 , for example with a Levenberg–Marquardt algorithm (Press et al. 1992). Because of the complex multimodal shape of the χ^2 distribution in parameter space, a Lomb–Scargle periodogram (Lomb 1978; Scargle 1982) is often used beforehand to fit circular orbits at a range of orbital periods, providing starting points for the Keplerian fit. Then, the reality of the signal is assessed by calculating the false-alarm probability (FAP) that the observed signal could arise due to noise fluctuations, typically using Monte Carlo simulations (Marcy et al. 2005; Cumming 2004). This has become particularly important as radial velocity surveys reveal planets with lower velocity amplitudes, comparable to the measurement uncertainties and other sources of noise. A related question is comparing different models for the data, for example deciding whether a two (or more) planet model is preferred over a single-planet model, or whether to include a long-term trend due to a long-period companion (see e.g. Robinson et al. 2007).

Uncertainties in the fitted orbital parameters are then calculated. A common technique is to scramble the residuals to the best-fitting Keplerian orbit, add them back to the predicted velocity curve, and refit the orbit. After repeating this many times, the distribution of fitted parameters gives an estimate of the uncertainty (Marcy et al.

2005). Another approach is to use Bayesian methods implemented with Markov Chain Monte Carlo (MCMC) simulations (Ford 2005). In the case of a non-detection, the upper limit on the planet mass as a function of orbital period is an important input for population studies (Walker et al. 1995; Cumming, Marcy & Butler 1999; Endl et al. 2002; Wittenmyer et al. 2006).

Often, all of these steps must be carried out for a given radial velocity data set. Many of them are based on Monte Carlo simulations involving fitting Keplerian orbits to synthetic data sets. These trials can become cumbersome for the large numbers of orbital frequencies that must be considered. For this reason, recent calculations have produced upper limits for circular orbits only (Cumming et al. 2008), relying on the fact that detectability of planets falls off with eccentricity only for $e \gtrsim 0.6$ (Endl et al. 2002; Cumming 2004), or on a sparse grid of orbital period values (O’Toole et al. 2008).

We focus in this paper on a Bayesian approach to the analysis of radial velocity data. The advantage is that, in principle, a Bayesian analysis answers all of the above questions with a single calculation, providing constraints on model parameters and odds ratios which can be used to decide which model best describes the data (Cumming 2004; Ford 2005; Gregory 2005b). This would simplify analysis of radial velocity data sets. The difficulty in practice is that the marginalization over parameters requires the evaluation of multidimensional integrals over parameter space.

Bayesian methods have been applied to planet searches, using sophisticated MCMC techniques to evaluate the integrals. Ford (2005) applied this technique to determine the constraints on orbital parameters, while Gregory in a series of papers (Gregory 2005b, 2007a,b)

*E-mail: cumming@physics.mcgill.ca

also considers model comparison. Ford (2006) investigated different proposal distribution functions to help speed convergence of MCMC chains. Whereas the chains used by Ford (2005) were 10^6 – 10^{10} steps in length, Ford (2006) found that it was possible to achieve convergence after only 10^4 – 10^6 steps by optimizing the directions in parameter space in which steps are taken. Recently, Balan & Lahav (2009) have developed an MCMC code `EXOFIT` based on the methodology of Ford & Gregory (2007) which is publicly available.¹

Despite this tremendous progress, the application of MCMC to radial velocity data is not yet routine, although it is commonly used to assess uncertainties in orbital parameters. As well as optimizing the steps in parameter space, one important difficulty in using the MCMC approach is assessing whether the chains have converged. Another is that when the signal-to-noise ratio is low and the distribution of χ^2 in parameter space is multimodal, the MCMC chain may miss minima in χ^2 . Gregory (2005b) introduced a parallel tempering scheme in which several MCMC chains are run simultaneously, each with a different temperature, hotter chains making larger jumps in parameter space and colder chains exploring local minima. As the calculation progresses, the chains exchange information in a way that preserves their statistical character. This scheme has been successfully applied to multiple planet systems (Gregory 2007a,b).

In this paper, we take a different approach. We consider models with one planet only, or one planet plus a long-term linear trend, and use a combination of grid-based numerical evaluation and exact and approximate analytic methods to evaluate the marginalization integrals. The idea is to look for ways in which the marginalization integrals can be evaluated more efficiently. As well as providing a useful tool for analysing radial velocity data for single-planet systems, it provides a check on the output of MCMC simulations, and may have application to making MCMC codes for the analysis of multiple-planet systems more efficient.

We start in Section 2 with an overview of the Bayesian approach, including how to write down the posterior probabilities for orbital parameters and how to use them to calculate FAPs. In Section 3, we discuss circular orbits, using analytic techniques to evaluate the marginalization integrals. In Section 4, we divide the parameters for eccentric orbits into fast (linear) and slow (non-linear) parameters, and use the analytic techniques for circular orbits to marginalize over the fast parameters. In Section 5, we compare our results to MCMC calculations, and traditional methods for evaluating FAPs and upper limits on companion mass.

2 OVERVIEW

The Bayesian analysis of radial velocity data has been discussed previously by several authors (Cumming 2004; Ford 2005; Gregory 2005a,b; Ford 2006; Ford & Gregory 2007; Gregory 2007a,b; Ford 2008). Here, we give a brief reminder of the basic ideas and introduce our notation, and show how the systemic velocity and noise uncertainty can be analytically marginalized.

2.1 Parameter estimation

We start with a model for the radial velocities with set of parameters \mathbf{a} . For example, a single Keplerian orbit has six parameters $\mathbf{a} = (K, P, e, \omega_p, t_p, \gamma)$, where K is the velocity amplitude, P is the orbital period, e is the eccentricity, ω_p and t_p are the argument and time of

pericentre and γ is the systemic velocity. The data consist of a set of N measured velocities v_i , observation times t_i and errors σ_i . Bayes' theorem allows us to calculate the probability distribution of the parameters \mathbf{a} given the data, also known as the posterior probability of \mathbf{a} ,

$$\mathcal{P}(\mathbf{a}|d) = \frac{\mathcal{P}(\mathbf{a})\mathcal{P}(d|\mathbf{a})}{\mathcal{P}(d)}, \quad (1)$$

where $\mathcal{P}(d)$ is a normalization factor. The term $\mathcal{P}(\mathbf{a})$ is the prior probability distribution for the parameters \mathbf{a} , which allows us to specify any knowledge of the parameter distribution that we have before the data are taken. If the errors are Gaussian distributed and uncorrelated, the likelihood of the data or probability of the data given a particular choice of model parameters is

$$\mathcal{P}(d|\mathbf{a}) = \frac{1}{\prod_i (2\pi)^{1/2} \sigma_i} \exp\left[-\frac{\chi^2(\mathbf{a})}{2}\right], \quad (2)$$

where

$$\chi^2(\mathbf{a}) = \sum_{i=1}^N w_i [v_i - V_i(\mathbf{a})]^2 \quad (3)$$

is the usual χ^2 statistic, written in terms of weights $w_i = 1/\sigma_i^2$. We write the model velocity at time t_i as V_i .

Often, we are interested in the probability distribution of a single parameter, or a subset of parameters. For example, a circular orbit has $\mathbf{a} = (K, P, \gamma, \phi)$, where K is the velocity amplitude, P is the orbital period, γ is the systemic velocity and ϕ is the orbital phase. It is likely that we are not interested in the particular values of γ or ϕ , but want to constrain the orbital period and velocity amplitude. The joint probability distribution for P and K can be obtained by marginalizing over the other parameters,

$$\mathcal{P}(P, K|d) = \int d\phi \int d\gamma \mathcal{P}(P, K, \phi, \gamma|d). \quad (4)$$

Marginalization amounts to performing a weighted average of the probability distribution over the unwanted parameters.

A number of other useful quantities can be obtained from $\mathcal{P}(P, K|d)$. Further integration over K gives $\mathcal{P}(P|d)$, or integration over P gives $\mathcal{P}(K|d)$. A confidence interval for K can be calculated from $\mathcal{P}(K|d)$. For example, if a planet is not detected in a given data set, an upper limit can be placed on the amplitude of undetected orbits. The 99 per cent upper limit K_{99} is given by $\int_0^{K_{99}} dK \mathcal{P}(K|d) / \int_0^\infty dK \mathcal{P}(K|d) = 0.99$.

For eccentric orbits, we will focus in this paper on obtaining $\mathcal{P}(P, e, K|d)$ by marginalizing over γ, ω_p and t_p .

2.2 The noise distribution

In equation (2), we assumed that the standard deviation of the noise for each observation σ_i was given. In reality, other noise sources may be present in the data that hinder the identification of planetary signals, for example intrinsic stellar ‘jitter’ (e.g. Wright 2005) due to rotation of spots across the surface of the star or changes in line profiles over time related to magnetic activity. This extra noise can be incorporated as an additional parameter of the model. A common choice (Gregory 2005b; Ford 2006) is to add the extra noise term in quadrature with the measurement errors σ_i .

Here, we instead multiply each value of σ_i by a noise scaling factor k (Cumming 2004; Gregory 2005a), and analytically marginalize over k (e.g. Sivia 1996)

$$\mathcal{P}(d|\mathbf{a}) \propto \int_0^\infty \frac{dk}{k} \frac{1}{k^N} \exp\left[-\frac{\chi^2(\mathbf{a})}{2k^2}\right] \propto [\chi^2(\mathbf{a})]^{-N/2}. \quad (5)$$

¹ Available at <http://www.http://zuserver2.star.ucl.ac.uk/lahav/exofit.html>.

The constant prefactor, which depends only on the weights w_i and the number of observations N , does not affect the shape of the posterior probability distributions, and cancels out when we calculate odds ratios. Therefore, we drop it and replace equation (2) with

$$\mathcal{P}(d|\mathbf{a}) = [\chi^2(\mathbf{a})]^{-N/2}. \quad (6)$$

This is a Student's t -distribution rather than Gaussian distribution (Sivia 1996).

We take an infinite range for k in equation (5), whereas in reality we likely have some information about the uncertainty in the noise level. For example, we may be able to estimate the size of the expected stellar jitter based on stellar properties (Wright 2005). Alternatively, we could keep k as a parameter, and evaluate the constraints on k from the data, $\mathcal{P}(k|d)$ (e.g. Ford 2006). We have tried marginalizing numerically over k with finite limits, and find that the results for realistic ranges of k are close to the analytic case with infinite limits. Therefore, we marginalize analytically over k and adopt equation (6) as the likelihood throughout this paper.²

2.3 Priors

The choice of appropriate prior probabilities $\mathcal{P}(\mathbf{a})$ for the various parameters has been discussed in depth in the literature (e.g. Gregory 2005b; Ford & Gregory 2007). We mostly follow this previous work. For circular orbits, we use uniform priors for γ and ϕ , and priors for K , P that are uniform in log (the Jeffreys prior). For eccentric orbits, we take uniform priors in γ , t_p , ω_p , e , and log-uniform priors in K , and P , i.e. $\mathcal{P}(P) = 1/[P \log(P_2/P_1)]$ and $\mathcal{P}(K) = 1/[K \log(K_2/K_1)]$. If a long-term trend is included in the model (a linear term βt_i ; see Appendix A), we take a uniform prior in the slope β .

In fact, Gregory (2005b) and Ford & Gregory (2007) use a modified Jeffreys prior for the noise term and the velocity amplitude rather than the standard Jeffreys prior. A modified Jeffreys prior is uniform in log above some scale, and uniform below that scale. For radial velocity amplitude K or extra noise term, the turnover scale is taken to be $\approx 1 \text{ m s}^{-1}$. The values of K we are interested in are typically larger than this, and so for simplicity we use a Jeffreys prior between our lower and upper limits in K . The ranges that we take are $K = 1 \text{ m s}^{-1}$ to $2\Delta v$, where Δv is the observed range of velocities, and $P = 1$ day to the time-span of the observations.

2.4 Model comparison and the false-alarm probability

Marginalizing over all the parameters of a model gives the total probability of that model. For example, given $\mathcal{P}(P, K|d)$ for circular orbits, we could calculate the total probability that a planet is present:

$$\mathcal{P}(1|d) = \int dP \int dK \mathcal{P}(P, K|d). \quad (7)$$

Similarly, by considering a model without a planet, we can calculate the probability that no planet is present given the data, $\mathcal{P}(0|d)$. We define the normalization $\mathcal{P}(d)$ in equation (1) so that the sum over the probabilities of all models is unity. For example, if we consider only two possible models, that there is or is not a planet present, we choose $\mathcal{P}(d)$ such that

$$\mathcal{P}(1|d) + \mathcal{P}(0|d) = 1. \quad (8)$$

² The techniques we develop below for rapidly marginalizing over parameters can also be applied to the case where k is kept as a parameter. This is discussed in Appendix B.

We can think of the posterior probability that there is no planet present $\mathcal{P}(0|d)$ as the FAP. It can be written without including the $\mathcal{P}(d)$ factors explicitly as

$$F = \mathcal{P}(0|d) = \frac{1}{1 + \Lambda}, \quad (9)$$

where Λ is the odds ratio

$$\Lambda = \frac{\mathcal{P}(1|d)}{\mathcal{P}(0|d)} \quad (10)$$

(the normalization factors $\mathcal{P}(d)$ cancel out when the ratio is taken). For $\Lambda \gg 1$, $F \approx \Lambda^{-1}$. The odds ratio is

$$\Lambda = \frac{\int dK \int dP \mathcal{P}(P, K|d)}{\mathcal{P}(0|d)} \quad (11)$$

for circular orbits, or

$$\Lambda = \frac{\int dK \int dP \int d\epsilon \mathcal{P}(P, K, \epsilon|d)}{\mathcal{P}(0|d)} \quad (12)$$

for eccentric orbits.

This approach can be generalized to more than two models. For example, later we will consider four possible models for a given star, the possible combinations of including or not including a Keplerian orbit with period less than the time-span of the observations, and including or not including a long-term trend. To calculate the FAP associated with the short-period planet, we define the odds ratio:

$$\Lambda = \frac{\mathcal{P}(1|d) + \mathcal{P}(1, t|d)}{\mathcal{P}(0|d) + \mathcal{P}(0, t|d)}, \quad (13)$$

where 1 or 0 indicate that the short-period planet is or is not included in the model, and t indicates that a long-term trend is included.

2.5 Probability that there is no planet $\mathcal{P}(0|d)$

The posterior probability of no planet $\mathcal{P}(0|d)$, where the ‘no-planet’ model is a constant velocity $V_i = \gamma$, can be calculated analytically. Using the likelihood of equation (6)

$$\mathcal{P}(0|d) = \frac{1}{\mathcal{P}(d)\Delta\gamma} \int_{\gamma_1}^{\gamma_2} d\gamma [\chi^2(\gamma)]^{-N/2}, \quad (14)$$

where $\Delta\gamma = \gamma_2 - \gamma_1$ is the range of values of γ considered, and we assume a uniform prior for γ in that range. Minimizing χ^2 with respect to γ , we find the best-fitting value $\gamma_0 = \sum w_i v_i / \sum w_i$. In terms of γ_0 , we can write

$$\chi^2(\gamma) = \chi^2(\gamma_0) + (\gamma - \gamma_0)^2 \sum w_i. \quad (15)$$

The fact that the distribution of $\chi^2(\gamma)$ is analytic is mentioned in Ford (2006). For clarity, we drop the subscript i on the sum in equation (15) and in the remainder of the paper, a sum over the observations with i running from 1 to N is implied.

The quadratic form of χ^2 allows the integral over γ to be carried out analytically when the limits $\gamma_1 \rightarrow -\infty$ and $\gamma_2 \rightarrow \infty$. In that limit, the normalization factor diverges, $\Delta\gamma \rightarrow \infty$. However, the values of $\Delta\gamma$ cancel when we form an odds ratio, as does the normalization factor $\mathcal{P}(d)$. Therefore, we can drop the prefactor after integrating, giving the final result

$$\mathcal{P}(0|d) = [\chi^2(\gamma_0)]^{-(N-1)/2}. \quad (16)$$

An alternative ‘no-planet’ model is a linear trend in the radial velocities over time, $V_i = \gamma + \beta t_i$. A linear term is often included (and needed) in radial velocity fits to account for additional companions with long orbital periods. A similar formula for $\mathcal{P}(0|d)$ can be derived in that case. For clarity, we leave this to Appendix A, along with how to add a linear term to the circular and Keplerian orbit fits, and consider only the constant-velocity no-planet model in the main text.

3 CIRCULAR ORBITS

In the previous section, we saw that a calculation of $\mathcal{P}(\mathbf{a}|d)$ followed by successive marginalization provides constraints on all model parameters and a measure of the FAP. The difficulty in practice is in performing the integrals over parameter space. We first consider circular orbits, which have a simple sinusoidal velocity curve, and introduce some analytic approximations that allow us to rapidly carry out these integrals. Apart from being a testing ground for these techniques which we will then apply to eccentric orbits, fitting circular orbits is actually quite useful since sinusoid fits are sufficient to detect orbits even with moderate eccentricities ($e \lesssim 0.5$; Endl et al. 2002; Cumming 2004).

For a circular orbit, the model for the velocities is

$$V_i = \gamma + K \sin(\omega t_i + \phi) \quad (17)$$

which has four parameters: γ is the systemic velocity; K is the velocity semi-amplitude; ϕ is the phase and $\omega = 2\pi/P$ is the orbital frequency, P is the orbital period. Our aim in this section is to obtain $\mathcal{P}(P, K)$, marginalizing over γ and ϕ . We first marginalize analytically over γ to obtain $\mathcal{P}(\phi, K, P|d)$, and then present two different methods for efficiently marginalizing over the parameters K and ϕ . The methods are summarized and applied to an example data set in Section 3.5.

3.1 Analytic marginalization of the systemic velocity

To integrate over γ , we note again that χ^2 depends quadratically on γ around the best-fitting value, as given by equation (15), where this time $\gamma_0(\phi, K, P)$ is the best-fitting systemic velocity at each ϕ, P and K , i.e. $\gamma_0(\phi, K, P)$ is the value of γ that minimizes χ^2 at each ϕ, P and K , and $\chi^2(\gamma_0)$ is the corresponding minimum value of χ^2 . The best-fitting systemic velocity can be calculated from $\partial\chi^2/\partial\gamma = 0$, giving

$$\gamma_0 = \sum w_i [v_i - K \sin(\omega t_i + \phi)] / \sum w_i. \quad (18)$$

Adopting a uniform prior for γ and integrating for $\Delta\gamma \rightarrow \infty$, we find

$$\mathcal{P}(d|\phi, K, P) = (\chi^2[\gamma_0, \phi, K, P])^{-(N-1)/2}, \quad (19)$$

where $\gamma_0(\phi, K, P)$ is given by equation (18), and we have set the prefactor equal to unity as in Section 2.5.

3.2 Evaluation of $\mathcal{P}(\phi, K, P|d)$ on a grid

Next, we describe a method for rapidly evaluating $\mathcal{P}(\phi, K, P|d)$ numerically for a grid of values of ϕ, K and P . We introduce the averages

$$\begin{aligned} \langle v \rangle &= \sum w_i v_i / \sum w_i \\ \langle C \rangle &= \sum w_i \cos(\omega t_i) / \sum w_i \\ \langle S \rangle &= \sum w_i \sin(\omega t_i) / \sum w_i \\ \langle vC \rangle &= \sum w_i v_i \cos(\omega t_i) / \sum w_i \\ \langle vS \rangle &= \sum w_i v_i \sin(\omega t_i) / \sum w_i \\ \langle C^2 \rangle &= \sum w_i \cos^2(\omega t_i) / \sum w_i \\ \langle S^2 \rangle &= \sum w_i \sin^2(\omega t_i) / \sum w_i \\ \langle SC \rangle &= \sum w_i \cos(\omega t_i) \sin(\omega t_i) / \sum w_i. \end{aligned} \quad (20)$$

In this notation, equation (18) can be written as $\gamma_0 = \langle v \rangle - K \langle C \rangle \sin \phi - K \langle S \rangle \cos \phi$. Substituting this expression into χ^2 and simplifying, we find

$$\begin{aligned} \frac{\chi^2(\phi, K, P)}{\sum w_i} &= \langle \langle v^2 \rangle \rangle - 2K [\langle \langle vC \rangle \rangle \sin \phi + \langle \langle vS \rangle \rangle \cos \phi] \\ &\quad + K^2 [\langle \langle C^2 \rangle \rangle \sin^2 \phi + \langle \langle S^2 \rangle \rangle \cos^2 \phi \\ &\quad + 2\langle \langle SC \rangle \rangle \sin \phi \cos \phi], \end{aligned} \quad (21)$$

where $\langle \langle fg \rangle \rangle = \langle (f - \langle f \rangle)(g - \langle g \rangle) \rangle = \langle fg \rangle - \langle f \rangle \langle g \rangle$.

Equation (21) allows efficient calculation of χ^2 for multiple values of the parameters ϕ, K and P . Given three vectors – a vector of K values, a vector of ϕ values (and corresponding values of $\sin \phi$ and $\cos \phi$) and a vector of orbital periods and the corresponding averages over the data (terms in angle brackets) – a three-dimensional (3D) matrix of χ^2 values can be quickly generated. The advantage is that the sums over the data need to be calculated only once, rather than being re-evaluated for each new choice of K and ϕ .

Marginalizing over ϕ is then straightforward, since the integral

$$\begin{aligned} \mathcal{P}(d|K, P) &= \frac{1}{2\pi} \int_0^{2\pi} d\phi \mathcal{P}(d|\phi, K, P) \\ &= \frac{1}{2\pi} \int_0^{2\pi} d\phi [\chi^2(\phi, K, P)]^{-(N-1)/2} \end{aligned} \quad (22)$$

can be calculated using a quadrature method based on the values of ϕ in the grid. To calculate the odds ratio, we should compare this with the probability for a no-planet model, which has $V_i = \gamma$ only. In this case, $\gamma_0 = \langle v \rangle$ and $\chi_0^2 / \sum w_i = \langle \langle v^2 \rangle \rangle$, so that

$$\mathcal{P}(0|d) = \left(\langle \langle v^2 \rangle \rangle \sum w_i \right)^{-(N-1)/2}, \quad (23)$$

which can be used in equation (11) for Λ .

3.3 Analytic marginalization of ϕ and K

The reason that we could analytically integrate over γ is that the model V_i is linear in γ . Now in fact, we can perform a similar analytic integration over K and ϕ by rewriting equation (17) in terms of the linear parameters A and B ,

$$V_i = \gamma + A \sin \omega t_i + B \cos \omega t_i \quad (24)$$

where $A = K \cos \phi$ and $B = K \sin \phi$. In seminal papers on Bayesian signal detection, Bretthorst (1988) carried out analytic integration over A and B , and we follow the same approach here (see also Ford 2008).

To perform the integration, we use the fact that the quadratic shape of χ^2 that we found for γ (equation 15) generalizes to an arbitrary linear model $V_i = \sum_k a_k g_k(t_i)$. It is straightforward to show that³

$$\chi^2(\mathbf{a}) = \chi^2(\mathbf{a}_0) + \delta \mathbf{a} \cdot \boldsymbol{\alpha} \cdot \delta \mathbf{a}, \quad (25)$$

³ There is an approximation known as the Laplace approximation (Sivia 1996) in which the quadratic form in equation (25) is assumed close to the minimum χ^2 value. Ford (2008) applied this approximation to circular orbit fits at specified orbital periods, but in fact as we have noted here the approximation is exact in this case because the model is linear. We have tried applying the Laplace approximation to carry out the integral in ϕ in equation (22). However, we find that this approximation does not perform well at low K , where $\mathcal{P}(d|\phi)$ is bimodal, and in addition it is not convenient numerically as it requires a search for the peak in $\mathcal{P}(d|\phi)$ at each value of K .

where the matrix α is the inverse of the correlation matrix (Press et al. 1992), and has components $\alpha_{kl} = (1/2)(\partial^2 \chi^2 / \partial a_k \partial a_l) = \sum w_i g_k(t_i) g_l(t_i)$. The marginalization integral with uniform priors for the parameters can be done analytically⁴

$$\int d^m \mathbf{a} (\chi^2)^{-N/2} = \frac{(\chi_0^2)^{-\frac{N-m}{2}} \pi^{m/2} \Gamma(\frac{N-m}{2})}{\sqrt{\det \alpha} \Gamma(\frac{N}{2})}, \quad (26)$$

where m is the number of parameters integrated over. We use subscript zero to indicate the best-fitting value of parameters, or the corresponding minimum value of χ^2 .

Applying this result to the integration over A and B gives

$$\begin{aligned} \mathcal{P}(P|d) &= \frac{1}{P} \int \frac{dA dB d\gamma}{\Delta A \Delta B \Delta \gamma} [\chi^2(A, B, \gamma, P)]^{-N/2} \\ &= \frac{1}{P \Delta A \Delta B \Delta \gamma} \frac{(\chi_0^2)^{-\frac{N-3}{2}} \pi^{3/2} \Gamma(\frac{N-3}{2})}{\sqrt{\det \alpha} \Gamma(\frac{N}{2})}. \end{aligned} \quad (27)$$

The values of χ_0^2 and $\det \alpha$ can be calculated as a function of P as follows. First, by minimizing χ^2 with respect to A , B and γ , the best-fitting values of γ , $A = K \cos \phi$ and $B = K \sin \phi$ are

$$A_0 = \frac{\langle \langle vS \rangle \rangle \langle \langle C^2 \rangle \rangle - \langle \langle vC \rangle \rangle \langle \langle SC \rangle \rangle}{\langle \langle C^2 \rangle \rangle \langle \langle S^2 \rangle \rangle - \langle \langle SC \rangle \rangle^2} \quad (28)$$

$$B_0 = \frac{\langle \langle vC \rangle \rangle \langle \langle S^2 \rangle \rangle - \langle \langle vS \rangle \rangle \langle \langle SC \rangle \rangle}{\langle \langle C^2 \rangle \rangle \langle \langle S^2 \rangle \rangle - \langle \langle SC \rangle \rangle^2} \quad (29)$$

$$\gamma_0 = \langle v \rangle - A_0 \langle S \rangle - B_0 \langle C \rangle, \quad (30)$$

and the minimum value of χ^2 is

$$\begin{aligned} \frac{\chi_0^2(P)}{\sum w_i} &= \langle \langle v^2 \rangle \rangle - 2A_0 \langle \langle vS \rangle \rangle - 2B_0 \langle \langle vC \rangle \rangle \\ &\quad + A_0^2 \langle \langle S^2 \rangle \rangle + B_0^2 \langle \langle C^2 \rangle \rangle + 2A_0 B_0 \langle \langle SC \rangle \rangle \end{aligned} \quad (31)$$

and

$$\frac{\det \alpha}{(\sum w_i)^3} = \langle \langle S^2 \rangle \rangle \langle \langle C^2 \rangle \rangle - \langle \langle SC \rangle \rangle^2. \quad (32)$$

This allows us to easily calculate $\mathcal{P}(P|d)$.

The only remaining question is what to choose for the prior ranges ΔA and ΔB (the prior range in gamma $\Delta \gamma$ cancels when we form the odds ratio). Unfortunately, the analytic evaluation of the integral in equation (27) is only possible for a uniform prior in A and B . Since $dA dB = K dK d\phi$, a uniform prior in A and B corresponds to a prior $\mathcal{P}(K) \propto K$ rather than the logarithmic prior $\mathcal{P}(K) \propto 1/K$ that we assumed in the grid-based calculation (see the discussion in Bretthorst 1988 who chose a different prior to Jaynes 1987). Therefore, the analytic marginalization gives more weight to large- K solutions, whereas the grid-based approach gives more weight to small- K solutions. We correct for this in an approximate way by choosing the normalization appropriately. We find that the choice

$$\Delta A \Delta B = K_0(P) K_{0,av} \log(K_2/K_1), \quad (33)$$

where $K_{0,av}$ is the best-fitting velocity amplitude averaged over all frequencies reproduces the normalization of the grid-based calculation, with final odds ratios typically within a factor of 2.

⁴ To prove equation (26), follow the method given in the Appendix of Sivia (1996), where a similar result is derived for a likelihood $\propto \exp(-\chi^2/2)$.

3.4 The probability distribution of K at each orbital period

Analytical marginalization over the linear parameters A and B is convenient, but in doing so we have thrown away information about the velocity amplitude K . It turns out that we can get it back very easily using an analytic approximation for the shape of $\mathcal{P}(d|K)$ due to Jaynes (1987).⁵ The idea is to assume the parameters A and B are uncorrelated,⁶ giving

$$\mathcal{P}(A, B|d) \propto \exp \left[-\frac{(A - A_0)^2}{2\sigma_A^2} - \frac{(B - B_0)^2}{2\sigma_B^2} \right], \quad (34)$$

where A_0 and B_0 are the best-fitting values and σ_A and σ_B are the errors in determining A and B from the data. Now writing $A = K \cos \phi$ and $B = K \sin \phi$, we find

$$\mathcal{P}(K, \phi|d) \propto \exp \left[-\frac{K^2}{2\sigma_K^2} + \frac{K K_0}{\sigma_K^2} \cos(\phi + \phi_0) \right], \quad (35)$$

where ϕ_0 is a constant that can be determined (the precise value is not important here), the best-fitting amplitude is $K_0 = (A_0^2 + B_0^2)^{1/2}$ and we assume $\sigma_A^2 = \sigma_B^2 = \sigma_K^2$. If the variance of the noise is s^2 , we expect to be able to determine the amplitude K to an accuracy $\sigma_K^2 \approx 2s^2/N$. Using this approximation, together with the integral representation of the modified Bessel function

$$I_0(z) = \frac{1}{2\pi} \int_0^{2\pi} dt e^{z \cos t} \quad (36)$$

gives

$$\mathcal{P}(K|d) \propto \exp \left(-\frac{NK^2}{4s^2} \right) I_0 \left(\frac{NK K_0}{2s^2} \right). \quad (37)$$

Since we want a prior for K of $\mathcal{P}(K) \propto 1/K$, we divide the area element $dA dB = K dK d\phi$ by K^2 , giving the final result

$$\mathcal{P}(K|d) dK \propto \exp \left(-\frac{NK^2}{4s^2} \right) I_0 \left(\frac{NK K_0}{2s^2} \right) \frac{dK}{K}. \quad (38)$$

Comparing to the results of our grid search, we find that equation (38) reproduces the distribution of K values at each orbital period remarkably well. We estimate s^2 as the mean square deviation of the residuals to the best-fitting sinusoid, or $s^2(P) = \chi_0^2(P) / \sum w_i$. We normalize the distribution of K at each P so that $\int \mathcal{P}(K, P|d) dK = \mathcal{P}(P|d)$, where $\mathcal{P}(P|d)$ is determined from the analytic marginalization over A and B (equation 27). This choice of normalization as a function of P gives the best agreement with the grid code.

3.5 Summary and example

Let us summarize the main results of this section. We have discussed two methods for evaluating $\mathcal{P}(P, K|d)$ for circular orbits. First, equations (21) and (22) can be used to calculate $\chi^2(P, K, \phi)$ for many different values of P , K and ϕ , and from there $\mathcal{P}(P, K|d)$ obtained by integration over ϕ . No approximations are made in this approach, which we refer to as the ‘grid-based approach’. Secondly,

⁵ We follow a slightly different argument than Jaynes (1987), but with the same spirit. The same approach was used by Groth (1975) to derive the statistical distribution of periodogram powers in the presence of a signal plus Gaussian noise, and recently Shen & Turner (2008) made a similar approximation to derive the shape of the probability density for eccentricity in a Keplerian orbit fit.

⁶ This is a good approximation for large N . The covariance between A and B is $\propto \sum w_i \sin \omega t_i \cos \omega t_i$ which averages to zero for large N .

equations (27–33) provide a method for evaluating $\mathcal{P}(P|d)$ using analytic marginalization over the linear parameters A and B and therefore K and ϕ . The analytic marginalization requires that we assume a prior $\mathcal{P}(K) \propto K$ rather than $1/K$, but by choosing the normalization appropriately (equation 33) we approximately recover the results corresponding to $\mathcal{P}(K) \propto 1/K$. Next, given the best-fitting amplitude $K_0 = (A_0^2 + B_0^2)^{1/2}$ at each period, $\mathcal{P}(P, K|d)$, can be calculated for a grid of K values using the analytic approximation of equation (38). We refer to this second approach as the ‘analytic approach’.

As an example, we consider the 23 radial velocities for the star HD 4203 made available in the Butler et al. (2006) catalogue of nearby exoplanets (see Vogt et al. 2004 for the original discovery of this planet). The orbital parameters given by Butler et al. (2006) are $P = 431.88 \pm 0.85$ d, $K = 60.3 \pm 2.2$ m s⁻¹ and $e = 0.519 \pm 0.027$. They also include a linear long-term trend of -4.38 ± 0.71 m s⁻¹ yr⁻¹. The rms of the residuals to this solution is 4.1 m s⁻¹.

We use both of the techniques described above to fit a circular orbit plus constant to this data. We consider orbital periods between 1 day and the time-span of the data, $T = 2000$ d. We evaluate $4N_f$ frequencies, where $N_f = (\Delta f)T$ is the estimated number of independent frequencies in the frequency range Δf (Cumming 2004). The values of K considered range from 1 m s⁻¹ to $2\Delta v$, where Δv is the range of the measured velocities. For the grid-based approach, we find that the typical time required on a 2–3 GHz CPU is $\sim 10^{-7}$ s per set of parameters (ϕ , K , P), so that for example 3000 periods, 100 values of K and 30 phases, or 10^7 total combinations, take 1 s to evaluate. For ϕ , we align the grid with the best-fitting phase ϕ_0 at each P . In this way, we guarantee that the best-fitting value of ϕ is included on the grid, which reduces the number of grid points we need to use in ϕ . The analytic marginalization technique requires $\sim 5 \times 10^{-7}$ s per P and K value, so that a search of 3000 periods, keeping track of 100 values of K takes ~ 0.1 s. We use the routine `BESS10` from Press et al. (1992) to calculate the Bessel function in equation (38).

Fig. 1 compares the two techniques. The red curves show the results of the analytic marginalization, the black curves show the results of the grid-based calculation. The FAPs are 0.14 (grid) and 0.060 (analytic) (odds ratios 6.3 and 16, respectively). The distribution of K agrees well between the two techniques, although the probability curve is shifted to larger values of K for the analytic approach compared to the grid approach, consistent with the different priors. The FAP ~ 0.1 means that this would not count as a detection. This is an example of a case in which the large eccentricity $e > 0.5$ prevents detection by fitting circular orbits. The best-fitting amplitude $K \approx 30$ m s⁻¹ for circular orbits is significantly smaller than for the Keplerian orbit fit of Butler et al. (2006). Using 100 values of K between 1 and 60 m s⁻¹, we find the 99 per cent upper limit on K is 41.2 m s⁻¹ (analytic) or 41.3 m s⁻¹ (grid). The bottom panel in Fig. 1 compares the probability distribution of K at two different periods obtained from the grid-based approach and the analytic approach. This shows that equation (38) reproduces the distribution from the grid-based calculation well.

4 ECCENTRIC ORBITS

We now consider full Keplerian fits to the data. The techniques we developed in the previous section for circular orbits can be readily applied to Keplerian orbits, because the Keplerian model is linear in a subset of parameters which can therefore be treated analytically, as we now describe.

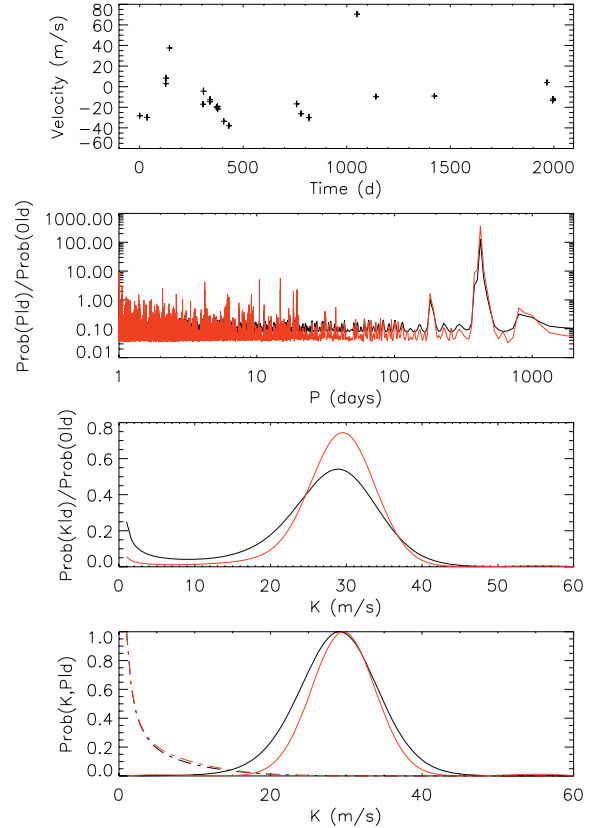


Figure 1. Results of circular orbit fitting to data for HD 4203, using analytic marginalization over K and ϕ , and reconstructing $\mathcal{P}(K)$ using equation (38) (red curves) and by calculating $\mathcal{P}(\phi, K, P)$ on a grid (black curves). $\mathcal{P}(K)$ in the third panel is normalized such that each curve has the same area beneath the curve. The bottom panel compares the $\mathcal{P}(K)$ obtained for periods 420.1 days (solid curves, close to the best-fitting frequency) and 19.0 days (dot-dashed curves, no significant fit at this frequency) for the analytic and grid-based approaches.

4.1 Calculation of $\mathcal{P}(P, K, e|d)$

For a Keplerian orbit, the radial velocity can be written as

$$V = \gamma + K[\cos(\theta + \omega_p) + e \cos \omega_p], \quad (39)$$

where K is the velocity amplitude, e is the eccentricity of the orbit, ω_p is the argument of periastron.⁷ The true anomaly θ is a function of the time t and the three parameters e , P and t_p , where t_p is the time of periastron passage [acting as an overall phase for $V(t)$]. To calculate $\theta(t; e, P, t_p)$, we must solve the relations

$$\tan\left(\frac{\theta}{2}\right) = \left(\frac{1+e}{1-e}\right)^{1/2} \tan\left(\frac{E}{2}\right) \quad (40)$$

$$E - e \sin E = M = \frac{2\pi}{P}(t - t_p), \quad (41)$$

where E is the eccentric anomaly and M is the mean anomaly.

The first point to note is that the six orbital parameters, $\mathbf{a} = (\gamma, K, \omega_p, P, e, t_p)$ can be divided into two groups, ‘slow’ and ‘fast’ parameters, $\mathbf{a}_s = (P, e, t_p)$ and $\mathbf{a}_f = (\gamma, K, \omega_p)$, respectively. Each time we change a value of the slow parameters, we must resolve equations (40) and (41) to calculate the values of θ , whereas when

⁷ We write it as ω_p to distinguish it from the orbital frequency $\omega = 2\pi/P$.

we change a value of the fast parameters only we do not need to recalculate the values of θ . This is reminiscent of the division into fast and slow parameters in analysis of cosmic microwave background data (e.g. Lewis & Bridle 2002; Tegmark et al. 2004). We can use this division to increase the speed of the parameter search.

For a given set of the slow parameters, we can find the best-fitting fast parameters with a linear least-squares fit, since we can write

$$V = A \sin \theta + B \cos \theta + \tilde{\gamma} \quad (42)$$

with $A = -K \sin \omega_p$, $B = K \cos \omega_p$ and $\tilde{\gamma} = \gamma + K e \cos \omega_p$. A linear least-squares fit returns the best-fitting values of A , B and $\tilde{\gamma}$, and therefore K ($K^2 = A^2 + B^2$), ω_p ($\tan \omega_p = -B/A$) and γ . This halves the number of parameters that we need to search to find the best-fitting solution.

The fact that the fast parameters \mathbf{a}_f can be obtained from a linear fit means that we can directly apply the techniques we developed for circular orbits in Section 3 to marginalize over them. For the grid-based approach, equation (21) should be replaced by

$$\frac{\chi^2(K, \omega_p)}{\sum w_i} = \langle \langle v^2 \rangle \rangle + 2K [\langle \langle vS \rangle \rangle \sin \omega_p - \langle \langle vC \rangle \rangle \cos \omega_p] + K^2 [\langle \langle C^2 \rangle \rangle \cos^2 \omega_p + \langle \langle S^2 \rangle \rangle \sin^2 \omega_p - \langle \langle SC \rangle \rangle 2 \sin \omega_p \cos \omega_p], \quad (43)$$

where ω_p now plays the same role as ϕ for circular orbits, and the sums over the data involve θ_i rather than ωt_i . For example, the definition of $\langle S \rangle$ in equation (20) should be replaced by $\langle S \rangle = \sum w_i \sin \theta_i / \sum w_i$.

Similarly, since equations (24) and (42) are of the same form, the analytic integration over A and B can be applied directly to the Keplerian case, giving $\mathcal{P}(P, e, t_p|d)$ analytically from equations (27–33). As for circular orbits, the distribution of velocity amplitude at each (P, e, t_p) , $\mathcal{P}(K, P, e, t_p|d)$, can be recovered, being well approximated by equation (38).

4.2 Example

As an example, we return to the HD 4203 data considered previously. We first calculate $\mathcal{P}(P, e|d)$ for a grid in P and e . The integration over t_p is carried out using a simple algorithm in which we double the number of equally spaced t_p values until the required accuracy is obtained. For each combination of P , e and t_p considered, we analytically integrate over γ , K and ω_p , and at the same time use equation (38) to keep track of $\mathcal{P}(K; P, e, t_p|d)$. We use Newton's method to solve Kepler's equation, taking advantage of the fact that the required derivative can be calculated analytically. Our implementation of this algorithm takes $\approx 5 \times 10^{-5}$ s per P , e and t_p value considered, with 30 K values tracked through the calculation. For an average 200 values of t_p , 10 eccentricities and 3000 periods, the total time needed is ≈ 30 s for a scan of parameter space. We have also implemented the grid-based approach, and find that it is about 10 times slower than the analytic approach. The results agree well between both techniques.

The results for HD 4203 are shown in Figs 2 and 3. We first run a coarse scan of the parameter space for a single Keplerian orbit plus a linear trend. We calculate $4N_f \approx 8000$ frequencies, corresponding to the period range 1 day to ≈ 2000 days (the time-span of the data), 10 eccentricities between 0 and 0.9, 10 velocities between 1 and 216 m s^{-1} (twice the velocity span of the data). The resulting constraints on P , e and K are shown in Fig. 2. The odds ratio is 4×10^4 for the Keplerian orbit plus linear trend compared to a constant-velocity model. We show the results including a linear trend, because




Figure 2. Results of Keplerian fits to the HD 4203 data from Butler et al. (2006), including a linear trend. In this coarse scan of parameter space, $\mathcal{P}(P, e, K|d)$ is calculated for 10 eccentricities between 0 and 0.9, 10 velocities between 1 and 217 m s^{-1} (twice the velocity span of the data) and 7978 periods between 1 and 1996 days (the time-span of the data).

the best-fitting model presented by Butler et al. (2006) includes a trend, but in fact our results at this stage do not require a trend. The odds ratio for a similar search but without the linear term is 5×10^4 .

We then carry out a more detailed calculation of the parameter space near the best-fitting model corresponding to the peak in $\mathcal{P}(P|d)$ at ≈ 440 d in Fig. 2. The results are shown in Fig. 3. The odds ratio is 9.5×10^{10} for a Keplerian orbit plus trend compared to a constant model (for the ranges of parameters shown in Fig. 3). The much larger value of the odds ratio compared to our coarse calculation is because the parameter space considered is smaller and the peak in $\mathcal{P}(P, e, K|d)$ has now been resolved. We can renormalize the odds ratio to correspond to the full range of parameter space considered in the coarse search by multiplying by the ratio of $\log P_2/P_1$ and $\log K_2/K_1$ in each calculation. Doing this, we find an odds ratio 7.4×10^7 . Without the linear trend, the odds ratio is 100 times smaller, 7×10^5 , normalized to the full range of parameters. This indicates that a model with a linear trend is strongly preferred given this data. Without the linear trend, the probability peaks at similar values of P and K , but with a larger eccentricity, $e \approx 0.7$.

The dotted curves in Fig. 3 show Gaussian distributions with the central values and standard deviations given by Butler et al. (2006) for K , P and e . Overall, there is a good agreement with the central values and widths.

Repeating the calculation shown in Fig. 3 with the grid-based method for marginalizing over K and ω_p gives almost identical constraints on orbital parameters, but a smaller odds ratio by a factor of 2, 3.5×10^7 compared to 7.4×10^7 . We have also checked that other peaks in $\mathcal{P}(P|d)$ that can be seen in Fig. 2 do not contribute significantly to the odds ratio. The next most important is the peak

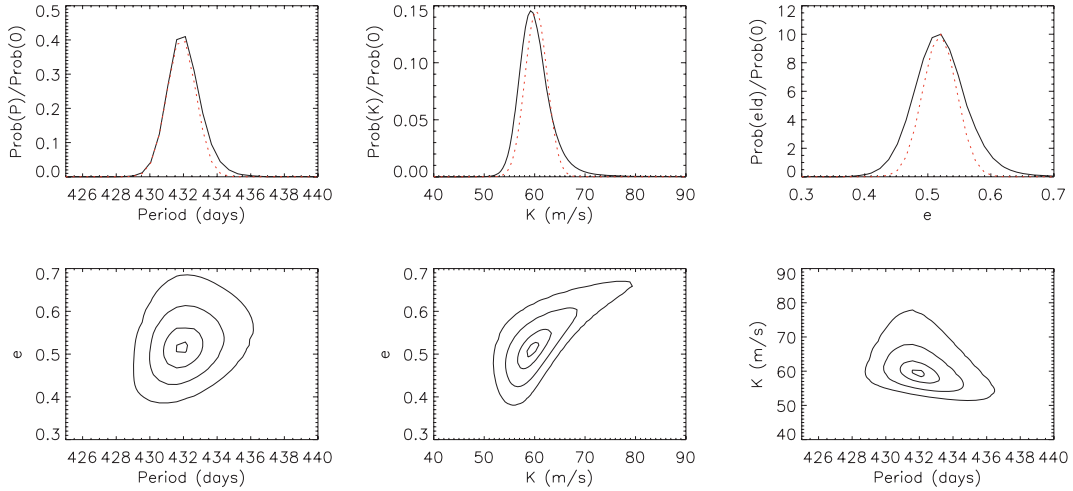


Figure 3. Results of Keplerian fits to the HD 4203 data from Butler et al. (2006), including a long-term trend. The dotted curves show Gaussian distributions with central values and standard deviations matching those given by Butler et al. (2006). 100 values of K , 30 eccentricities and 30 periods were calculated in the range shown. The contours enclose 10, 50, 90 and 99 per cent of the probability.

at $P \approx 800$ d, but its odds ratio is 400 times smaller than the peak at 432 d shown in detail in Fig. 3.

The coarse sampling for HD 4203 gave an odds ratio that was a factor of 400 smaller than the final odds ratio obtained by zooming in on the most significant peak. We find that increasing the period sampling by a factor of 2 to $8T\Delta f$ gives an odds ratio from the coarse search in good agreement with the odds ratio from zooming in on the peak.

5 COMPARISON WITH PREVIOUS WORK

In Section 4, we presented an algorithm that can efficiently compute $\mathcal{P}(P, K, e|d)$ for a radial velocity data set. As described in Section 2, this contains information about the constraints on P , K , and e and also allows a FAP to be calculated. We now use our algorithm to recalculate results in the literature from MCMC and other techniques and compare.

5.1 Orbital parameter constraints from MCMC calculations

Ford (2005) used a MCMC calculation to study the constraints on orbital parameters from radial velocity data, and this paper has been followed by several others (Gregory 2005b; Ford 2006; Ford & Gregory 2007; Gregory 2007a,b; Ford 2008; Balan & Lahav 2009). We have calculated the constraints on orbital parameters for the different single-planet cases considered in these papers, and overall the agreement is excellent.

One difference is that in several published cases the posterior probability for eccentricity drops towards zero at low eccentricities, whereas we find $\mathcal{P}(e|d)$ is approximately constant as e goes to zero. For HD 76700, this difference appears to be because of the different prior assumed by Ford (2005). The MCMC calculations in that paper take steps in $e \cos \omega_p$ and $e \sin \omega_p$ in such a way that the assumed prior is uniform in $d(e \cos \omega_p) d(e \sin \omega_p)$ giving a prior $e de d\omega_p \propto e$. In Fig. 4, we allow for this different prior by plotting $e\mathcal{P}(e|d)$, and the result compares favourably with Fig. 2 of Ford (2005). (Ford 2005 discusses the use of importance sampling, in which the samples are weighted $\propto 1/e$ to give an effective prior uniform in e , but this does not seem to have been applied in Fig. 2 of that paper.)




Figure 4. The eccentricity distribution derived for HD 76700, using data from Tinney et al. (2003). The solid curve is for analytical marginalization over the noise scaling parameter k , the dotted curve is for $k = 1$, and the dashed curve shows $e\mathcal{P}(e|d)$, corresponding to a uniform prior in $d(e \cos \omega_p) d(e \sin \omega_p)$.

For HD 72659, marginalization over the extra noise source opens up considerable parameter space at low eccentricity. In Fig. 5, we show the constraints on eccentricity and period with k fixed at $k = 1$ and with k marginalized over. Ford (2005), unlike later papers (e.g. Ford 2006), does not include an additional noise term, and our results for $k = 1$ compare well with Figs 4 and 5 of that paper.

5.2 Odds ratios from Gregory's parallel tempering MCMC approach

In a series of papers, Gregory has developed a MCMC code which uses parallel tempering to exchange information between chains running with different 'temperatures'. Combining the results of different chains gives the total posterior probability for the model, allowing calculations of odds ratios and therefore model comparisons.

Gregory (2005b) analysed 18 radial velocities for HD 73526 from Tinney et al. (2003). The period range was from 0.5 to 3732 days, and velocities from 0 to 400 m s^{-1} using a Jeffrey's prior with a break at 1 m s^{-1} . An additional noise term was added which was

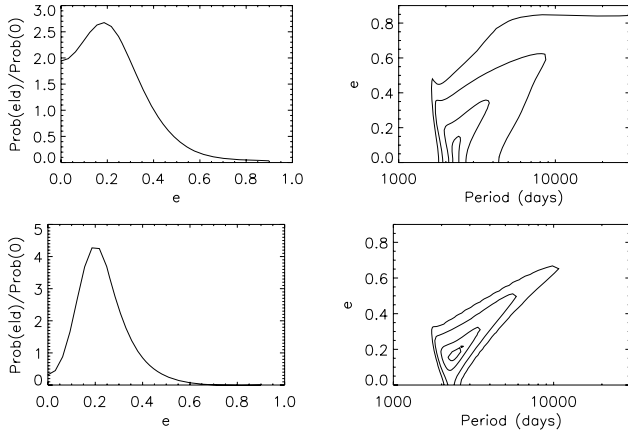


Figure 5. The eccentricity distribution and joint eccentricity–orbital period distribution derived for HD 72659. The top panels are for analytic marginalization over the noise parameter k , whereas the bottom panels take $k = 1$ as in Ford (2005). Contours enclose 10, 50, 90 and 99 per cent of the probability.

allowed to range between 0 and 100 m s^{-1} . He pointed out that there were two additional possible solutions with $P \approx 128$ and 376 d besides the previously obtained solution at $P \approx 191 \text{ d}$. A chain covering the entire parameter space did not converge, and so separate chains were run focussing on each of the three probability peaks. The odds ratio for a planet compared to a constant velocity was found to be 9.3×10^5 (table 5 of Gregory 2005b). We ran a calculation with the same period and velocity range as Gregory (2005b) (except that we take the lower bound in K to be 1 m s^{-1} with a Jeffrey’s prior) (9940 frequencies, 30 eccentricities and 30 velocities). The odds ratio was 3.3×10^6 . Zooming in on the three peaks gives probability distributions for e , P and K that are very similar to the results of Gregory (2005b). The odds ratios for the $P \approx 128$, 190 , and 376 d peaks are 2.4×10^4 , 1.1×10^5 , and 1.0×10^6 (assuming the full prior range so that these numbers can be compared). The sum of these, 1.1×10^6 agrees well with the odds ratio found by Gregory (2005b) whose odds ratio includes only these three peaks. The relative probabilities of the three peaks are 2, 10 and 88 per cent. Gregory (2005b) found relative probabilities of 4, 3 and 93 per cent.

Gregory (2007a) found evidence for a second planet in HD 208487; we compare to their odds ratio and posterior probability for a one-planet fit. The posterior probability distributions were calculated for the 35 velocities from Butler et al. (2006). We find an excellent agreement with the distributions of P , e and K shown in fig. 7 of Gregory (2007a). The odds ratio for a single-planet model for this data (table 6 of Gregory 2007a) was $1.7\text{--}2.6 \times 10^4$ for two different choices of the turnover in the modified Jeffrey’s prior for the extra noise scale. For the parameter ranges in fig. 7 of Gregory (2007a), we find an odds ratio of 1.4×10^8 . Rescaling to a velocity range $1\text{--}2129 \text{ m s}^{-1}$, and period range 1 day to 1000 years, this becomes 6.1×10^4 , a factor of 3 times greater than Gregory (2007a). (The details of the priors were different, e.g. the upper limit on velocity in Gregory 2007a’s prior depended on period and eccentricity, but we expect this to give only a small difference.)

Gregory (2007b) presented evidence for three planets in HD 11964 from 87 radial velocities in the Butler et al. (2006) catalogue. The odds ratio reported for the single-planet model is 3×10^9 (table 4 of Gregory 2007b). We find an odds ratio in good agreement, 2×10^9 . Although Gregory (2007b) does not show posterior probability distributions for orbital parameters for the one-planet

model, the distributions of P , e and K we find compare well with those for the $P \approx 2000 \text{ d}$ signal in the three-planet model of Gregory (2007b). For this data, Butler et al. (2006) include a linear term. We find the odds ratio for a linear versus constant no-planet model to be 1300 . Including a linear term in the planet model gives an odds ratio of 3×10^6 , much smaller than the odds ratio for a planet model only. Therefore, we find that a single-planet model with $P \approx 2000 \text{ d}$ is preferred over a linear trend only or planet plus linear trend by a large factor (in agreement with Wright et al. 2007 who also concluded that the trend reported by Butler et al. 2006 was likely spurious).

5.3 False-alarm probabilities

Marcy et al. (2005) discuss the calculation of FAPs using a scrambled velocity method in which the residuals to the best-fitting Keplerian orbit are used as an estimate of the noise distribution. In that paper, they announced five new planets from the Keck Planet Search. FAPs were calculated for two cases that looked marginal, HD 45350 (FAP < 0.1 per cent scrambled, 4×10^{-5} F -test) and HD 99492 (FAP ≈ 0.1 per cent scrambled, 3×10^{-4} F -test). For HD 99492, we find odds ratios scaled to 1.0 for no planet are 0.33 for a linear trend but no planet, 1.66 for a planet, 200.0 for a planet plus linear trend. Therefore, a linear trend is preferred in this case. The FAP using equation (13) for the odds ratio is 7×10^{-3} . For HD 45350, we find odds ratios: 1.0, 0.18, 6.7×10^5 , 4.7×10^5 , giving FAP $\approx 10^{-6}$. As Marcy et al. (2005) noted, the evidence for a linear trend in this source is marginal (the odds ratios are similar with and without a trend).

Cumming (2004) described a quick estimate of the FAP based on an F -test at each independent frequency. Generalizing the Lomb–Scargle periodogram to eccentric orbits, the idea is to define a power at each frequency

$$z = \frac{(N-5)\Delta\chi^2}{4\chi_{\text{Kep}}^2} = \frac{(N-5)(\chi_0^2 - \chi_{\text{Kep}}^2)}{4\chi_{\text{Kep}}^2}. \quad (44)$$

For Gaussian noise, z follows the $F_{4,N-5}$ distribution,⁸ which allows a calculation of $\text{Prob}(z > z_{\text{max}})$ for an observed maximum power z_{max} . The FAP is then

$$\text{FAP} = 1 - [1 - \text{Prob}(z > z_{\text{max}})]^{N_f} \approx N_f \text{Prob}(z > z_{\text{max}}). \quad (45)$$

The number of independent frequencies N_f can be estimated as $N_f \approx T\Delta f$.

We have used this approach to calculate the FAP for the 84 stars with published radial velocities as part of the Butler et al. (2006) catalogue of exoplanets. To find χ_{Kep}^2 , we follow the automated procedure used by Cumming et al. (2008), which involves using the top two well-separated peaks in the Lomb–Scargle periodogram as starting periods for full Keplerian fits. To compare with the Bayesian odds ratios, we convert the F -test FAP into an odds ratio by inverting equation (9). To find Bayesian odds ratios, we run a coarse sampling of the parameter space with $8T\Delta f$ periods for each of these 84 stars, with and without a long-term linear trend.

The results are shown in Fig. 6. In the lower panel, we use the minimum value of χ^2 found in the Bayesian calculation to calculate the F -test FAP. In this case, the odds ratios are well correlated,

⁸ Assuming that the no-planet model being compared to is a constant-velocity model. If a linear trend is included in the no-planet model and the planet model, z is defined with a factor of $N-6$ replacing $N-5$, and then follows the $F_{4,N-6}$ distribution.

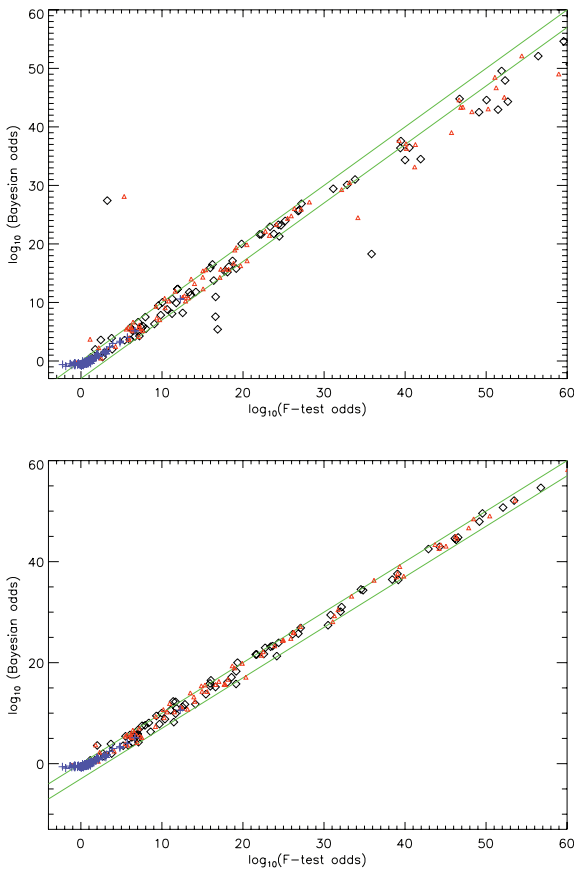


Figure 6. Comparison between logarithm of the odds ratio $\log_{10} \Lambda$ from the Bayesian calculation and the F -test. The Bayesian odds ratios are from coarse sampling ($8T \Delta f$ periods, 10 eccentricities) of the 84 radial velocity data sets from Keck, Lick and Anglo-Australian Telescope published as part of the Butler et al. (2006) catalogue. We compare with analytic F -test FAPs, converted to odds ratios using the relation $\Lambda = (1/F) - 1$. The crosses are the odds ratios for a linear trend versus constant velocity, the diamonds are odds ratios for a planet versus constant and the triangles are for planet plus long-term trend versus long-term trend only. The upper panel uses a Keplerian fitting routine to determine the F -test FAP, whereas in the lower panel we use the minimum χ^2 found in the Bayesian routine to calculate the analytic FAP.

although with the Bayesian odds ratio between 1 and 1000 times smaller than the F -test odds ratio. In the upper panel, there is more scatter. This arises from differences between the minimum χ^2 values found by the Keplerian fitting routine and the Bayesian routine. For example, the two points above and to the left of the upper panel of Fig. 6 are for HD 80606, which has a very eccentric orbit. Our Keplerian fitting routine, which uses circular orbit fits as its starting point failed to find the best-fitting solution, whereas the Bayesian routine with its systematic scan of parameter space did find it. Generally, the scatter is downwards, indicating that the Bayesian routine sometimes find a larger minimum χ^2 than the Keplerian fitting routine. Likely this is due to the finite period sampling, whereas the Keplerian fitting routine can adjust the period to lower χ^2 .

The fact that the Bayesian odds ratios tend to be lower than the F -test odds ratios indicates that the Bayesian calculation is more conservative than the F -test. In fact, this is expected. Cumming (2004) showed that the Bayesian odds ratio is closely related to the F -test (periodogram), but with a different definition for the number of independent frequencies. In the Bayesian calculation, the

number of trials not only counts the frequencies, but also the range of the other parameters (Cumming 2004). In this way, the Bayesian calculation penalizes models with larger ranges of parameters, for all parameters, not just frequency.

5.4 2D periodograms

Wright et al. (2007) investigate the constraints that can be placed on the orbital parameters of long-period orbits that have been only partially observed. They calculated the minimum χ^2 at points across the $m \sin i - P$ plane. Similarly, O’Toole et al. (2008) introduced a ‘2D Keplerian Lomb–Scargle periodogram (2DKLS)’ in which the periodogram power is evaluated on a grid of P and e , with a full Keplerian fit carried out at each point. O’Toole et al. (2008) discuss the considerable computing resources being used to conduct simulations of detectability using this new 2D periodogram. The techniques we discuss earlier for rapid evaluation of multiple χ^2 values could prove useful in more efficiently evaluating the 2DKLS periodogram. The constraints on $P-e$ or $P-K$ calculated in this paper differ from Wright et al. (2007) and O’Toole et al. (2008) in that for each choice of P , e or P , K all values of the other parameters are taken into account, weighted by their probability, rather than finding the best-fitting values of the other parameters. This is the standard difference between Bayesian and frequentist approaches.

O’Toole et al. (2008) mention that one of the reasons for looking at the periodogram power as a function of P and e is to help with detection of highly eccentric orbits. They consider the $e = 0.97$ planet around HD20782 as an example. Their best fit has $e = 0.97 \pm 0.01$, $P = 591.9 \pm 2.8$ and $K = 185.3 \pm 49.7$. Our results for this data are shown in Fig. 7. The discrete nature of the K distribution is due to the finite sampling of the grid in eccentricity. The O’Toole et al. (2008) solution lies on our contours, but towards the edge. The Bayesian calculation, which averages over the marginalized parameters, opens up a wider parameter space than the best fit and error bars from O’Toole et al. (2008) suggest.

5.5 HD 5319

HD 5319 has a planet with minimum mass $1.9 M_J$ in a 675-d low-eccentricity orbit (Robinson et al. 2007). This is an interesting example to compare to because the analysis of Robinson et al. (2007) used several different statistical methods. First, they used Monte Carlo simulations of data sets with noise only (simulated by selecting with replacement from the observed velocities) to assess the FAP, finding 1.3×10^{-3} . They used both a scrambled velocity Monte Carlo simulations and an MCMC Bayesian calculation to estimate the uncertainties in the derived orbital parameters. They used an F -test to test the significance of including a linear trend in their model, finding a FAP of 3×10^{-4} indicating that a linear term is strongly preferred.

The results of our calculation are shown in Fig. 8. The dotted lines show the best-fitting parameters and the errors found by Robinson et al. (2007), assuming Gaussian distributions, and agree well both in terms of central values and widths. Interestingly, the MCMC simulations run by Robinson et al. (2007) did not agree as well with their scrambled velocity approach, whereas we find good agreement. The odds ratio for a trend in the no-planet model is 0.9. For models with a planet, the odds ratios are 9.0×10^8 (with trend) and 1.0×10^6 (without trend). The model with a trend therefore has greater odds by a factor of 10^3 , in good agreement with the F -test FAP of 3×10^{-4} found by Robinson et al. (2007). However, the

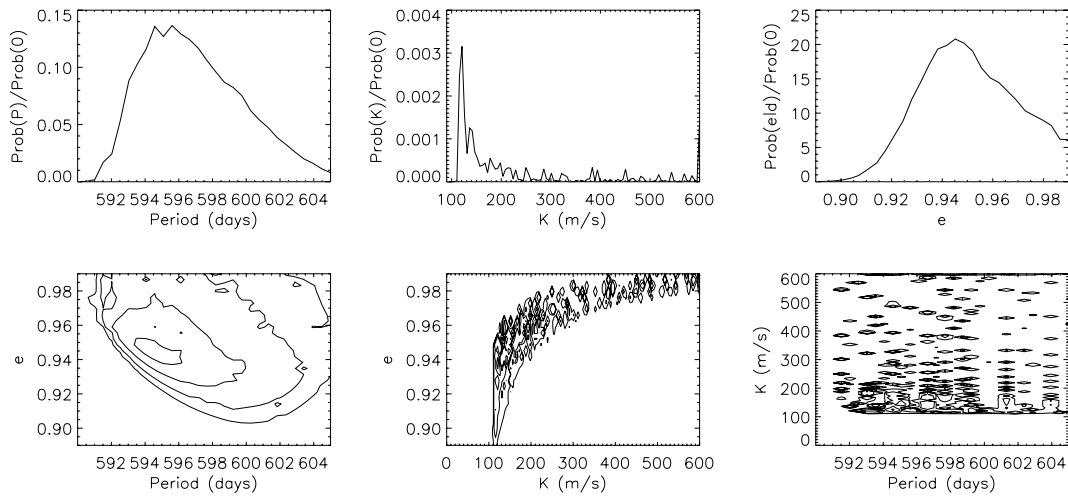


Figure 7. Results of Keplerian fits to the HD 20782 data from O’Toole et al. (2008). 30 eccentricities, 30 periods and 100 K values were calculated in the ranges shown. Contours enclose 10, 50, 90 and 99 per cent of the total probability.

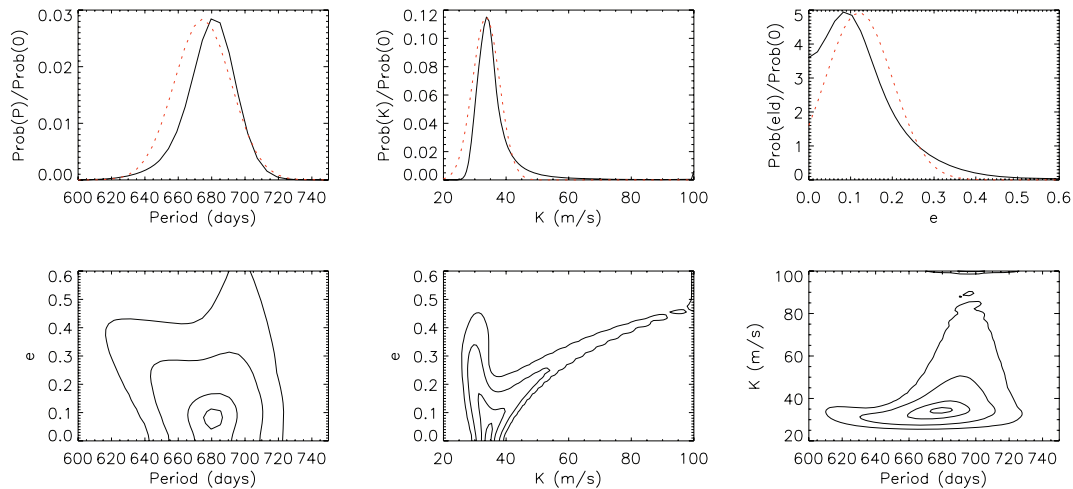


Figure 8. Results of Keplerian fits to the 5319 data including a linear trend. 30 eccentricities, 30 periods and 100 K values were calculated in the ranges shown. Dotted curves show Gaussian distributions with central values and standard deviations taken from Robinson et al. (2007). Contours enclose 10, 50, 90 and 99 per cent of the total probability.

overall FAP we find $\sim 10^{-9}$ is much smaller than the simulations of Robinson et al. (2007) suggested, $\sim 10^{-3}$.

5.6 Upper limits on K

In an analysis of the Lick Planet Search, Cumming et al. (1999) calculated upper limits for 63 stars with non-detections. They used a Monte Carlo approach, in which simulated data sets with a circular orbit plus noise were analysed and the velocity amplitude determined which resulted in detection 99 per cent of the time. We have reanalysed the same data using our Bayesian scheme, first with circular orbits and then with eccentric orbits. We calculate the 99 per cent upper limit K_{99} by $\int_0^{K_{99}} dK \mathcal{P}(K|d) = 0.99$ [where $\mathcal{P}(K|d)$ is normalized so that the total probability is unity].

The results are shown in Fig. 9. The triangles are for circular orbit models, and the circles are for eccentric orbits. For the seven stars found to have a significant linear trend by Cumming et al. (1999), we include a linear trend in the model. Overall, the agreement is good. Cumming et al. (1999) (and Cumming et al. 2008) calculate upper limits for circular orbits to reduce the computational time

needed. Based on the calculations of the effect of eccentricity on detectability of Endl et al. (2002) and Cumming (2004), they proposed that K_{99} for circular orbits would be a good estimate of K_{99} for orbits with $e \lesssim 0.5$. We can test that here by calculating K_{99} from the partial K distribution

$$\mathcal{P}(K|d) = \int_0^{e_{\text{cut-off}}} de \mathcal{P}(e, K|d) \quad (46)$$

with different cut-offs $e_{\text{cut-off}}$. For $e_{\text{cut-off}} = 0.9$, we find that the value of K_{99} is generally much greater than K_{99} for circular orbits, due to a tail of large- K , large-eccentricity solutions. However, for $e_{\text{cut-off}} = 0.5$, the agreement is very good. This is shown in Fig. 9, where we show results for $e_{\text{cut-off}} = 0.5$ (red symbols) and $e_{\text{cut-off}} = 0.7$ (green symbols). The $e_{\text{cut-off}} = 0.7$ values of K_{99} are significantly greater than the circular orbit or $e_{\text{cut-off}} = 0.5$ values.

6 SUMMARY AND CONCLUSIONS

In this paper, we consider the Bayesian analysis of radial velocity data. An advantage of this kind of analysis over traditional methods




Figure 9. Comparison between the 99 per cent upper limits on K for circular orbits determined by Cumming et al. (1999) for 63 stars from the Lick Planet Search and the 99 per cent upper limits on K from a Bayesian analysis of the same data. Solid triangles or circles include a linear trend in the fits [these are the seven stars that Cumming et al. (1999) found to have a significant slope], whereas open triangles or circles do not include a linear trend. The dotted line indicates a 1:1 correspondence between the two calculations of the upper limit. The black triangles are for circular orbits, the red circles are for eccentric orbits with $e < 0.5$ and the green circles are for eccentric orbits with $e < 0.7$.

is that a single calculation gives the FAP and the probability distributions of orbital period, eccentricity and velocity amplitude, allowing error bars or upper limits on these quantities to be determined. Using periodogram methods, separate calculations are required for each of these quantities, typically requiring many Monte Carlo trials.

Previous work on the Bayesian analysis of radial velocities has used MCMC techniques (although see Ford 2008 who used analytic techniques to partially carry out the marginalization for circular orbits). Our approach has been to apply some exact and approximate analytic results (based on previous work by Jaynes 1987; Bretthorst 1988) to the marginalization integrals for Keplerian fits to radial velocity data. In particular, we analytically integrate over the linear model parameters for each combination of P , e and t_p , and use an analytic approximation (equation 38) to reconstruct the probability distribution of K . An implementation of this algorithm in IDL is available on request from the authors.

With this approach, a full search of parameter space for a single Keplerian orbit takes several minutes on a 2–3 GHz processor, or several seconds for circular orbits, making it applicable to data sets from large-velocity surveys. Constraints on orbital parameters (which involve surveying smaller regions of parameter space) can be calculated in seconds, competitive with MCMC techniques.⁹ Our calculation can certainly be improved further. For example, we have focused on the marginalization over the linear parameters in this paper, and used the simplest approach of evaluation on an evenly spaced grid to integrate over the remaining parameters P , K and e .

We compared our results with previous calculations. The constraints on orbital parameters and odds ratios agree well with

⁹ In Ford (2006), the computer time needed was $\sim 10^{-6} s N_{\text{obs}} N_p L_c N_c$, where N_{obs} is the number of observations, N_p is the number of planets, L_c the length of each chain and N_c is the number of chains considered. For 30 observations, 1 planet, 10 chains each of length 10^4 (multiple chains are required to assess convergence; Ford 2006), the total time required is ≈ 3 s.

MCMC results. We find that the Bayesian odds ratios are systematically lower than F -test odds ratios by a factor between 1 and 1000. This is due to the different accounting of trials in the two calculations (Cumming 2004), with the Bayesian calculation including an Occam’s razor penalty which accounts for the range of all parameters rather than only the frequency range. The techniques we have developed for rapidly calculating χ^2 may have application to other techniques, such as the 2D periodograms of Wright et al. (2007) and O’Toole et al. (2008). We find good agreement with the upper limits on velocity amplitude K calculated for circular orbits by Cumming et al. (1999) if we restrict our attention to $e \lesssim 0.5$. More eccentric orbits give rise to a tail of solutions at large K . This shows that characterizing the K distribution with a single parameter (e.g. the 99 per cent upper limit; Cumming et al. 2008) is not appropriate for population analyses with highly eccentric orbits included. On the other hand, for low-to-moderate eccentricity orbits ($e \lesssim 0.5$), upper limits can be derived from circular orbit fits which is much less numerically intensive.

The division of Keplerian parameters into ‘fast’ and ‘slow’ may prove useful in MCMC simulations. At the least, the systemic velocity does not need to be included as a parameter; it can be quickly evaluated for each set of the other parameters, and used to evaluate χ^2 (this was also noted by Ford 2006). One possible complication is that Ford (2005) takes steps in a mixture of fast and slow parameters, $e \cos \omega$ and $e \sin \omega$, to help speed convergence. Separating the slow and fast parameters could potentially reduce efficiency in this case. Further investigations are needed.

ACKNOWLEDGMENTS

We thank Tyler Dodds for some early work on this problem during the summer of 2005 and Gil Holder for useful comments. AC acknowledges support from the National Sciences and Engineering Research Council of Canada, Le Fonds Québécois de la Recherche sur la Nature et les Technologies and the Canadian Institute for Advanced Research. AC is an Alfred P. Sloan Research fellow.

REFERENCES

- Balan S. T., Lahav O., 2009, MNRAS, 394, 1936
- Bretthorst L. G., 1988, Lecture Notes in Statistics, Vol. 48, Bayesian Spectrum Analysis and Parameter Estimation. Springer-Verlag, Berlin
- Butler R. P. et al., 2006, ApJ, 646, 505
- Cumming A., 2004, MNRAS, 354, 1165
- Cumming A., Marcy G. W., Butler R. P., 1999, ApJ, 526, 890
- Cumming A., Butler R. P., Marcy G. W., Vogt S. S., Wright J. T., Fischer D. A., 2008, PASP, 120, 531
- Endl M., Kürster M., Els S., Hatzes A. P., Cochran W. D., Dennerl K., Döbereiner S., 2002, A&A, 392, 671
- Ford E. B., 2005, AJ, 129, 1706
- Ford E. B., 2006, ApJ, 642, 505
- Ford E. B., 2008, AJ, 135, 1008
- Ford E. B., Gregory P. C., 2007, in Babu G. J., Feigelson E. D., eds, ASP Conf. Ser. Vol. 371. Statistical Challenges in Modern Astronomy IV. Astron. Soc. Pac., San Francisco, p. 189
- Gregory P. C., 2005a, Bayesian Logical Data Analysis for the Physical Sciences: A Comparative Approach with Mathematica Support. Cambridge Univ. Press, Cambridge
- Gregory P. C., 2005b, ApJ, 631, 1198
- Gregory P. C., 2007a, MNRAS, 374, 1321
- Gregory P. C., 2007b, MNRAS, 381, 1607
- Groth E. J., 1975, ApJS, 29, 285

- Jaynes E. T., 1987, in Smith C. R., Erickson G. J., eds, *Maximum Entropy and Bayesian Spectral Analysis and Estimation Problems*. Reidel, Dordrecht, p. 1
- Lewis A., Bridle S., 2002, *Phys. Rev. D*, 66, 103511
- Lomb N. R., 1976, *Ap&SS*, 39, 447
- Marcy G. W., Butler R. P., Vogt S. S., Fischer D. A., Henry G. W., Laughlin G., Wright J. T., Johnson J. A., 2005, *ApJ*, 619, 570
- O'Toole S. J., Tinney C. G., Jones H. R. A., Butler R. P., Marcy G. W., Carter B., Bailey J., 2008, *MNRAS*, 1406
- Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 1992, *Numerical Recipes: The Art of Scientific Computing*, 2nd edn. Cambridge Univ. Press, Cambridge
- Robinson S. E. et al., 2007, *ApJ*, 670, 1391
- Scargle J. D., 1982, *ApJ*, 263, 835
- Shen Y., Turner E. L., 2008, *ApJ*, 685, 553
- Sivia D. S., 1996, *Data Analysis: A Bayesian Tutorial*. Oxford Univ. Press, Oxford
- Tegmark M. et al., 2004, *Phys. Rev. D*, 69, 103501
- Tinney C. G., Butler R. P., Marcy G. W., Jones H. R. A., Penny A. J., McCarthy C., Carter B. D., Bond J., 2003, *ApJ*, 587, 423
- Vogt S. S., Butler R. P., Marcy G. W., Fischer D. A., Pourbaix D., Apps K., Laughlin G., 2002, *ApJ*, 568, 352
- Walker G. A. H., Walker A. R., Irwin A. W., Larson A. M., Yang S. L. S., Richardson D. C., 1995, *Icarus*, 116, 359
- Wittenmyer R. A., Endl M., Cochran W. D., Hatzes A. P., Walker G. A. H., Yang S. L. S., Paulson D. B., 2006, *AJ*, 132, 177
- Wright J. T., 2005, *PASP*, 117, 657
- Wright J. T. et al., 2007, *ApJ*, 657, 533

APPENDIX A: INCLUDING A LINEAR TERM (LONG-TERM TREND)

In the main text, the ‘no-planet’ model that we have compared the sinusoid and Kepler fits to was a constant-velocity model. Often, a linear term is included in the fit to account for long time-scale trends in the data. Since adding a linear trend adds one extra linear term to the model, we can analytically marginalize over the slope in the same way as we marginalize over the constant term. In this Appendix, we give the formulae to do that.

A1 Is there evidence for a long-term trend?

First, consider a constant versus a linear model. Minimizing χ^2 as a function of γ for $V_i = \gamma$, we find the best-fitting constant term is

$$\gamma_0 = \langle v \rangle, \quad (\text{A1})$$

the corresponding minimum value of χ^2 is

$$\frac{\chi_{\text{const}}^2}{\sum w_i} = \langle \langle v^2 \rangle \rangle \quad (\text{A2})$$

and

$$\det \alpha = \sum w_i. \quad (\text{A3})$$

Inserting these expressions into equation (26) with the number of parameters $m = 1$ gives the posterior probability for a fit of a constant.

For a straight line fit, $V_i = \gamma + \beta t_i$, we find

$$\gamma_0 = \frac{\langle v \rangle \langle t^2 \rangle - \langle vt \rangle \langle t \rangle}{\langle \langle t^2 \rangle \rangle} \quad (\text{A4})$$

$$\beta_0 = \frac{\langle \langle vt \rangle \rangle}{\langle \langle t^2 \rangle \rangle} \quad (\text{A5})$$

$$\frac{\det \alpha}{(\sum w_i)^2} = \langle \langle t^2 \rangle \rangle \quad (\text{A6})$$

$$\frac{\chi_{\text{line}}^2}{\sum w_i} = \langle v^2 \rangle - 2\gamma_0 \langle v \rangle - 2\beta_0 \langle vt \rangle + \gamma_0^2 + 2m\beta_0\gamma_0 \langle t \rangle + \beta_0^2 \langle t^2 \rangle. \quad (\text{A7})$$

Using equation (26), the odds ratio is

$$\Lambda = \frac{(\chi_{\text{line}}^2)^{-(N-2)/2}}{(\chi_{\text{const}}^2)^{-(N-1)/2}} \left[\frac{\pi}{(\sum w_i) \langle \langle t^2 \rangle \rangle} \right]^{1/2} \frac{\Gamma[(N-2)/2]}{\Gamma[(N-1)/2]} \frac{1}{\Delta\beta}, \quad (\text{A8})$$

where $\Delta\beta$ is the prior range for β (the prior range for γ is the same in both models, and cancels). Here, we take β to lie between $\pm\Delta v/T = \pm(v_{\text{max}} - v_{\text{min}})/T$, giving $\Delta\beta = 2\Delta v/T$, i.e. we use the range of velocity amplitudes that we consider and the time of the observations to set the range of slopes.

There is an important issue to mention here (we thank the referee for raising it) that the prior range of parameters should not depend on the data (the prior probability should reflect our state of knowledge before the data were taken.). That is not true here since the range of observed velocities is used to determine what range of velocity amplitudes to search. Strictly, the normalization of the prior should not reflect this but be completely independent of the data. For example, the range of slopes could be set by looking at the range of slopes in previous planet discoveries (e.g. in Butler et al. 2006 the reported slopes extend to $\approx 100 \text{ m s}^{-1} \text{ yr}^{-1}$) or the range of velocity amplitudes extend up to a maximum set by the amplitude induced by a $\approx 10 M_J$ companion (Gregory 2005b; Ford & Gregory 2007). However, the final odds ratios are not very sensitive to the exact choice of prior range. The range of velocity amplitudes enters the normalization logarithmically (since the prior is taken to be uniform in log). The range of slopes has the largest effect since it enters linearly, but we find that using a different choice, e.g. a range of β from -100 to $+100 \text{ m s}^{-1} \text{ yr}^{-1}$ changes the odds ratios by factors of a few to several only.

A2 Including a trend in the circular or Keplerian orbit fit

Consider the model

$$V_i = \gamma + \beta t_i + A \sin \theta_i + B \cos \theta_i, \quad (\text{A9})$$

where $\theta_i = 2\pi t_i/P$ for a circular orbit fit. Minimizing χ^2 with respect to the four parameters γ, β, A, B , we find that their best-fitting values can be written in a concise way by defining a new average:

$$\overline{xy} \equiv \langle \langle xy \rangle \rangle - \frac{\langle \langle xt \rangle \rangle \langle \langle yt \rangle \rangle}{\langle \langle t^2 \rangle \rangle}. \quad (\text{A10})$$

Using this notation,

$$\gamma_0 = \langle v \rangle - \beta_0 \langle t \rangle - A \langle S \rangle - B \langle C \rangle \quad (\text{A11})$$

$$\beta_0 = \frac{\langle \langle vt \rangle \rangle - A \langle \langle St \rangle \rangle - B \langle \langle Ct \rangle \rangle}{\langle \langle t^2 \rangle \rangle} \quad (\text{A12})$$

$$A_0 = \frac{\overline{vS} \overline{C^2} - \overline{vC} \overline{SC}}{\overline{C^2} \overline{S^2} - \overline{SC}^2} \quad (\text{A13})$$

$$B_0 = \frac{\overline{vC} \overline{S^2} - \overline{vS} \overline{SC}}{\overline{C^2} \overline{S^2} - \overline{SC}^2}. \quad (\text{A14})$$

The expressions for A_0 and B_0 are the same as previously, but with the new averages. We also find

$$\frac{\det \alpha}{(\sum w_i)^4} = \langle \langle t^2 \rangle \rangle \left[\overline{S^2 C^2} - (\overline{S C})^2 \right] \quad (\text{A15})$$

and

$$\begin{aligned} \frac{\chi_0^2}{\sum w_i} &= \langle \langle v^2 \rangle \rangle - 2(A_0 \langle \langle vS \rangle \rangle + B_0 \langle \langle vC \rangle \rangle) \\ &\quad + A_0^2 \langle \langle S^2 \rangle \rangle + B_0^2 \langle \langle C^2 \rangle \rangle + 2A_0 B_0 \langle \langle SC \rangle \rangle \\ &\quad - \beta_0^2 \langle \langle t^2 \rangle \rangle \\ &= \overline{v^2} - 2(A_0 \overline{vS} + B_0 \overline{vC}) \\ &\quad + A_0^2 \overline{S^2} + B_0^2 \overline{C^2} + 2A_0 B_0 \overline{SC}. \end{aligned} \quad (\text{A16})$$

Equations (A11–A16) replace equations (28–32) when a long-term trend is included. They are essentially the same, but with the average \overline{xy} used instead of $\langle \langle xy \rangle \rangle$. Equation (26) with $m = 4$ then allows marginalization over the four parameters A , B , β and γ .

Similarly, for the grid-based approach, the expression for χ^2 is of the same form as equation (21), but with the averages calculated as \overline{xy} instead of $\langle \langle xy \rangle \rangle$.

APPENDIX B: LIKELIHOOD FOR FIXED NOISE SCALING PARAMETER K

In the main text, we integrated over the noise scaling parameter k , giving the likelihood in equation (6) (t -distribution) rather than equation (2) (exponential). As we argued in Section 2.3, the analytic marginalization over an infinite range of k is a good approximation

for a reasonable spread in k . However, it could be that we are able to predict k quite accurately, e.g. if the level of stellar jitter has been predetermined for a particular star, in which case we might want to carry out a calculation for fixed k . Also, this would allow a calculation of the posterior probability for k .

For fixed k , we have $\mathcal{P}(d|\mathbf{a}) \propto k^{-N} \exp[-\chi^2(\mathbf{a})/2k^2]$. The normalization over the constant term is then, taking circular orbits as an example,

$$\mathcal{P}(d|\phi, K, P) \propto \int_{-\infty}^{\infty} d\gamma k^{-N} \exp\left(-\frac{\chi^2}{2k^2}\right), \quad (\text{B1})$$

where $\chi^2(\gamma)$ has the quadratic form of equation (15). Therefore, we can take

$$\mathcal{P}(d|\phi, K, P) = k^{-(N-1)} \exp\left(-\frac{\chi^2[\gamma_0, \phi, K, P]}{2k^2}\right) \quad (\text{B2})$$

as a replacement for equation (19), where we set the prefactor to unity as it cancels when we form the odds ratio.

Similarly, the analytic result giving marginalization over m parameters for a general linear model (equation 26) becomes

$$\begin{aligned} &\int d^m \mathbf{a} k^{-N} \exp\left(-\frac{\chi^2}{2k^2}\right) \\ &= \frac{(2\pi)^{m/2} k^{-(N-m)}}{\sqrt{\det \alpha}} \exp\left(-\frac{\chi_0^2}{2k^2}\right) \end{aligned} \quad (\text{B3})$$

As a check, marginalization over k at this stage takes us back to equation (26).

This paper has been typeset from a $\text{T}_{\text{E}}\text{X}/\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$ file prepared by the author.