Bayesian Analysis (2011) **6**, Number 3, pp. 475–500

A Bayesian Surrogate Model for Rapid Time Series Analysis and Application to Exoplanet Observations

Eric B. Ford*, Althea V. Moorhead† and Dimitri Veras‡

Abstract. We present a Bayesian surrogate model for the analysis of periodic or quasi-periodic time series data. We describe a computationally efficient implementation that enables Bayesian model comparison. We apply this model to simulated and real exoplanet observations. We discuss the results and demonstrate some of the challenges for applying our surrogate model to realistic exoplanet data sets. In particular, we find that analyses of real world data should pay careful attention to the effects of uneven spacing of observations and the choice of prior for the "jitter" parameter.

Keywords: model comparison, model selection, periodogram, numerical methods, exoplanets, radial velocities, transit timing variations

1 Motivation & Overview of Exoplanet Observations

Since the 1990s, nearly 500 extrasolar planets (or exoplanets) have been discovered around other stars in our galaxy, yet only a few have been observed directly. In all other cases, the planet's presence has been inferred from its influence on the light of its host star.

1.1 Doppler Observations

The most productive such method to date has been observing the Doppler shift of the star light due to the gravitational perturbations of the planet. For a single planet on a circular orbit, the Doppler signature arises from a sinusoidal variation in the star's velocity along the line of sight. More generally, a single planet causes a periodic variation in the stellar velocity ($v(t)$) that follows the shape predicted by Kepler's laws of planetary motion, which can usually be well-approximated by the first few terms of a Fourier expansion (Konacki and Maciejewski 1999). If the amplitude of the fundamental term (with frequency $f = 2\pi/P$, where P is the orbital period) in the Fourier expansion of $v(t)$ is K_0, then the coefficients of harmonic terms (with frequency nf) are of order $e^n K_0$, where e is the orbital eccentricity. Given the typical signal-to-noise of detections and the typical exoplanet eccentricity of less than 0.3 (and often less than 0.1), the use of as few as two terms in this Fourier expansion is often accurate to within observational

*Department of Astronomy, University of Florida, Gainesville, FL, mailto:eford@astro.ufl.edu
†Department of Astronomy, University of Florida, Gainesville, FL, mailto:altheam@astro.ufl.edu
‡Department of Astronomy, University of Florida, Gainesville, FL and Institute of Astronomy, University of Cambridge, Cambridge, UK, mailto:veras@ast.cam.ac.uk

 DOI:10.1214/11-BA619

uncertainties. If a star has multiple planets, its Doppler signature can be much more complex. In some cases, the observed stellar velocities (v_k) can be well approximated as the linear superposition of multiple planets on Keplerian orbits. In other cases, the planet-planet interactions cause effects comparable to the overall amplitude of the signal and thus must be considered. In any case, the observational signature is quasi-periodic. A common question is whether a planet model with N_p planets is sufficient to describe the available observations or whether the data demand at least N_p+1 planets. While an exploration of the full physical parameter space would be computationally prohibitive, a lower dimensional surrogate model can be quite useful for analyzing such a system (Veras et al. 2011).

1.2 Transit Timing Variations

Recently, an extremely promising new method of detecting exoplanets has burst on to the scene. As a planet passes in front of its host star (i.e., a transit), the star's brightness appears to decrease. In an idealized two-body system, the mid-times of the transits (t_k) are strictly periodic at the orbital period (P), so the kth transit occurs at time $t_k = t_0 + k \times P$. If there are additional planets, then the times of the transits deviate from a linear ephemeris (Agol et al. 2005; Holman and Murray 2005; Ford and Holman 2007). The perturbations of an additional planet can cause deviations of the transit times that are simply sinusoidal in time, a periodic non-sinusoidal pattern or a very complex quasi-periodic waveform (Figures 1 and 2; Ford and Holman 2007; Nesvorný and Morbidelli 2008; Nesvorný 2009; Nesvorný and Beaugé 2010). While it is impractical to explore all the parameters of a full physical model, a lower dimensional surrogate model may be able to help identify regions of parameter space that merit further investigation with a full physical model.

1.3 Relation to Previous Research

In this manuscript, we present a new method for analyzing periodic time series data using a computationally efficient Bayesian surrogate model. The details of our model are chosen to facilitate the analysis of exoplanet observations. We test our model by analyzing Doppler and transit timing data sets. Thanks to the computational efficiency of our algorithm, it is possible to apply it to a large library of simulated data sets to understand how the model performs for different types of planetary systems.

Astronomers routinely apply Markov chain Monte Carlo (MCMC) techniques to perform Bayesian parameter estimation when analyzing Doppler observations of an exoplanet host star (e.g., Ford 2005; Gregory 2005). For multiple planet systems, MCMC methods are computationally intensive, even when the model evaluation is performed ignoring the gravitational interaction between the planets. While this is a good approximation for analyzing the Doppler observations of many systems, mutual planetary interactions can be quite significant for planetary systems near a mean motion resonance (e.g., planets b and c in the GJ 876 planetary system; Laughlin et al. 2005). Full n-body integrations to account for these mutual planetary interactions are much more computa-

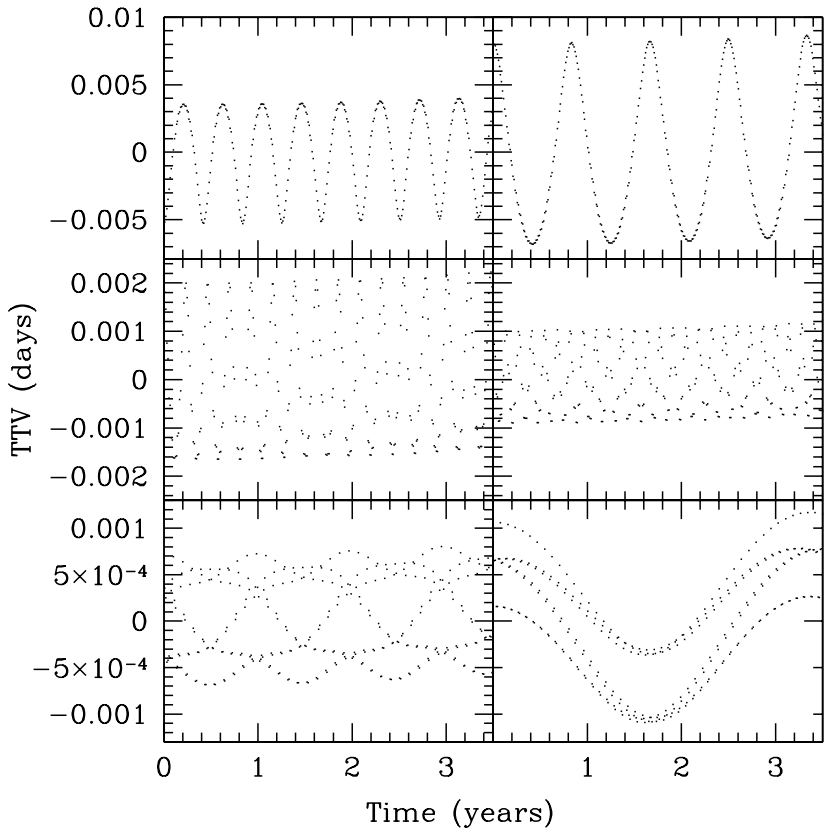

Figure 1: In this figure, we show the transit timing signature for a series of two planet systems, each extending for 3.5 years, the nominal mission lifetime for NASA's *Kepler* space observatory. In each case, transit times are for a 20 Earth-mass planet following an initially circular orbit with a star-planet separation of approximately 0.05AU. (One AU is the average distance between the Earth and the Sun). In each case, there is an additional 2 Earth-mass planet (which we assume is not observed to transit) following a slightly eccentric orbit ($e_2 = 0.1$) and an initial mean star-planet separation of 0.078AU (upper left), 0.080AU (upper right), 0.080AU (middle left), 0.082AU (middle right), 0.084AU (lower left) or 0.086AU (lower right). These separations are near the location of the 1:2 mean motion resonance ($\simeq 0.0794$AU). Note that the vertical axis scale changes from row to row. Even these small changes in the orbital separation result in qualitative changes in short- and long-term structure of the transit timing variations. Here we show simulated data from full n-body integrations with no data gaps and no measurement uncertainties. In practice, even *Kepler* misses some transits (e.g., due to data downlink with Earth, spacecraft abnormalities) and the transit timing measurements have noise of $\sim 10^{-3} - 10^{-2}$ days, depending primarily on the brightness of the target star and the size of the planet. For some planetary systems (e.g., Kepler-9 b&c, Holman et al. 2010) or triple star systems (Carter et al. 2011; Slawson et al. 2011; Steffen et al. 2011), the amplitude of the transit (eclipse) timing variations is much larger than *Kepler's* measurement uncertainties. For other planetary systems (e.g., Kepler-11 b-f, Lissauer et al. 2011b), the magnitude of transit timing variations are comparable to the measurement uncertainties. We expect the Bayesian approach and our surrogate model to be most useful for such systems, once a sufficient number and timespan of observations have been collected. For many systems with no detectable transit timing variations, statistical methods such as those described here will play an important role in establishing the significance of non-detection and the implied upper limits for the mass of any perturbing planets (e.g., Steffen and Agol 2005).

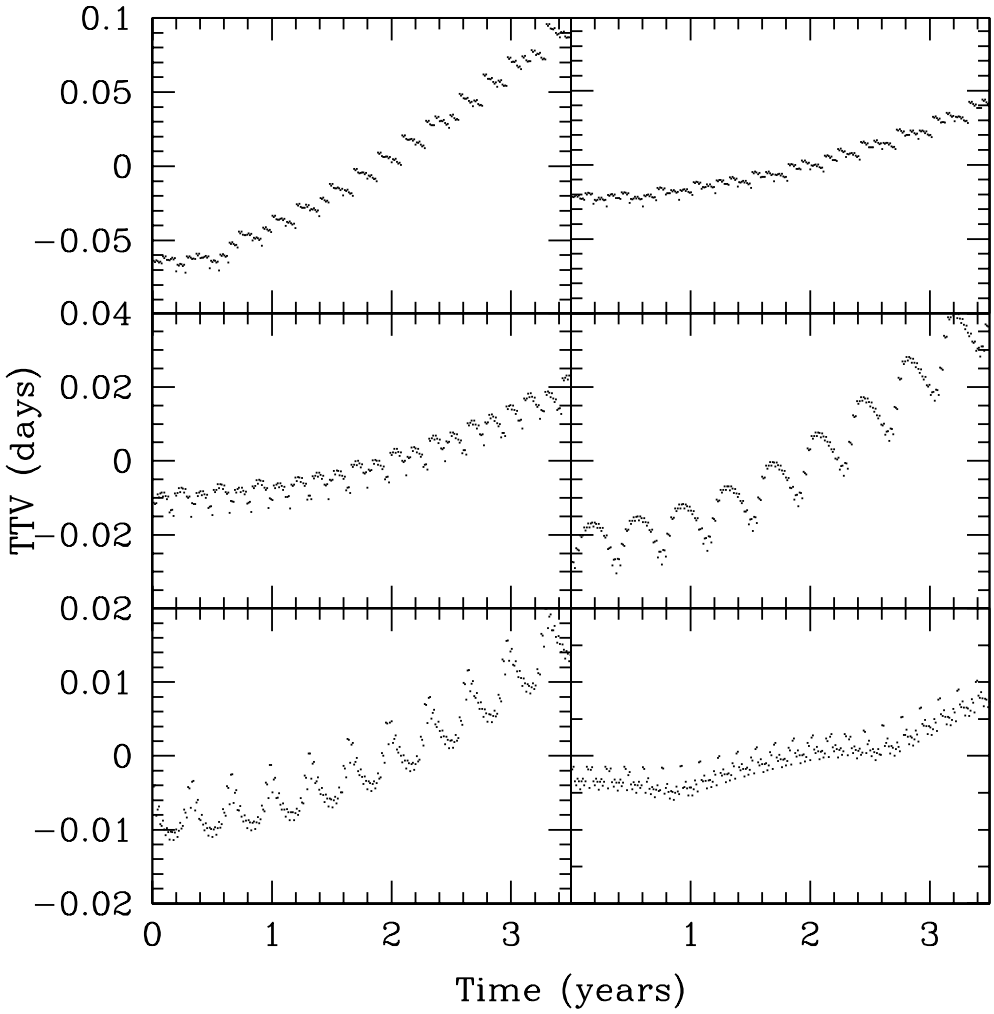

Figure 2: This figure is similar to Figure 1, except the the outer planet has been moved closer to the transiting planet. In each case, there is a 2 Earth-mass planet (which we assume is not observed to transit) following a slightly eccentric orbit ($e_2 = 0.1$) and an initial mean star-planet separation of 0.062AU (upper left), 0.063AU (upper right), 0.064AU (middle left), 0.065AU (middle right), 0.066AU (lower left) or 0.067AU (lower right). These separations are near the location of the 2:3 mean motion resonance ($\simeq$ 0.0655AU). Note that the vertical axis scale changes from row to row. Again, even these small changes in the orbital separation result in qualitative changes in short- and long-term structure of the transit timing variations.

tionally demanding. One approach is to develop and parallelize computationally efficient methods that allow one to use full n-body integrations (e.g., Johnson et al. 2011). While the GJ 876 data set has many high signal-to-noise observations, most exoplanet host stars have fewer observations and lower signal-to-noise, resulting in weak constraints on many physical model parameters and making it even more difficult to sample these parameter spaces using MCMC. The use of a lower-dimensional surrogate model has the potential to contribute to the analysis of such systems by identifying the periodicities that are statistically significant without introducing several additional parameters that are poorly constrained. Similar methods are routinely applied in a frequentist context to identify planet candidates from Doppler observations (Cumming et al. 2008). Our Bayesian surrogate model can be thought of as a Bayesian generalization of the Lomb-Scargle periodogram (Cumming 2004) that has been further generalized to allow for multiple frequencies, perhaps due to the perturbations from additional planets or perhaps due to significant eccentricity (Konacki and Maciejewski 1999). Previously, a much more restricted version of the surrogate model ($N_{f,\max} = 1$, $N_{d,\max} = 0$, see §2.2) was used to evaluate strategies for scheduling Doppler observations (Ford 2008). The generalization in this manuscript allows for identifying multiple periodicities, as is necessary for application to eccentric and/or multiple planet systems.

We are optimistic that the surrogate model has even more potential for contributing to the analysis of transit timing variation data. In the transit timing variation method, the entire signal is due to the mutual gravitational perturbations. Given the highly non-linear nature of the problem (particularly near resonances), a physical model requires performing computationally expensive n-body integrations. While it might be practical to perform MCMC sampling around one mode of the posterior distribution while using full n-body integrations, it is not practical to perform a brute-force global search of the high-dimensional parameter space while using full n-body integrations (Veras et al. 2011). The evaluation of our surrogate model is orders of magnitude faster than an n-body integration. Additionally, the surrogate model is linear in most of its model parameters, allowing for efficient identification of the modes and integration over linear parameters, once we condition on the non-linear parameters. (We perform integration over non-linear parameters via brute force, as described in the supplementary materials.) The speed of the surrogate model makes it well-suited to exploring a broad range of possible orbital configurations. Once the surrogate model identifies significant periodicities, n-body integrations can be used to perform a more detailed exploration of the full physical models in regions which could produce the periodicities identified by the surrogate model. We present results of our Bayesian surrogate model applied to simulated transit timing data and discuss the implications of our results for the prospects of transit timing-based planet searches.

2 Bayesian Surrogate Model for Analysis of Quasi-Periodic Time Series

2.1 Data

First, we describe our general model for analysis of quasi-periodic time series. In the case of Doppler observations, the independent variable (x) would correspond to time (t) and y would correspond to the star's velocity ($v(t)$). In the case of transit timing observations, x would correspond to the transit number and y would correspond to the mid-transit time. Each observation (y_k) is accompanied by an estimate of the measurement uncertainty (σ_k). The independent variables (i.e., transit number or time of each observation) are assumed to be known precisely.

2.2 Model

We explore the use of a surrogate model given by

$$y(x_k;\theta) = \sum_{i=1}^{N_f} \left[S_i \sin\left(2\pi\ f_i x_k\right) + C_i \cos\left(2\pi\ f_i x_k\right)\right] + \sum_{i=0}^{N_d} D_i x_k^i, \tag{1}$$

where x_k is the independent variable for the kth observation and and $y(x_k;\theta)$ is the predicted value of the observable based on the surrogate model with parameters θ. In the final sum, the superscript i is the exponent. The surrogate model parameters θ include 1) the number of frequencies N_f in the surrogate model, 2) the frequencies f_i, 3) the amplitudes S_i and C_i for those periodicities, 4) the order N_d of the polynomial terms, and 5) the polynomial coefficients D_i. This model exactly describes time series which are the superposition of polynomial and sinusoidal signals. The surrogate model can be used for Bayesian model comparison to determine how complex a model (i.e., how many frequencies and/or polynomial terms) is justified for a given data set.

In principle, one could consider alternative basis functions. We favor the use of sinusoids since they exactly describe the gravitational perturbation of a planet following a circular orbit (and non-interacting with other planets). Further, a sinusoid often provides a good first approximation to the perturbation of a planet on an eccentric orbit, given typical eccentricities and measurement precision. In practice, we will truncate the model to use just a few frequencies and polynomial terms, so as to provide an acceptable model for the observations while facilitating the efficient evaluation of the model. If observations span a sufficiently long period of time, then there is no need for polynomial terms. In practice, the orbital period of an outer planet (e.g., Saturn 29.5 years, Neptune 165 years) can be much longer than the time span of observations (typically ∼1-10 years). If the orbital period is much longer than the time span of observations then a simple linear term (constant radial acceleration for Doppler data) is all that can be discerned from the available data. In the case of a planet on a circular orbit, the "jerk" (i.e., derivative of acceleration) becomes significant as the time span of observations approaches a quarter of the orbital period. For the general case of a planet

on an eccentric orbit, the jerk can become evident after a greater or lesser duration, depending upon the orientation of the orbit and phase of the observations. In either case, we cannot infer the physical parameters with just two derivatives measured, as these are insufficient to characterize an orbit that requres more parameters to fully specify (seven physical parameters in general, and at least three even if we arbitrarily assume a circular, coplanar orbit). In these cases, using only one or two polynomial terms dramatically increases the computational efficiency, since they eliminate the need to explore a large number of frequencies. Another advantage of the polynomial basis is that a simple linear term corresponds to the Doppler perturbation from a body sufficiently distant that its orbital motion is insignificant over the time span of observations. A quadratic term corresponds to physical motion of the perturber over the time span of observations.

In cases where the time span of observations is sufficient to discern a third polynomial term, it becomes more difficult to justify a polynomial model. On one hand, a cubic term is still much less computationally demanding than considering an additional harmonic term (see §2.5). On the other hand, according to our physical model, all signals are periodic on sufficiently long timescales and three derivatives are sufficient to constrain the range of plausible periods and amplitudes. While three derivatives are not sufficient to infer the shape of the waveform, in many cases (e.g, a distant planet on nearly circular orbit) a simple sinusoid is a reasonable first approximation. Therefore, we favor considering a model with an additional harmonic term over including three or more polynomial terms.

2.3 Likelihood

We take the likelihood of each observation to be

$$L_k(x_k, y_k, \sigma_k|\theta) = N(y_k - y(x_k;\theta), \sigma_k^2 + \sigma_j^2) \tag{2}$$

where $N(\mu, \sigma^2)$ is a normal distribution with mean μ and variance σ^2. Here $\theta = \left\{N_f, f_1, ...f_{N_f}, S_1, ...S_{N_f}, C_1, ...C_{N_f}, N_d, D_0, ...D_{N_d}, \sigma_j\right\}$. Note that we have expanded the set of model parameters (θ) to include σ_j, a "jitter" parameter that is related the amount of scatter that is not accounted for by the observational uncertainties. The origin of the jitter need not be specified. It could be due to inaccurate estimates of the observational uncertainties or physical effects that are not included in a particular model (e.g., star spots, p-modes, additional planets).

We assume that the observational errors are uncorrelated in time, so the likelihood for a data set of N_{obs} observations is given by

$$L(\theta) = \prod_{k=1}^{N_{obs}} L_k(x_k, y_k, \sigma_k|\theta)\,. \tag{3}$$

2.4 Priors

In principle, one might aim to choose priors that are based on the distribution of masses and orbital periods for exoplanets. Of course, characterizing those distributions is one of the primary motivations to conduct exoplanet searches. While exoplanet searches have detected hundreds of planets, most of these have masses comparable to Saturn or greater and orbital periods of a few years or less (Cumming et al. 2008). Inevitably, astronomers pushing the frontiers of knowledge (e.g., searching for less massive planets) will not know the intrinsic distribution of those planets' orbital properties. Thus, we choose broad priors based on physical intuition and mathematical principles, as outlined in this section.

For the number of frequencies in the surrogate model, we adopt the following prior: $p(N_f) = (1 - 2\alpha + \alpha^N_{f,\max})/(1-\alpha)$ for $N_f = 0$, $p(N_f) = \alpha^{N_f}$ for $0 < N_f \leq N_{f,\max}$, and $p(N_f) = 0$ for $N_f > N_{f,\max}$, where α parametrizes our prior belief about the likelihood of multiple frequencies in the signal. The maximum number of frequencies to be considered ($N_{f,\max} \geq 1$) is chosen so as to provide enough complexity to model the data while keeping the model evaluation practical. The choice of a geometric probability for an increasing number of frequencies has the advantage that the prior ratio $p(N_f = n+1)/p(N_f = n)$ is independent of n (as long as $n + 1 \leq N_{f,\max}$). If we associate each frequency with one planet and consider planet detection serially (i.e., first look for evidence of one planet, next look for evidence of another planet), then the minimum Bayes factor to have a significant detection of each successive planet is constant.

The prior for the frequencies themselves are given by $p(\log f_i) \sim U(\log f_{\min}, \log f_{\max})$, where $U(a, b)$ is a uniform distribution between a and b. For our applications, there are physical limits on the range of viable frequencies. For example, $f_{\max}$ could be set by the shortest orbital period in which a planet would be able to survive for the age of the star, and $f_{\min}$ could be set by the longest orbital period in which a planet would be able to remain bound for the age of the star given perturbations from the galactic tidal field and passing stars. In most cases, the time span of observations (T_{obs}) will not be sufficient to distinguish such long-period signals from a low-order polynomial. Thus, we typically set $f_{\min} \sim 2/T_{obs}$, as the surrogate model is still able to model slow variations and this aids in the rapid evaluation of the surrogate model. Our choice of a prior that is flat in the logarithm of the frequency is motivated by the maximum entropy prior for a scale parameter. We also experimented with a modified Jeffreys prior. We found no significant difference in the results, as long as there is a significant detection. This can be understood simply in terms of the characteristic width of the peaks in the likelihood which are much smaller than the domain of the frequencies. Thus, as long as the prior for frequency has significant support across the entire domain, the choice of the prior for frequency has minimal impact on the shape of the posterior (in locations with significant posterior probability) unless there is not significant evidence for any periodicities. While we do not attempt to justify our choice of prior based on the distribution of exoplanet orbital periods, we note that this alternative approach would result in a fairly similar choice of prior, at least for frequencies between $\sim$1/(2 days) and $\sim$1/(2000 days) (Cumming et al. 2008).

Each of the frequencies (f_i) has an amplitude $A_i = \sqrt{S_i^2 + C_i^2}$ and a phase $\phi_i = \mathrm{atan2}(-S_i, C_i)$. We choose a uniform prior for each phase, $p(\phi_i) = 1/(2\pi)$. Physically, this corresponds to time invariance; i.e., the other planetary system does not know what time we choose to label as $t = 0$. For A_i, the total amplitude at frequency i, we adopt a modified Jeffreys prior, $p(A_i) \sim (1 + A_i/A_o)^{-1}$ for $0 \leq A_i \leq A_{\rm max}$, where $A_{\rm max}$ is the maximum plausible amplitude. For application to Doppler observations, $A_{\rm max}$ could be set by the maximum velocity perturbation by a planet (or binary star). For application to transit timing variations, $A_{\rm max}$ could be set by the orbital period. Note that, in principle, $A_{\rm max}$ could be a function of f_i. The other parameter, A_o, prevents a divergence of the prior at small amplitudes. For some applications A_o may be physically motivated. For our applications, we choose A_o based on the minimum detectable signal based on the available data set, e.g., $A_o \sim \left\langle 1/\sigma_i^2 \right\rangle^{-1/2}$, where $\left\langle 1/\sigma_i^2 \right\rangle^{-1/2}$ is effectively the weighted average measurement precision. In practice, we find that any sufficiently small choice of A_o and large choice of $A_{\rm max}$ gives similar results for parameter estimation for a given N_f. The choice of A_o, $A_{\rm max}$ and α do affect the marginalized posterior probability ratio for N_f.

For our typical applications, the total signal amplitude is proportional to the mass of the planet. While we do not attempt to justify our choice of prior based on the distribution of exoplanet masses, we note that this alternative approach would result in a fairly similar choice of prior, at least for readily detectable planets (e.g., more massive than Saturn for Doppler surveys). Present observations are only beginning to provide significant constraints on the distribution of planet masses for low-mass planets at small short orbital periods ($\leq$50 days) (Howard et al. 2010; Borucki et al. 2011; Youdin 2011).

We implement the above by using a "two-dimensional modified Jeffreys prior" for each pair of amplitudes S_i and C_i,

$$p(S_i, C_i) = 1/\left[4\pi \sqrt{S_i^2 + C_i^2} \log\left(1 + A_{\rm max}/A_o\right) \left(1 + \sqrt{S_i^2 + C_i^2}/A_o \right) \right] \quad \text{for} \sqrt{S_i^2 + C_i^2} \leq A_{\rm max}, \tag{4}$$

and $p(S_i, C_i) = 0$ for $\sqrt{S_i^2 + C_i^2} > A_{\rm max}$. Inspection of Eqn. 4 shows that it is essentially a modified Jeffreys prior for the total amplitude A_i.

For the number of polynomial terms in the surrogate model, we adopt a prior similar to that for the number of frequencies: $p(N_d) = (1 - 2\beta + \beta^{N_{d,\rm max}-1})/(1-\beta)$ for $N_d = 1$, $p(N_d) = \beta^{N_d - 1}$ for $1 < N_d \leq N_{d,\rm max}$, and $p(N_d) = 0$ for $N_d > N_{d,\rm max}$, where β parametrizes our prior belief about the likelihood of higher order polynomial terms being present in the signal. For our applications, the maximum polynomial order considered ($N_{d,\rm max} \geq 1$) is typically set to 1 and rarely more than 2. The choice of a geometric probability for an increasing number of polynomial terms has the advantage that the prior ratio $p(N_d = n+1)/p(N_d = n)$ is independent of n (as long as $n + 1 \leq N_{d,\rm max}$).

For both the jitter and the polynomial coefficients, i.e., each of $B \in (D_i,\, \sigma_j)$, we adopt a modified Jeffreys prior,

$$p(B) = 1/\left[2\left|B\right| \log\left(1 + B_{\rm max}/B_o\right) \left(1 + \left|B\right|/B_o\right) \right], \qquad \left|B\right| \leq B_{\rm max}, \tag{5}$$

where B_o and $B_{\rm max}$ are analogous to A_o and $A_{\rm max}$. We adopt the same values as for A_o and $A_{\rm max}$. For some data sets, it may be wise to choose a smaller B_o, as the minimum measurable magnitudes for the D_i depend on the effective measurement precision, the number of observations and the time-span of observations. In many cases, the amplitude of long-term trends is proportional to the perturbing planet's mass, motivating a prior $p(D_i)$ that is concentrated at small signals and that has a similar shape to the prior for the amplitudes. We also tried using a uniform prior, $p(D_i) = 1/(2A_{\rm max})$ for $|D_i| \leq A_{\rm max}$. In practice, the likelihood is very sharply peaked in D_1, so results are not sensitive to the choice of its prior. We have only begun to experiment with the surrogate model for $N_d = 2$ (based on discussion at end of §2.2), and caution that further experimentation may be needed for models with $N_d \geq 2$.

In this manuscript, we present results based on a modified Jeffreys prior, as the jitter is a scale parameter and we have physical reason to limit $p(\sigma_j)$ at small values of σ_j, once it becomes small compared to the measurement precision or astrophysical effects that can cause non-gravitational perturbations to the Doppler or transit timing signal (e.g., non-uniform star spots). However, we found that in some cases our results could be sensitive to the prior for the jitter (Payne and Ford 2011). For example, either imposing an upper cutoff on the prior for σ_j at ~ 10m/s or using a normal distribution for $p(\sigma_j)$ can significantly increase the posterior probability for $N_f = n + 1$ relative to that of $N_f = n$ or narrow the range of allowed amplitudes. Previous studies of the empirical distribution of Doppler jitter are based on relatively small samples sizes (Wright 2005), so they have little to say about the tails of the distribution. Based on our results, we encourage further observations and statistical analyses that could inform the choice of prior for σ_j.

2.5 Numerical Evaluation of Model and Posterior PDF

The surrogate model was designed to provide a good approximation to Doppler or transit timing observations and is likely to provide a reasonable approximation of many other time series. A second important feature of the surrogate model is that it permits efficient evaluation. Several tricks to perform an efficient brute force integration are described in the supplementary material. A key feature of the model is that for given values of N_f, N_d, the f_i and σ_j, the model is linear. Thus, the integration over the remaining parameters can be performed via linear algebra and the Laplace approximation. Marginalizing over σ_j and the f_i can be performed using standard numerical techniques. By evaluating the model conditioned on N_f and N_d and marginalizing over the remaining parameters, one can compare the marginal posterior probabilities to identify values of N_f and N_d that provide a good model for the observations without introducing more model parameters than are justified given the available data. For our typical applications, the posterior is dominated by small values of N_f and N_d, so higher values need not be considered explicitly. Since the surrogate model can be integrated numerically, it provides a quantitative basis for Bayesian model comparison and/or model selection.

3 Discussion of Application to Exoplanet Observations

3.1 Doppler Data

One common challenge for exoplanet searches is deciding when the available data provides sufficient evidence to constitute a planet discovery. This is particularly challenging in the case of stars with multiple planets, as the data can always be better modeled by adding additional planets. The Fourier decomposition of the Doppler signature of a planet is dominated by power at the frequency $(1/P)$ corresponding to the orbital period (P). The power at the harmonic frequencies, $2/P$, $3/P$, ..., λ/P, decreases as $e^{\lambda-1}$, where e is the orbital eccentricity. The eccentricity for an ellipse is constrained to $(0, 1)$ and most known exoplanets have eccentricities smaller than 0.15. Thus, applying the Bayesian surrogate model to Doppler data sets is expected to result in a posterior distribution for the most significant frequency (f_1) near one over the orbital period of the planet which dominates the Doppler signature. The second most prominent frequency (f_2) could correspond to a harmonic of the first planet or to the orbital period of a second planet. Some planetary systems contain two planets with orbital periods that differ by a factor of nearly 2 (e.g., Laughlin et al. 2005; Wright et al. 2011; Lissauer et al. 2011b). For certain planet masses and eccentricities, the Doppler signature of two such planets can be mimicked by one planet with a more eccentric orbit (Giuppone et al. 2009; Anglada-Escudé et al. 2010; Moorhead and Ford 2010). If there is actually only one planet, then one would expect the next most prominent frequency to correspond to twice the fundamental frequency ($f_2 = 2 \times f_1$) to within measurement precision. On the other hand, for systems that actually contain two planets, the sidereal orbital periods often deviate from exact resonance (Lissauer et al. 2011b). One might hope that in cases where the ratio of orbital periods differed from two, that the surrogate model could recognize this difference ($\delta \equiv f_2 - 2 \times f_1$), so one could infer that the observations were due to two planets, rather than one planet on a more highly eccentric orbit (Giuppone et al. 2009; Anglada-Escudé et al. 2010; Moorhead and Ford 2010).

To explore this possibility, we applied the Bayesian surrogate model to several sets of Doppler observations of exoplanet host stars. We focused on exoplanets with large velocity amplitudes and believed to have a significant eccentricity, as those systems provide the best prospects for measuring the harmonic frequencies precisely. In particular, we choose systems with $Ke^2 > 3$m/s, where K is the velocity amplitude and e is the orbital eccentricity. In all cases, the surrogate model efficiently found the fundamental frequency corresponding to the orbital period. As expected, f_1 is very tightly constrained and the marginalized posteriors for additional frequencies are significantly broader. In some cases, we found that the marginalized posterior for f_2 did not correspond to $2 \times f_1$.

In order to determine how often f_2 deviated from $2 \times f_1$ by chance, we performed a similar analysis on several simulated data sets. Each simulation was modeled on simulated velocities that were calculated according to the best-fit orbital period, amplitude, eccentricity, arrangement of pericenter and orbital phase. We added Gaussian measurement noise with a scale set by the claimed measurement uncertainty. The results for

one case (HD 162020) are shown in Figures 3 and 4. If we use the same observation times and uncertainties as the actual observations, then the marginal posterior distributions for $P_2 = 1/f_2$ and $P_1/2 = 2/f_1$ do not overlap. If we generate a similar data set, but with random observation times, then the marginal posterior distributions for $P_2 = 1/f_2$ and $P_1/2 = 2/f_1$ do overlap. We conclude that one must be very cautious of aliasing due to unevenly spaced observations when interpreting the marginal posterior distributions for the f_i for realistic data sets.

Our results suggest that the problems of aliasing could be reduced by obtaining regularly (or randomly) spaced observations. Unfortunately, this is impractical for observations made from the surface of the Earth. First, observations can not be made when the Sun is above the horizon (or even less than $\sim 12^\circ$ below the horizon) due to scattering of sunlight by Earth's atmosphere. As the Earth revolves around the Sun, the time of day at which a given star can be observed from a given site changes. For most stars, there are multiple months each year when high-precision Doppler observations are not possible, since the star appears too close to the sun (after projecting onto the sky). Thus, for most stars, Doppler observations are prone to aliasing at frequencies associated with the solar day and the solar year. (In principle, observations of stars near the North or South pole from an observatory in the Arctic or Antarctic could avoid aliasing near the day. However, there are no observatories with high-precision Doppler capability near either pole due to logistical issues.) Second, while planet host stars are faint compared to the daytime sky brightness, they are much brighter than distant galaxies. Therefore, time allocation committees assign the vast majority of observing time near new Moon to extragalactic astronomers. Exoplanet searches are typically assigned observing time near full Moon, introducing aliasing at frequencies associated with the lunar month ($\simeq$ 29.5 days). In practice, the best way to avoid this is to dedicate an observatory to Doppler observations (e.g., the HARPS instrument at the European Southern Observatory's 3.6m telescope in La Silla, Chile). Of course, this requires considerable resources and is not an option for the world's largest telescopes. Third, astronomers often attempt to optimize the efficiency of their observations on a given night by observing each star when it is near the greatest (angular) altitude in the sky that night, as this minimizes the amount of absorption of starlight by the Earth's atmosphere. This strategy introduces yet another frequency associated with the sidereal day (i.e., the time for a star to return to the nearly same point on the sky from a given location, roughly 23 hours, 56 minutes and 4 seconds). While these may seem like picky details, each of these timescales can be found in public data and in some cases contribute to qualitative ambiguities in the orbital solutions (Dawson and Fabrycky 2010). While attention to scheduling can help improve efficiency of planet searches, ultimately weather (i.e., cloudy skies) will lead to data gaps and prevent optimal experimental design for any Earth-based observatory. Space-based observatories are orders of magnitude more expensive to construct and operate. Therefore, we must develop tools to analyze realistic data sets. Studies such as this will help us to use those tools responsibly and reduce the risk of making erroneous claims.

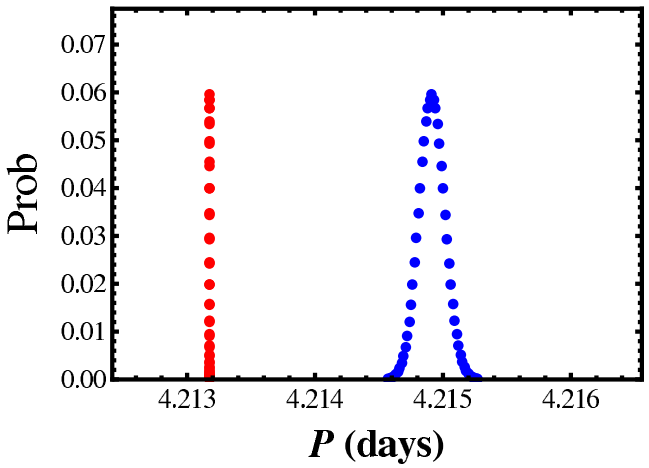

Figure 3: In this figure, we consider a simulated data set based on the best-fit orbital parameters for the exoplanet HD 162020. We show the marginalized posterior probability distributions for $P_1/2 = 2/f_1$ (red; very narrow distribution) and $P_2 = 1/f_2$ (blue; broad distribution) from a surrogate model with $N_f = 2$. Based on the Fourier expansion of the Doppler signature of a planet on a Keplerian orbit, one would expect that the surrogate model would yield marginalized posterior probability distributions for f_2 which overlap marginalized posterior probability distributions for $f_1/2$. We find that this is not necessarily the case, even for this simulated data set with high signal-to-noise, uncorrelated Gaussian measurement errors and accurate estimates of the measurement uncertainties. For this calculation, we have used the actual times of observations.

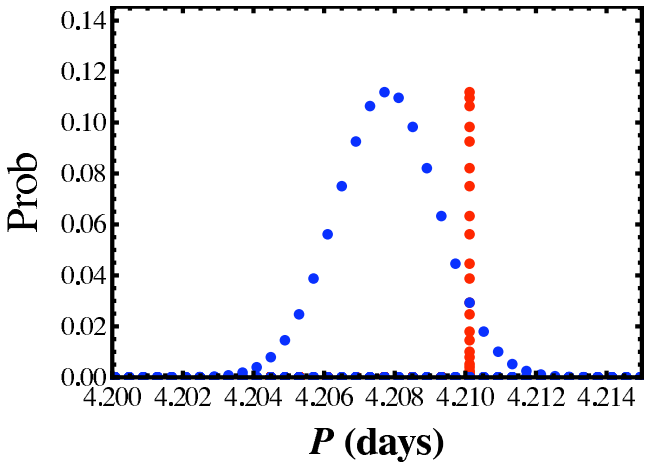

Figure 4: This figure is analogous to Figure 3, except that the observation times are chosen randomly. In this case, the posterior for $P_2 = 1/f_2$ (blue; broad distribution) overlaps with the posterior for $P_1/2 = 2f_1$ (red; narrow distribution), as expected for a Keplerian orbit. Contrast this with Figure 3 which uses actual observation times. We conclude that for realistic Doppler data sets, aliasing due to the limited number of the unevenly spaced observations can result in the marginal posterior distributions for $f_1/2$ and f_2 not overlapping. This poses a considerable challenge to testing the one-planet model based on the dominant frequencies.

3.2 Transit Timing Variations

The rapidly increasing number of known exoplanets that transit their host star has led several observers to measure transit times in hopes of detecting deviations from a linear ephemeris due to perturbations by another (potentially non-transiting) planet (e.g., Miller-Ricci et al. 2008; Maciejewski et al. 2010). In most previous studies, the number of precisely measured transit times has been insufficient to fit a physical model. The prospects for the transit timing variation method are poised to improve dramatically thanks to NASA's *Kepler* mission, which is observing over 100,000 stars nearly continuously for 3.5 years (Borucki et al. 2011; Steffen et al. 2010). Indeed, *Kepler* recently reported the first strong evidence for transiting timing variations in Kepler-9 and Kepler-11 systems (Holman et al. 2010; Lissauer et al. 2011a). Interestingly, in both of these cases, transit timing variations were used to confirm planet candidates that had already been identified via transit and to constrain their masses and orbits. *Kepler* has also identified dozens of planet candidates with putative transit timing variations (Ford et al. 2011). However, further observations are needed in order to infer the mass and orbital properties of the perturbing bodies.

In principle, one could model the full light curve (i.e., observed brightness versus time; Carter et al. 2011). However, this would vastly increase the computational time required and most of the information is contained in the transit time. (Using data from NASA's *Kepler* mission, there are typically $\sim 10^3 - 10^7$ brightness measurements for each transit time.) The transit time is the most precisely measured parameter for each transit, and the transit times are sensitive to whether the transiting planet is slightly ahead or behind "schedule" due to gravitational perturbations from other planets. The next best-measured parameter is the transit duration which depends on the orbit of the transiting planet and the stellar radius (Moorhead et al. 2011). The depth of each transit is primarily determined by the relative sizes of the planet and host star and is not affected by gravitational perturbations of other planets. The detailed shape of each transit also depends on detailed stellar properties ("limb darkening parameters"). For a system of non-coplanar planets, transit duration variations may be detectable. We have focused our analysis in this paper on coplanar systems viewed edge-on, so as to reduce the dimensionality of the parameter space to be explored. While more parameters are required to describe non-coplanar systems, in some cases it may be possible to derive additional constraints on the orbits based on transit duration variations, or lack thereof (Holman et al. 2010). Given the computational cost of modeling the full light curve, we recommend future research to develop tools to analyze a series of transit times and transit durations.

To explore the potential for transit timing variations to enable the detection of a non-transiting planet, we have generated $\sim 10^7$ simulated data sets with a wide variety of orbital periods, eccentricities and angles (Veras et al. 2011). We apply the surrogate model to simulated data sets to identify the dominant frequency and its amplitude. We find that the surrogate model can provide an accurate model for some data sets, particularly those very near a mean motion resonance, a regime which is particularly difficult to approximate analytically (Nesvorný and Morbidelli 2008; Nesvorný 2009;

Nesvorný and Beaugé 2010). In other cases, the transit timing variation signature is more complex and would require several frequencies to model adequately. For extended data sets, this can be challenging, both due to computation time and available memory.

While the surrogate model can provide a reasonable approximation to many transit timing signatures, the inferred model parameters depend sensitively on the orbital phases, as well as more basic physical parameters such as the planet mass and orbital period. Further, as one increases the number of observations, the inferred parameters often change significantly. This significantly complicates the interpretation of the surrogate model outputs. While the inferred model parameters will eventually stabilize with a sufficient number and time span of observations, we find that inferred parameters can continue to change even after several years of observations. We conclude that it may not always be practical to invert transit timing variations and infer the mass and orbital properties of a non-transiting planet (Ragozzine and Holman 2010). A more extended discussion of implications for transit timing planet searches is presented separately (Payne et al. 2010; Veras et al. 2011). Of course, our results do not prove that other analysis techniques can not invert transit timing variations. However, our surrogate model was designed to capture the most important aspects of the problem. Thus, our results are suggestive that this problem may be more general. Since the original submission of this paper, the first confirmations of exoplanets via the transit timing variation methods were published (Holman et al. 2010; Lissauer et al. 2011a). In these cases, each of the detected planets transit the star, so transit timing variations were used to confirm planet candidates that had already been identified by the standard transit technique. Ford et al. (2011) showed that *Kepler* can be expected to measure transit timing variations for at least 12 systems with multiple transiting planet candidates. Based on analysis of the frequency of multiple transiting planet candidate systems (Lissauer et al. 2011b), we expect that even more planets with transit timing variations will be significantly perturbed by a non-transiting planet. Indeed, based on the first four months of observations, *Kepler* has identified dozens of planet candidates with prospective transit timing variations, most in systems with only a single transiting planet candidate. Both our results and Ford et al. (2011) suggest that further observations will be necessary before the masses and orbits of putative additional planets can be determined.

3.3 Conclusions

We developed a Bayesian surrogate model for analysis of time series data in general and applied this model to two types of exoplanet search data, Doppler and transit timing variations. The surrogate model can be evaluated very rapidly and we describe a method for efficiently integrating over most model parameters. This allows for calculating the (properly normalized) marginalized posterior probability and assessing the posterior probability for a given periodicity.

One strength of the surrogate model is for exploratory data analysis. For example, astronomers routinely use the Lomb-Scargle periodogram to search Doppler data for a periodic signature of a planet and to identify the range of periods that should

be explored with a more detailed model. One limiting case of the surrogate model ($N_{f,\max} = 1$ and $N_{d,\max} = 0$) directly corresponds to the Bayesian generalization of the Lomb-Scargle periodogram (Cumming 2004). When a planet has a large eccentricity or one star hosts multiple planets, the Lomb-Scargle periodogram typically reveals multiple significant periodicities. Previously, astronomers have dealt with this by applying the Lomb-Scargle periodogram to the residuals after subtracting the best-fit sinusoidal or Keplerian model. This approach can bias subsequent results, since the subtracted model is not exact. Further, assessing the significance of peaks in the periodogram of residuals is nontrivial. Most authors use a blind approach when searching for additional periodicities, but others favor using information about the frequencies previously identified (e.g., Konacki and Maciejewski 1999; Dawson and Fabrycky 2010). In practice, this can lead to a cumbersome decision tree in a frequentist context. Our surrogate model (with $N_{f,\max} > 1$) represents a Bayesian generalization of iterative frequentist methods for analyzing periodogram of residuals. The Bayesian surrogate model provides a rigorous basis for calculating Bayes factors of the marginalized posterior probability for the number of significant frequencies $B_{n+1,n} = p(N_f = n+1|x, y, \sigma)/p(N_f = n|x, y, \sigma)$. Another advantage of the surrogate model is that by marginalizing over the other model parameters (e.g., frequencies, amplitudes, jitter, polynomial terms), a spurious false positive should be less likely than analyzing residuals to only the best-fit model.

Finally, for many systems the surrogate model can provide a lower-dimensional model that still captures the important (i.e., observable) physical effects. For example, in a system of multiple low-mass planets, a full physical model has a dimension of $\simeq 7N_p$, where N_p is the number of planets. If the system has planetary orbits with small eccentricities and/or inclinations, then several of the model parameters may have no observable effect. For such systems, the surrogate model would be able to describe the system accurately using a lower-dimensional parameter space ($\simeq 3N_p$), greatly increasing computational efficiency and perhaps increasing the sensitivity for detecting additional planets (due to the less extreme Occam's factor).

The surrogate model is not meant to replace other tools for Bayesian parameter estimation and model selection. It is still beneficial to apply MCMC (and variants) for parameter estimation using a more physical model (e.g., Ford 2005; Gregory 2005; Ford 2006; Johnson et al. 2011). Similarly, tools such as restricted Monte Carlo, importance sampling and parallel tempering can be helpful for calculating Bayes factors using a more physical model (e.g., Ford and Gregory 2007; Gregory 2011). Since these tools are much more computationally expensive than the surrogate model, they are most appropriate once a putative set of orbital periods has been identified (e.g., by periodogram analysis, surrogate model, or human inspection for sufficiently large signals).

Thanks to the computational efficiency of the surrogate model, we were able to analyze numerous simulated datasets including multiple planet systems, something that would not have been feasible using previously available Bayesian methods such as MCMC. For Doppler planet searches we find that realistic observing cadences can lead to significant aliasing that prevents precisely testing whether there is a harmonic relationship between measured frequencies. Given the close relationship of our surrogate model to the widely used Lomb-Scargle periodogram, our results also serve as caution

regarding the integration of results based on periodogram analyses (e.g., Konacki and Maciejewski 1999; Anglada-Escudé et al. 2010; Dawson and Fabrycky 2010).

For analyzing transit timing variations, we find the posterior distributions for the surrogate model parameters are sensitive to the exact orbital configuration. While the sensitivity to important physical parameters is advantageous, sensitivity to parameters that do not have dynamical significance makes interpretation of the posterior distribution for surrogate model parameters more challenging. Unfortunately, the transit timing signature often evolves on a timescale comparable to or longer than a realistic time span for observations (e.g., 3.5-10 years for *Kepler*). This makes it impractical to build a library of possible transit timing signatures and the corresponding surrogate model outputs. The surrogate model may still be useful for establishing the significance of putative periodicities and/or long-term trends in transit timing data. In addition to the advantages of a Bayesian approach, the computational efficiency of the surrogate model could be useful for analyzing large sets of simulated data sets to aid in interpretation of a transit timing variation planet search.

In both cases, we find that qualitative results (e.g., whether the Bayes factor, $B_{n,n+1}$, for the significance of an additional frequency is greater or less than unity) can depend on the choice of prior for the jitter parameter (σ_j). In this paper we used a mathematically motivated prior for σ_j. Our result suggests that practical application of the surrogate model would significantly benefit from further astronomical observations and statistical analyses to assess the empirical distribution of the stellar jitter for both Doppler (e.g., Wright 2005) and transit timing observations (e.g., Holman et al. 2010). Given the relationship of our surrogate model to non-Bayesian methods being employed by astronomers, our results also suggest caution in the interpretation of other results based on frequentist analyses (which typically assume a single fixed value of the jitter).

4 Supplementary Material

4.1 Practical Model Evaluation

Integration over Linear Parameters

The surrogate model is linear in the parameters S_i, C_i, and D_i. For a given choice of N_f, N_d, the f_i and σ_j, one can calculate the "best-fit" values of the S_i, the C_i, and the D_i via simple linear algebra. The same linear algebraic operations allow the integrals over these linear parameters to be evaluated efficiently via the Laplace approximation (i.e., we expand the exponent about the best-fit parameters, keep the constant and second order terms, and extend the limits of integration to infinity; Cumming (2004)). For our applications, the posterior probability is typically a smooth function of σ_j, but can vary extremely rapidly with the f_i. Thus, we recommend fixing the choice of N_f, N_d, and the f_i so that the integral over σ_j can be evaluated efficiently with standard numerical integration techniques. (In the special case of equal measurement uncertainties, the integrals can be calculated analytically.)

Brute Force Integration over Frequencies

Unfortunately, the integrals over the f_i must be evaluated numerically. We recommend discretizing these integrals and evaluating them via brute force. The number of frequencies to be evaluated (M) can be quite large (Cumming 2004). Fortunately, there are several computational tricks that can speed up the calculation. In particular, most of the trigonometric functions can be computed using trigonometric identities to improve performance. We find that brute force integration over one or even two f_i is practical for realistic data sets. (For large data sets, a large amount of RAM may be required for efficient evaluation.)

Approximations for Models with Many Frequencies

Unfortunately, brute force integration over more than two f_i rapidly become prohibitive. Thus, we introduce the following approximation. We start by performing a brute force evaluation of the model conditioned on there being one frequency and then successively approximate the posterior conditioned on two frequencies, three frequencies, etc. When approximating the model conditioned on there being N_f frequencies (and $N_f \geq 2$), we limit the set of f_i with $i < N_f$ that are considered to be those which contributed significantly to the posterior probability for the model conditioned on there being $N_f - 1$ frequencies. That is we set a threshold (e.g., $\epsilon = 10^{-4}$) and after evaluating the model conditioned on $N_f = 1$ frequencies, we store the m_1 frequencies which have the largest posterior probabilities and collectively sum to at least $1 - \epsilon$ of the posterior probability conditioned on there being 1 frequency. At this point, we have searched $M \times (1 + m_1)$ frequencies (rather than M^2). Next, we estimate the marginal posterior probability for $N_f = 3$ by considering only those combinations of f_1 and f_2 that have the largest posterior probabilities and collectively sum to at least $1 - \epsilon$ of the posterior probability conditioned on there being 2 frequencies. Thus, when evaluating the model conditioned on N_f frequencies, there are no more than $M \prod_{i=1}^{N_f - 1} m_i$ frequencies to be explicitly evaluated, much less than M^{N_f}. In principle, this approximation can be relaxed by performing a full search over two frequencies. In this case evaluating the model conditioned on N_f frequencies requires no more than $M^2 \prod_{i=1}^{N_f - 2} m_i$ frequencies to be explicitly evaluated, which may or may not be practical for a given data set.

Truncating the Number of Frequencies

For data sets which are well described by only one or a few frequencies, models which include a large number of frequencies will have a very small posterior probability due to the Occam's factor associated with the higher-dimensional model. Thus, there is little point in evaluating models with large N_f. Thus, we recommend successively calculating the posterior probability conditioned on N_f ($p(\theta|N_f, x_k, y_k, \sigma_k)$) and stopping at $N_{f,stop}$, such that $p(N_f = N_{f,stop}|x_k, y_k, \sigma_k) \ll \sum_{i=0}^{N_{f,stop}-1} p(N_f = i|x_k, y_k, \sigma_k)$
and approximate the remaining models as $p(N_f > N_{f,stop}|x_k, y_k, \sigma_k) \approx 0$.

Computational Cost

The surrogate model can be evaluated much more quickly than an n-body integration, explores a lower-dimensional parameter space, and takes advantage of the linearity of the model in most of the model parameters. Nevertheless, the strong sensitivity to frequency dictates that we must perform a fine sampling in frequency. For example, consider a case of a $\simeq 10 M_{Earth}$-mass transiting planet with an orbital period near 4 days and a small planet which is not observed to transit with a period near 8 days. With a modest eccentricity (0.1), the transit timing variations of the inner planet could be ~ 10 minutes, comparable to the timing precision for each transit for a typical *Kepler* planet host star (Ford et al. 2011). Over 7.5 years of observations (possible with an extended *Kepler* mission), one would observe roughly 680 transits, allowing for an easy detection of such a single. If we assume a timing precision of 10 minutes and set $f_{\rm max}$ to 1 day, then each integral over frequency requires considering $M \sim 180,000$ frequencies. With a single core of an AMD Opteron 275 processor (2.2 GHz), this takes $\sim$ 10, 16 or 37 seconds for models with $N_f = 1$, 2 or 3. Setting $\epsilon = 10^{-3}$ (as described above), we needed to compute 1 ($N_f = 1$), 11 ($N_f = 2$) and 16 ($N_f = 3$) scans over frequency for both $N_d = 0$ and 1, requiring a total of 12 CPU-minutes for one system. Thus, the brute force exploration is relatively fast for a single model. For systems with no detectable signal, the time required decreases significantly, as $N_f = 3$ (or even 2) need not be explored. On the other hand, the computation time per system required grows significantly as the signal-to-noise increases, since the number of frequencies sampled (M) must be increased to avoid missing a narrow peak in the posterior density. Of course, in these cases, the signal is sufficiently large that fancy statistical methods are not necessary to detect the dominant signal. The speed of the surrogate model allowed us to analyze millions of simulated data sets and to explore the complex parameter space, using a cluster with hundreds of AMD Opteron servers at the University of Florida High-Performance Computing Center.

References

Agol, E., Steffen, J., Sari, R., and Clarkson, W. (2005). "On detecting terrestrial planets with timing of giant planet transits." *Monthly Notices of the Royal Astronomical Society*, 359: 567–579. 476

Anglada-Escudé, G., López-Morales, M., and Chambers, J. E. (2010). "How Eccentric Orbital Solutions Can Hide Planetary Systems in 2:1 Resonant Orbits." *Astrophysical Journal*, 709: 168–178. 485, 492

Borucki, W. J., Koch, D. G., Basri, G., Batalha, N., Brown, T. M., Bryson, S. T., Caldwell, D., Christensen-Dalsgaard, J., Cochran, W. D., DeVore, E., Dunham, E. W., Gautier, T. N., III, Geary, J. C., Gilliland, R., Gould, A., Howell, S. B., Jenkins, J. M., Latham, D. W., Lissauer, J. J., Marcy, G. W., Rowe, J., Sasselov, D., Boss, A., Charbonneau, D., Ciardi, D., Doyle, L., Dupree, A. K., Ford, E. B., Fortney, J., Holman, M. J., Seager, S., Steffen, J. H., Tarter, J., Welsh, W. F., Allen, C., Buchhave, L. A., Christiansen, J. L., Clarke, B. D., Désert, J., Endl, M., Fabrycky, D., Fressin,

F., Haas, M., Horch, E., Howard, A., Isaacson, H., Kjeldsen, H., Kolodziejczak, J., Kulesa, C., Li, J., Machalek, P., McCarthy, D., MacQueen, P., Meibom, S., Miquel, T., Prsa, A., Quinn, S. N., Quintana, E. V., Ragozzine, D., Sherry, W., Shporer, A., Tenenbaum, P., Torres, G., Twicken, J. D., Van Cleve, J., and Walkowicz, L. (2011). "Characteristics of planetary candidates observed by Kepler, II: Analysis of the first four months of data." *Astrophysical Journal*, 736: 19–40. 483, 489

Carter, J. A., Fabrycky, D. C., Ragozzine, D., Holman, M. J., Quinn, S. N., Latham, D. W., Buchhave, L. A., Van Cleve, J., Cochran, W. D., Cote, M. T., Endl, M., Ford, E. B., Haas, M. R., Jenkins, J. M., Koch, D. G., Li, J., Lissauer, J. J., MacQueen, P. J., Middour, C. K., Orosz, J. A., Rowe, J. F., Steffen, J. H., and Welsh, W. F. (2011). "KOI-126: A Triply Eclipsing Hierarchical Triple with Two Low-Mass Stars." *Science*, 331: 562–565. 477, 489

Cumming, A. (2004). "Detectability of extrasolar planets in radial velocity surveys." *Monthly Notices of the Royal Astronomical Society*, 354: 1165–1176. 479, 491, 492, 493

Cumming, A., Butler, R. P., Marcy, G. W., Vogt, S. S., Wright, J. T., and Fischer, D. A. (2008). "The Keck Planet Search: Detectability and the Minimum Mass and Orbital Period Distribution of Extrasolar Planets." *Publications of the Astronomical Society of the Pacific*, 120: 531–554. 479, 482

Dawson, R. I. and Fabrycky, D. C. (2010). "Radial Velocity Planets De-aliased: A New, Short Period for Super-Earth 55 Cnc e." *Astrophysical Journal*, 722: 937–953. 486, 491, 492

Ford, E. B. (2005). "Quantifying the Uncertainty in the Orbits of Extrasolar Planets." *Astronomical Journal*, 129: 1706–1717. 476, 491

— (2006). "Improving the Efficiency of Markov Chain Monte Carlo for Analyzing the Orbits of Extrasolar Planets." *Astrophysical Journal*, 642: 505–522. 491

— (2008). "Adaptive Scheduling Algorithms for Planet Searches." *Astronomical Journal*, 135: 1008–1020. 479

Ford, E. B. and Gregory, P. C. (2007). "Bayesian Model Selection and Extrasolar Planet Detection." In G. J. Babu & E. D. Feigelson (ed.), *Statistical Challenges in Modern Astronomy IV*, volume 371 of *Astronomical Society of the Pacific Conference Series*, 189. 491

Ford, E. B. and Holman, M. J. (2007). "Using Transit Timing Observations to Search for Trojans of Transiting Extrasolar Planets." *Astrophysical Journal Letters*, 664: L51–L54. 476

Ford, E. B., Rowe, J. F., Fabrycky, D. C., Carter, J., Holman, M. J., Lissauer, J. J., Ragozzine, D., Steffen, J. H., Borucki, W. J., Batalha, N. M., Bryson, S., Caldwell, D. A., Gautier, T. N., III, Jenkins, J. M., Koch, D. G., Li, J., Lucas, P., Marcy, G. W., Mullally, F. R., Quintana, E., McCauliff, S., Thompson, S. E., Tenenbaum,

P., Still, M., and Twicken, J. D. (2011). "Transit Timing Observations from Kepler: I. Statistical Analysis of the First Four Months." *Astrophysical Journal Supplement*. 489, 490, 494

Giuppone, C. A., Tadeu dos Santos, M., Beaugé, C., Ferraz-Mello, S., and Michtchenko, T. A. (2009). "Detectability and Error Estimation in Orbital Fits of Resonant Extrasolar Planets." *Astrophysical Journal*, 699: 1321–1332. 485

Gregory, P. C. (2005). "A Bayesian Analysis of Extrasolar Planet Data for HD 73526." *Astrophysical Journal*, 631: 1198–1214. 476, 491

— (2011). "Bayesian Re-analysis of the Gliese 581 Exoplanet System." *Monthly Notices of the Royal Astronomical Society*, 415: 2523–2545. 491

Holman, M. J., Fabrycky, D. C., Ragozzine, D., Ford, E. B., Steffen, J. H., Welsh, W. F., Lissauer, J. J., Latham, D. W., Marcy, G. W., Walkowicz, L. M., Batalha, N. M., Jenkins, J. M., Rowe, J. F., Cochran, W. D., Fressin, F., Torres, G., Buchhave, L. A., Sasselov, D. D., Borucki, W. J., Koch, D. G., Basri, G., Brown, T. M., Caldwell, D. A., Charbonneau, D., Dunham, E. W., Gautier, T. N., Geary, J. C., Gilliland, R. L., Haas, M. R., Howell, S. B., Ciardi, D. R., Endl, M., Fischer, D., Fürész, G., Hartman, J. D., Isaacson, H., Johnson, J. A., MacQueen, P. J., Moorhead, A. V., Morehead, R. C., and Orosz, J. A. (2010). "Kepler-9: A System of Multiple Planets Transiting a Sun-Like Star, Confirmed by Timing Variations." *Science*, 330: 51–54. 477, 489, 490, 492

Holman, M. J. and Murray, N. W. (2005). "The Use of Transit Timing to Detect Terrestrial-Mass Extrasolar Planets." *Science*, 307: 1288–1291. 476

Howard, A. W., Marcy, G. W., Johnson, J. A., Fischer, D. A., Wright, J. T., Isaacson, H., Valenti, J. A., Anderson, J., Lin, D. N. C., and Ida, S. (2010). "The Occurrence and Mass Distribution of Close-in Super-Earths, Neptunes, and Jupiters." *Science*, 330: 653–655. 483

Johnson, J. A., Payne, M., Howard, A. W., Clubb, K. I., Ford, E. B., Bowler, B. P., Henry, G. W., Fischer, D. A., Marcy, G. W., Brewer, J. M., Schwab, C., Reffert, S., and Lowe, T. B. (2011). "Retired A Stars and Their Companions. VI. A Pair of Interacting Exoplanet Pairs Around the Subgiants 24 Sextanis and HD 200964." *Astronomical Journal*, 141: 16–26. 479, 491

Konacki, M. and Maciejewski, A. J. (1999). "Frequency Analysis of Reflex Velocities of Stars with Planets." *Astrophysical Journal*, 518: 442–449. 475, 479, 491, 492

Laughlin, G., Butler, R. P., Fischer, D. A., Marcy, G. W., Vogt, S. S., and Wolf, A. S. (2005). "The GJ 876 Planetary System: A Progress Report." *Astrophysical Journal*, 622: 1182–1190. 476, 485

Lissauer, J. J., Fabrycky, D. C., Ford, E. B., Borucki, W. J., Fressin, F., Marcy, G. W., Orosz, J. A., Rowe, J. F., Torres, G., Welsh, W. F., Batalha, N. M., Bryson, S. T., Buchhave, L. A., Caldwell, D. A., Carter, J. A., Charbonneau, D., Christiansen,

J. L., Cochran, W. D., Desert, J., Dunham, E. W., Fanelli, M. N., Fortney, J. J., Gautier, T. N., III, Geary, J. C., Gilliland, R. L., Haas, M. R., Hall, J. R., Holman, M. J., Koch, D. G., Latham, D. W., Lopez, E., McCauliff, S., Miller, N., Morehead, R. C., Quintana, E. V., Ragozzine, D., Sasselov, D., Short, D. R., and Steffen, J. H. (2011a). "A closely packed system of low-mass, low-density planets transiting Kepler-11." *Nature*, 470: 53–58. 489, 490

Lissauer, J. J., Ragozzine, D., Fabrycky, D. C., Steffen, J. H., Ford, E. B., Jenkins, J. M., Shporer, A., Holman, M. J., Rowe, J. F., Quintana, E. V., Batalha, N. M., Borucki, W. J., Bryson, S. T., Caldwell, D. A., Carter, J. A., Ciardi, D., Dunham, E. W., Fortney, J. J., Gautier, T. N., III, Howell, S., Koch, D. G., Latham, D. W., Marcy, G. W., Morehead, R. C., and Sasselov, D. (2011b). "Architecture and Dynamics of Kepler's Candidate Multiple Transiting Planet Systems." `arxiv:1102.0543 [astro-ph.EP]`. 477, 485, 490

Maciejewski, G., Dimitrov, D., Neuhaeuser, R., Niedzielski, A., Raetz, S., Ginski, C., Adam, C., Marka, C., Moualla, M., and Mugrauer, M. (2010). "Transit timing variation in exoplanet WASP-3b." *Monthly Notices of the Royal Astronomical Society*, 407: 2625–2631. 489

Miller-Ricci, E., Rowe, J. F., Sasselov, D., Matthews, J. M., Guenther, D. B., Kuschnig, R., Moffat, A. F. J., Rucinski, S. M., Walker, G. A. H., and Weiss, W. W. (2008). "MOST Space-based Photometry of the Transiting Exoplanet System HD 209458: Transit Timing to Search for Additional Planets." *Astrophysical Journal*, 682: 586–592. 489

Moorhead, A. and Ford, E. (2010). "Resolving the degeneracy between eccentric planets and 2:1 mean motion resonances." In K. Gożdziewski, A. Niedzielski, & J. Schneider (ed.), *EAS Publications Series*, volume 42 of *EAS Publications Series*, 161–164. 485

Moorhead, A. V., Ford, E. B., Morehead, R. C., Rowe, J., Borucki, W. J., Batalha, N. M., Bryson, S. T., Caldwell, D. A., Fabrycky, D. C., Gautier, T. N., III, Koch, D. G., Holman, M. J., Jenkins, J. M., Li, J., Lissauer, J. J., Lucas, P., Marcy, G. W., Quinn, S. N., Quintana, E., Ragozzine, D., Shporer, A., Still, M., and Torres, G. (2011). "The Distribution of Transit Durations for Kepler Planet Candidates and Implications for their Orbital Eccentricities." *Astrophysical Journal Supplement*. 489

Nesvorný, D. (2009). "Transit Timing Variations for Eccentric and Inclined Exoplanets." *Astrophysical Journal*, 701: 1116–1122. 476, 489

Nesvorný, D. and Beaugé, C. (2010). "Fast Inversion Method for Determination of Planetary Parameters from Transit Timing Variations." *Astrophysical Journal Letters*, 709: L44–L48. 476, 490

Nesvorný, D. and Morbidelli, A. (2008). "Mass and Orbit Determination from Transit Timing Variations of Exoplanets." *Astrophysical Journal*, 688: 636–646. 476, 489

Payne, M. J. and Ford, E. B. (2011). "An Analysis of Jitter and Transit Timing Variations in the HAT-P-13 System." *Astrophysical Journal*, 729: 98–112. 484

Payne, M. J., Ford, E. B., and Veras, D. (2010). "Transit Timing Variations for Inclined and Retrograde Exoplanetary Systems." *Astrophysical Journal Letters*, 712: L86–L92. 490

Ragozzine, D. and Holman, M. J. (2010). "The Value of Systems with Multiple Transiting Planets." `arxiv:1006.3727 [astro-ph.EP]`. 490

Slawson, R. W., Prsa, A., Welsh, W. F., Orosz, J. A., Rucker, M., Batalha, N. M., Doyle, L. R., Engle, S. G., Conroy, K., Coughlin, J., Ames Gregg, T., Fetherolf, T., Short, D. R., Windmiller, G., Fabrycky, D. C., Howell, S. B., Jenkins, J. M., Uddin, K., Mullally, F., Seader, S. E., Thompson, S. E., Sanderfer, D. T., Borucki, W., and Koch, D. (2011). "Kepler Eclipsing Binary Stars. II. 2165 Eclipsing Binaries in the Second Data Release." `arxiv:1103.1659 [astro-ph.SR]`. 477

Steffen, J. H. and Agol, E. (2005). "An analysis of the transit times of TrES-1b." *Monthly Notices of the Royal Astronomical Society*, 364: L96–L100. 477

Steffen, J. H., Batalha, N. M., Borucki, W. J., Buchhave, L. A., Caldwell, D. A., Cochran, W. D., Endl, M., Fabrycky, D. C., Fressin, F., Ford, E. B., Fortney, J. J., Haas, M. J., Holman, M. J., Isaacson, H., Jenkins, J. M., Koch, D., Latham, D. W., Lissauer, J. J., Moorhead, A. V., Morehead, R. C., Marcy, G., MacQueen, P. J., Quinn, S. N., Ragozzine, D., Rowe, J. F., Sasselov, D. D., Seager, S., Torres, G., and Welsh, W. F. (2010). "Five Kepler target stars that show multiple transiting exoplanet candidates." *Astrophysical Journal*, 725: 1226–1241. 489

Steffen, J. H., Quinn, S. N., Borucki, W. J., Brugamyer, E., Bryson, S. T., Buchhave, L. A., Cochran, W. D., Endl, M., Fabrycky, D. C., Ford, E. B., Holman, M. J., Jenkins, J., Koch, D., Latham, D. W., MacQueen, P., Mullally, F., Prsa, A., Ragozzine, D., Rowe, J. F., Sanderfer, D. T., Seader, S. E., Short, D., Shporer, A., Thompson, S. E., Torres, G., Twicken, J. D., Welsh, W. F., and Windmiller, G. (2011). "The architecture of the hierarchical triple star KOI 928 from eclipse timing variations seen in Kepler photometry." `arxi:1106.4530 [astro-ph.SR]`. 477

Veras, D., Ford, E. B., and Payne, M. J. (2011). "Quantifying the Challenges of Detecting Unseen Planetary Companions with Transit Timing Variations." *Astrophysical Journal*, 727: 74–97. 476, 479, 489, 490

Wright, J. T. (2005). "Radial Velocity Jitter in Stars from the California and Carnegie Planet Search at Keck Observatory." *Publications of the Astronomical Society of the Pacific*, 117: 657–664. 484, 492

Wright, J. T., Veras, D., Ford, E. B., Johnson, J. A., Marcy, G. W., Howard, A. W., Isaacson, H., Fischer, D. A., Spronck, J., Anderson, J., and Valenti, J. (2011). "The California Planet Survey. III. A Possible 2:1 Resonance in the Exoplanetary Triple System HD 37124." *Astrophysical Journal*, 730: 93–111. 485

Youdin, A. N. (2011). "The Exoplanet Census: A General Method, Applied to Kepler." `arxiv:1105.1782 [astro-ph.EP]`. 483

Acknowledgments

The authors wish to thank Malay Ghosh, Phil Gregory, Tom Loredo and Matthew Payne for valuable discussions and feedback. This material is based upon work supported by the National Science Foundation under Grant No. 0707203. Additional support for this work was provided by the National Aeronautics and Space Administration under grant NNX09AB35Gb issued through the Origins of Solar Systems program and grant NNX08AR04G issued through the Kepler Participating Scientist Program.