

Extra-solar Planets via Bayesian model fitting

Phil Gregory

University of British Columbia

Nov. 8, 2011

A rich harvest of over 600 extra-solar planets is motivating a significant effort to improve statistical tools to better understand parameter uncertainties and compute model probabilities.

Recent work has highlighted a Bayesian Markov chain Monte Carlo approach - applicable to many nonlinear model fitting problems.

Outline

- 1. Bayesian primer** 1
- 2. Status of extra-solar planets (exoplanets)** 2
 - Challenge of nonlinear models** 3
- 3. Introduction to Markov chain Monte Carlo (MCMC)** 4
 - Fusion MCMC** 5
 - Multi-planet Kepler periodogram** 6
- 4. Evidence for 3 planets in 47 Ursae Majoris** 7
- 5. Test of habitable zone planet Gliese 581g** 8
- 6. Conclusions** 9

What is Bayesian Probability Theory? (BPT)

BPT = a theory of extended logic

Deductive logic is based on Axiomatic knowledge.

In science we never know any theory of nature is true because our reasoning is based on incomplete information.

Our conclusions are at best probabilities.

Any extension of logic to deal with situations of incomplete information (realm of inductive logic) requires a theory of probability.

A new perception of probability has arisen in recognition that the mathematical rules of probability are not merely rules for manipulating random variables.

They are now recognized as valid principles of logic for conducting inference about any hypothesis of interest.

This view of, ``Probability Theory as Logic'', was championed in the late 20th century by E. T. Jaynes.

**``Probability Theory: The Logic of Science''
Cambridge University Press 2003**

It is also commonly referred to as Bayesian Probability Theory in recognition of the work of the 18th century English clergyman and Mathematician Thomas Bayes.

Logic is concerned with the truth of propositions.

A proposition asserts that something is true.

Examples of propositions:

$A \equiv$ “The newly discovered radio astronomy object is a galaxy.”

$B \equiv$ “The measured redshift of the object is 0.150 ± 0.005 .”

$A \equiv$ “Theory X is correct.”

$\bar{A} \equiv$ “Theory X is not correct.”

$A \equiv$ “The frequency of the signal is between f and $f + df$.”

We will need to consider compound propositions like A, B which asserts that propositions A and B are true

$A, B/C$ asserts that propositions A and B are true given that proposition C is true

Rules for manipulating probabilities

Sum rule : $p(A | C) + p(\bar{A} | C) = 1$

Product rule : $p(A, B | C) = p(A | C) p(B | A, C)$
 $= p(B | C) p(A | B, C)$

Re-arrange the two RH sides of product rule gives

Bayes theorem :

$$p(A | B, C) = \frac{p(A | C) p(B | A, C)}{p(B | C)}$$

How to proceed in a Bayesian data analysis?

Write down Bayes' theorem, identify the terms and solve.

$$p(H_i | D, I) = \frac{p(H_i | I) \times p(D | H_i, I)}{p(D | I)}$$

Posterior probability
that H_i is true, given
the new data D and
prior information I

Normalizing constant

Every item to the right of the
vertical bar $|$ is assumed to be true

The likelihood $p(D | H_i, I)$, also written as $\mathcal{L}(H_i)$, stands for the probability that we would have gotten the data D that we did, if H_i is true.

As a theory of extended logic BPT can be used to find **optimal answers** to well posed scientific questions for a given state of knowledge, in contrast to a numerical recipe approach.

Two basic problems

1. Model selection (discrete hypothesis space)

“Which one of 2 or more models (hypotheses) is most probable given our current state of knowledge?”

e.g.

- Hypothesis or model M_0 asserts that the star has no planets.
- Hypothesis M_1 asserts that the star has 1 planet.
- Hypothesis M_i asserts that the star has i planets.

2. Parameter estimation (continuous hypothesis)

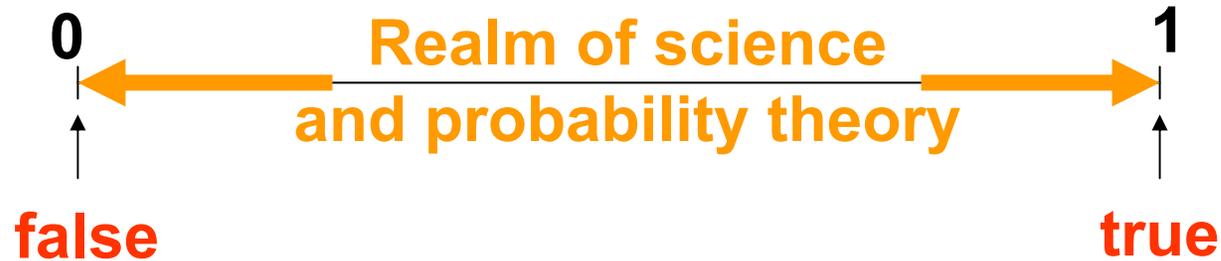
“Assuming the truth of M_1 , solve for the probability density distribution for each of the model parameters based on our current state of knowledge.”

e.g.

- Hypothesis H asserts that the orbital period is between P and $P+dP$.

Significance of this development

Probabilities are commonly quantified by a real number between 0 and 1.



The end-points, corresponding to absolutely false and absolutely true, are simply the extreme limits of this infinity of real numbers.

Bayesian probability theory spans the whole range.

Deductive logic is just a special case of Bayesian probability theory in the idealized limit of complete information.

Calculation of a simple Likelihood $p(D | M, \vec{X}, I)$

Let d_i represent the i^{th} measured data value . We model d_i by,

$$d_i = f_i(\vec{X}) + e_i$$

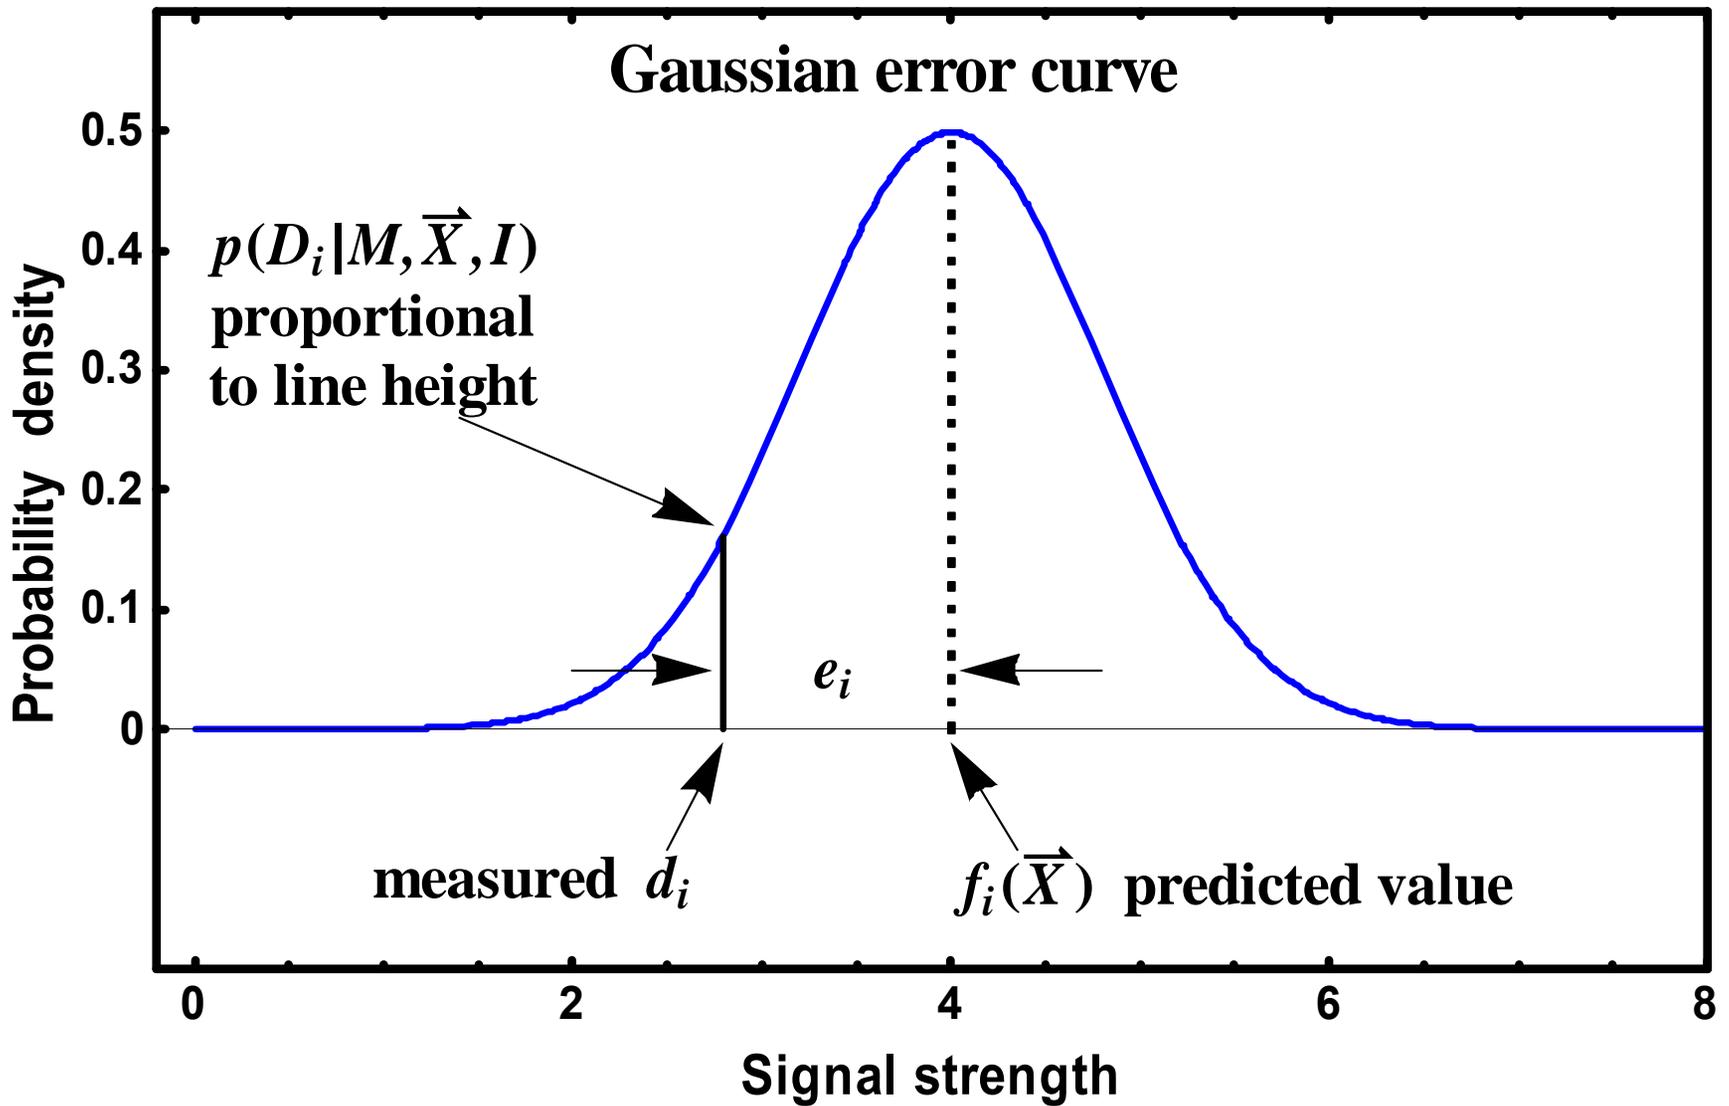
Model prediction for i^{th} data value
for current choice of parameters \vec{X}

where e_i represents the error component in the measurement.

Since M, \vec{X} is assumed to be true, if it were not for the error e_i , d_i would equal the model prediction f_i .

Now suppose prior information I indicates that e_i has a Gaussian probability distribution. Then

$$\begin{aligned} p(D_i | M, \vec{X}, I) &= \frac{1}{\sigma_i \sqrt{2\pi}} \text{Exp}\left[-\frac{e_i^2}{2\sigma_i^2}\right] \\ &= \frac{1}{\sigma_i \sqrt{2\pi}} \text{Exp}\left[-\frac{(d_i - f_i(\vec{X}))^2}{2\sigma_i^2}\right] \end{aligned}$$



Probability of getting a data value d_i a distance e_i away from the predicted value f_i is proportional to the height of the Gaussian error curve at that location.

Calculation of a simple Likelihood $p(D | M, \vec{X}, I)$

For independent data the likelihood for the entire data set $D=(D_1, D_2, \dots, D_N)$ is the product of N Gaussians.

$$p(D | M, \vec{X}, I) = (2\pi)^{-N/2} \left\{ \prod_{i=1}^N \sigma_i^{-1} \right\} \text{Exp} \left[-0.5 \sum_{i=1}^N \frac{(d_i - f_i(\vec{X}))^2}{\sigma_i^2} \right]$$

The familiar χ^2 statistic used in least-squares

Maximizing the likelihood corresponds to minimizing χ^2

Recall: Bayesian posterior \propto prior \times likelihood

Thus, only for a uniform prior will a least-squares analysis yield the same solution as the Bayesian posterior.

Simple example of when not to use a uniform prior

In the exoplanet problem the prior range for the unknown orbital period P is very large from $\sim 1/2$ day to 1000 yr (upper limit set by perturbations from neighboring stars).

Suppose we assume a uniform prior probability density for the P parameter. This would imply that we believed that it was $\sim 10^4$ times more probable that the true period was in the upper decade (10^4 to 10^5 d) of the prior range than in the lowest decade from 1 to 10 d.

$$\frac{\int_{10^4}^{10^5} p(P | M, I) dP}{\int_1^{10} p(P | M, I) dP} = 10^4$$

Usually, expressing great uncertainty in some quantity corresponds more closely to a statement of scale invariance or equal probability per decade. The Jeffreys prior has this scale invariant property.

$$p(\ln P | M, I) d \ln P = \frac{d \ln P}{\ln(P_{max} / P_{min})}$$

Integration not minimization

A full Bayesian analysis requires integrating over the model parameter space. Integration is more difficult than minimization.

However, the Bayesian solution provides the most accurate information about the parameter errors and correlations without the need for any additional calculations, i.e., Monte Carlo simulations.

**Shortly discuss an efficient method for
Integrating over a large parameter space
called Markov chain Monte Carlo (MCMC).**

PHIL GREGORY

Bayesian Logical Data Analysis for the Physical Sciences

A Comparative Approach with
Mathematica Support



CAMBRIDGE

Resources and solutions

This title has free
Mathematica based support
software available

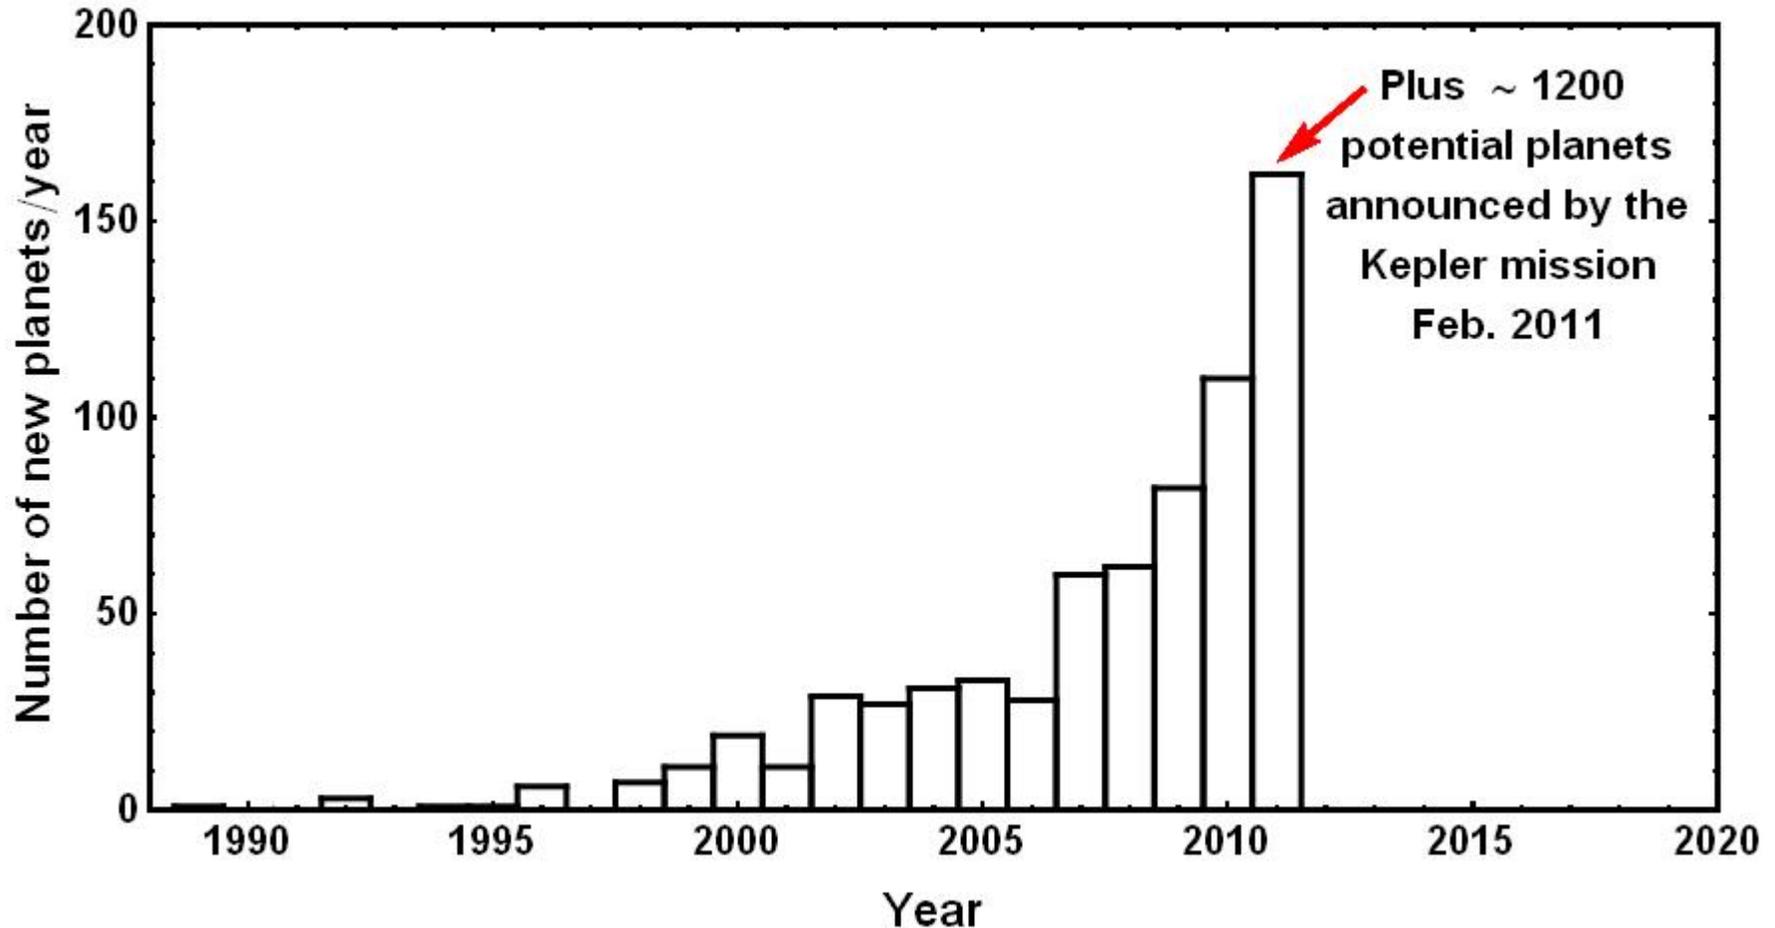
Chapters

1. Role of probability theory in science
2. Probability theory as extended logic
3. The how-to of Bayesian inference
4. Assigning probabilities
5. **Frequentist statistical inference**
6. **What is a statistic?**
7. **Frequentist hypothesis testing**
8. Maximum entropy probabilities
9. Bayesian inference (Gaussian errors)
10. Linear model fitting (Gaussian errors)
11. Nonlinear model fitting
12. Markov chain Monte Carlo
13. Bayesian spectral analysis
14. Bayesian inference (Poisson sampling)

Introduces statistical inference in the larger context of scientific methods, and includes 55 worked examples and many problem sets.

End section 1

Pace of Extra-solar Planet Discoveries as of Sept. 2011



Of the ~ 1200 Kepler mission candidates, **30 have a radius \leq earth.**

These need to be confirmed by RV measurements or transit timing variation analysis.

Extra-solar planets detected as of 28 Oct. 2011

695 exoplanets planets found in 571 planetary systems

The majority within a distance of 100 ly

81 multiple planet systems (one 7 planet system)

Detection method	No. of planets
Radial velocity	644
Transits	185
Microlensing	13
Imaging	26
Timing (includes 4 pulsar planets)	12

The Extrasolar Planets Encyclopaedia
Jean Schneider
(CNRS-LUTH, Paris Observatory)

My talk concerned with radial velocity data sets.

The radial velocity equation, a nonlinear model

$$f_i = \text{model prediction} = V + K (\cos(\theta[t_i + \chi P] + \omega) + e \cos \omega)$$

V = systematic velocity

$$K = \text{velocity amplitude} = \frac{2 \pi a \sin i}{P \sqrt{1 - e^2}} ;$$

(**a** = semi - major axis, **i** = inclination)

e = orbital eccentricity

ω = longitude of periastron

χ = fraction of orbit prior to data reference time that periastron occurred at

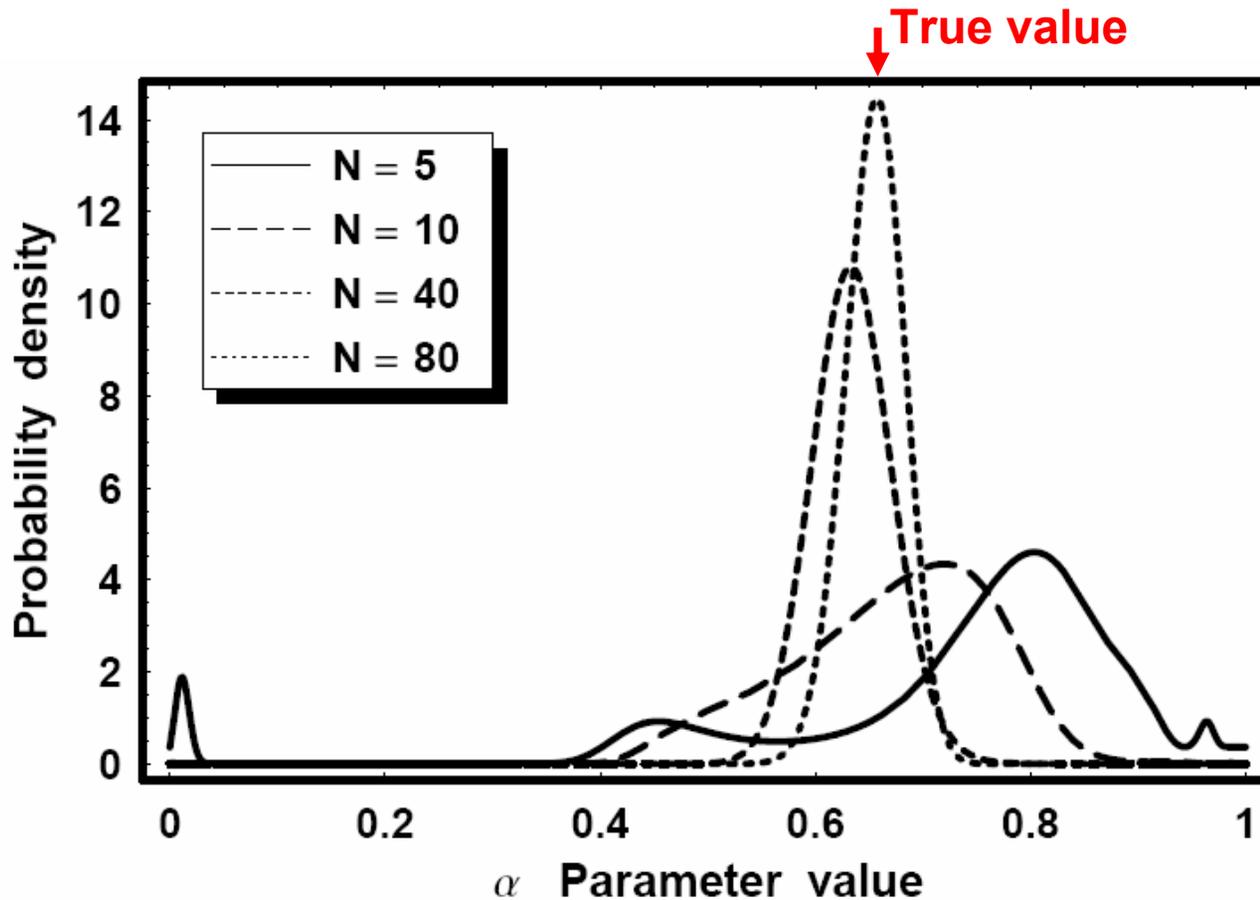
$\theta(t_i + \chi P)$ = true anomaly

= angle of star in orbit at time t_i relative to periastron

θ_i and **t_i** are related by the conservation of angular momentum equation

$$\partial_t \theta[t] - \frac{2\pi}{P(1 - e^2)^{3/2}} (1 + e \cos[\theta])^2 = 0$$

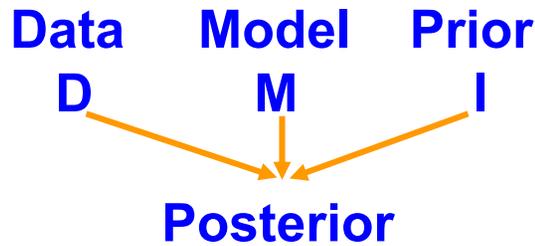
The challenge of nonlinear models, multiple peaks



The Bayesian posterior density for a nonlinear model with a single parameter, α , for 4 simulated data sets of different size ranging from $N = 5$ to $N = 80$. The $N = 5$ case has the broadest distribution and exhibits 4 maxima.

Asymptotic theory says that the maximum likelihood estimator becomes more unbiased, more normally distributed and of smaller variance as the sample size becomes larger.

**Numerical tools
for Bayesian
model fitting**



Linear models (uniform priors)

**Posterior has a single peak
(multi-dimensional Gaussian)**

Parameters given
by the normal equations
of linear least-squares

No integration required
solution very fast
using linear algebra

(chapter 10)

Nonlinear models

+ linear models (non-uniform priors)

Posterior may have multiple peaks

**Brute force
integration**

**For some
parameters
analytic
integration
sometimes
possible**

**Asymptotic
approx.'s**

peak finding
algorithms

- (1) Levenberg-
Marquardt
- (2) Simulated
annealing
- (3) Genetic
algorithm

Laplace
approx.'s

(chapter 11)

**Moderate
dimensions**

quadrature

randomized
quadrature

adaptive
quadrature

**High
dimensions**

nested sampling

MCMC

(chapter 12)

Why is integration important for Bayesian analysis?

Parameter estimation: to find the marginal posterior probability density function (PDF) for the orbital period P , we need to integrate the joint posterior over all the other parameters.

$$\underbrace{p(P | D, M_1, s, I)}_{\text{Marginal PDF for } P} = \int dK dV d\chi d\mathbf{e} d\omega ds \underbrace{p(P, K, V, \mathbf{e}, \phi, \omega, \mathbf{s} | D, M_1, s, I)}_{\text{Joint posterior probability density function (PDF) for the parameters}}$$

Marginal PDF
for P

Joint posterior probability
density function (PDF) for
the parameters

Markov chain Monte Carlo (MCMC) algorithms provide a powerful means for efficiently computing integrals in many dimensions to within a constant factor. This factor is not required for parameter estimation.

After an initial burn-in period, the MCMC produces an equilibrium distribution of samples in parameter space such that the density of samples is **proportional** to the **joint posterior PDF**.

It is very efficient because, unlike straight Monte Carlo integration, it doesn't waste time exploring regions where the joint posterior is very small.

Starting point: Metropolis-Hastings MCMC algorithm

$P(X|D,M,I)$ = target posterior probability distribution
(X represents the set of model parameters)

1. Choose X_0 an initial location in the parameter space . Set $t = 0$.

2. Repeat {

– Obtain a new sample Y from a proposal distribution $q(Y|X_t)$ that is easy to evaluate . $q(Y|X_t)$ can have almost any form.

I use a Gaussian proposal distribution. i.e., Normal distribution $N(X_t, \sigma)$

– Sample a Uniform (0, 1) random variable U .

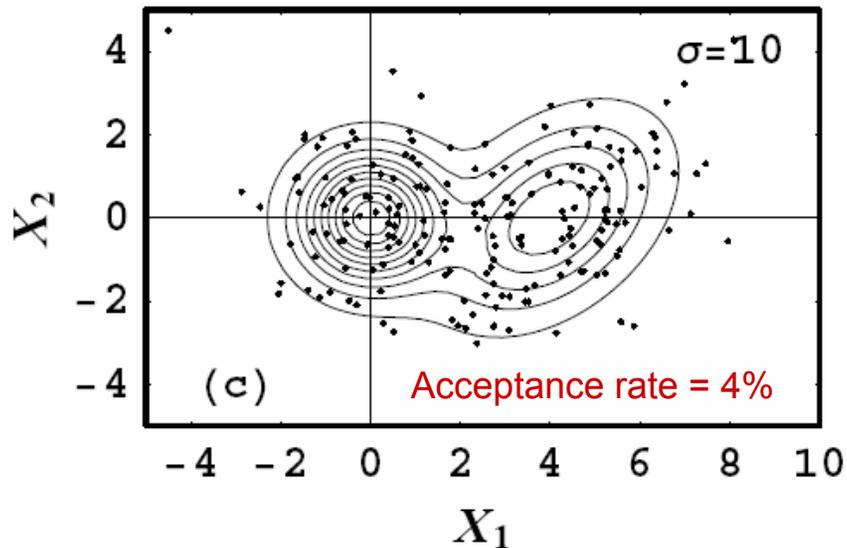
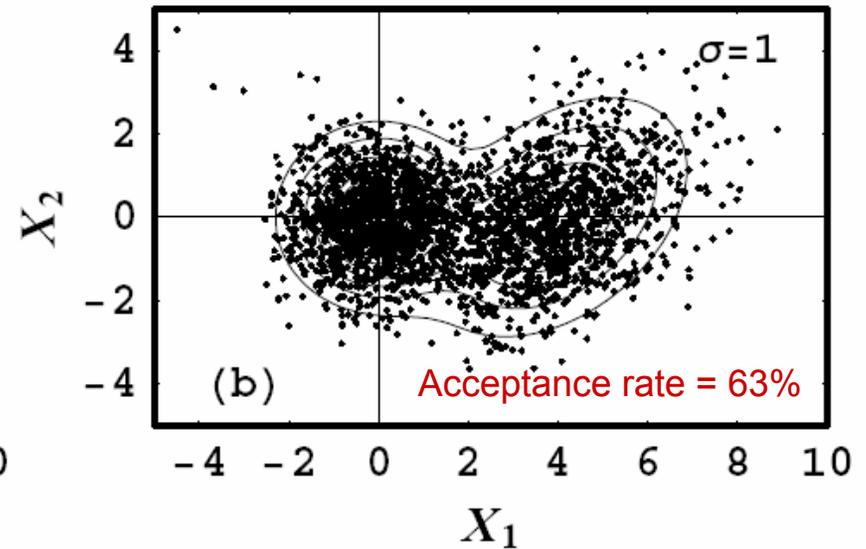
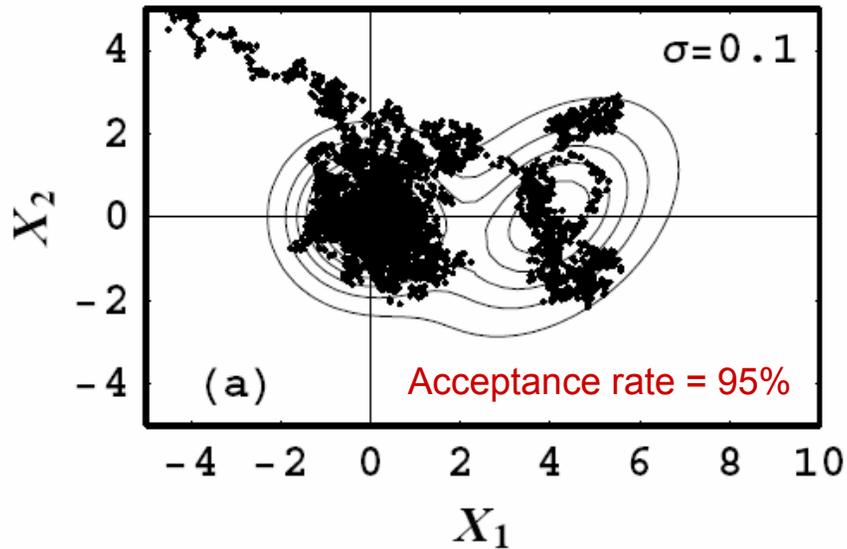
– If $U \leq \frac{p(Y|D, I)}{p(X_t|D, I)} \times \frac{q(X_t|Y)}{q(Y|X_t)}$ then set $X_{t+1} = Y$

otherwise set $X_{t+1} = X_t$

This factor =1
for a symmetric proposal
distribution like a Gaussian

– Increment t }

Toy MCMC simulations: the efficiency depends on tuning proposal distribution σ 's. Can be a very difficult challenge for many parameters.



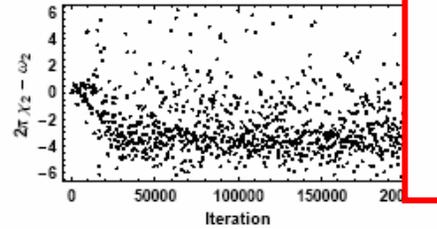
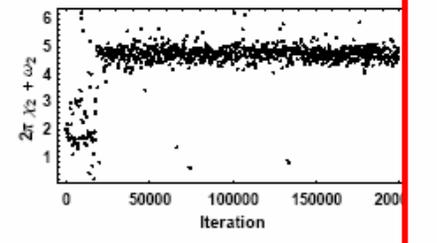
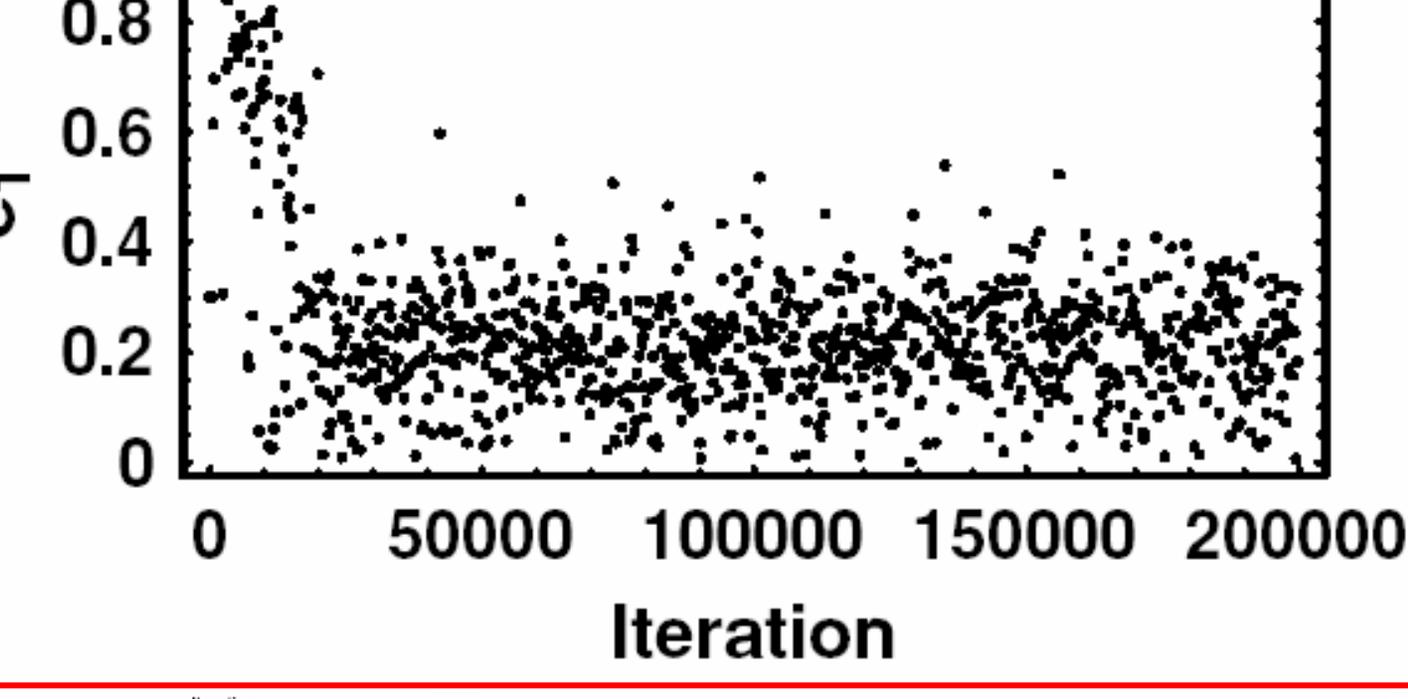
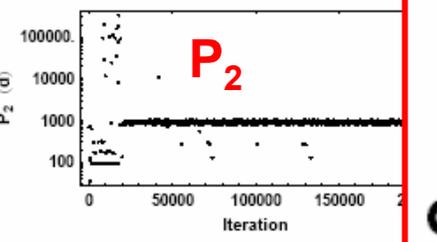
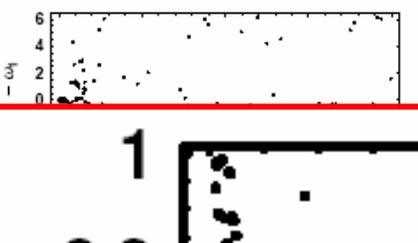
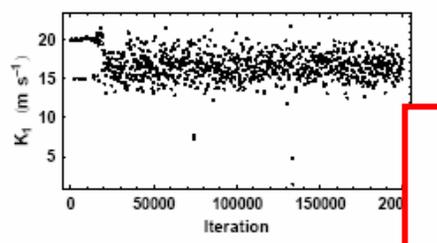
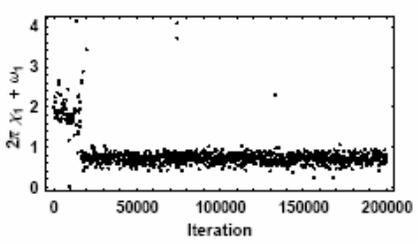
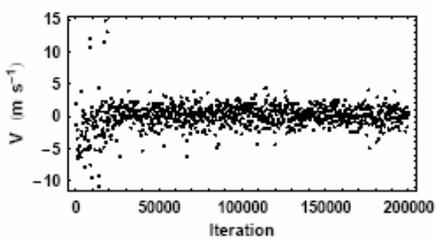
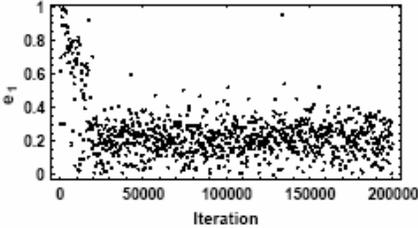
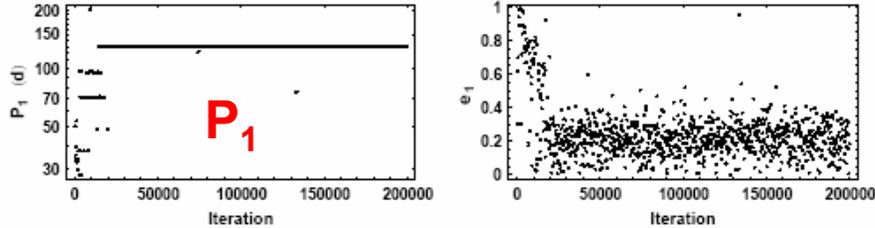
In this example the posterior probability distribution consists of two 2 dimensional Gaussians indicated by the contours

MCMC parameter samples for 2 planet model.

MNRAS, 374, 1321, 2007

P. C. Gregory

Title: **A Bayesian Kepler
Periodogram Detects a
Second Planet in HD 208487**



Parallel tempering MCMC

The simple Metropolis-Hastings MCMC algorithm can run into difficulties if the probability distribution is multi-modal with widely separated peaks. It can fail to fully explore all peaks which contain significant probability, especially if some of the peaks are very narrow.

One solution is to run multiple Metropolis-Hastings simulations in parallel, employing probability distributions of the kind

$$\pi(\mathbf{X} | \mathbf{D}, \mathbf{M}, \beta, \mathbf{I}) = p(\mathbf{X} | \mathbf{M}, \mathbf{I}) p(\mathbf{D} | \mathbf{X}, \mathbf{M}, \mathbf{I})^\beta \quad (0 < \beta \leq 1)$$

Typical set of β values = 0.09, 0.15, 0.22, 0.35, 0.48, 0.61, 0.78, 1.0
 $\beta = 1$ corresponds to our desired target distribution. The others correspond to progressively flatter probability distributions.

At intervals, a pair of adjacent simulations are chosen at random and a proposal made to swap their parameter states. The swap allows for an exchange of information across the ladder of simulations.

In the low β simulations, radically different configurations can arise, whereas at higher β , a configuration is given the chance to refine itself.

Final results are based on samples from the $\beta = 1$ simulation. Samples from the other simulations provide one way to evaluate the Bayes Factor in model selection problems.

Controlled Statistical Fusion

Next describe progress in the development of a new fusion Markov chain Monte Carlo algorithm which incorporates:

- a) **parallel tempering (also known as Exchange Monte Carlo),**
- b) **simulated annealing and,**
- c) **genetic crossover operations.**

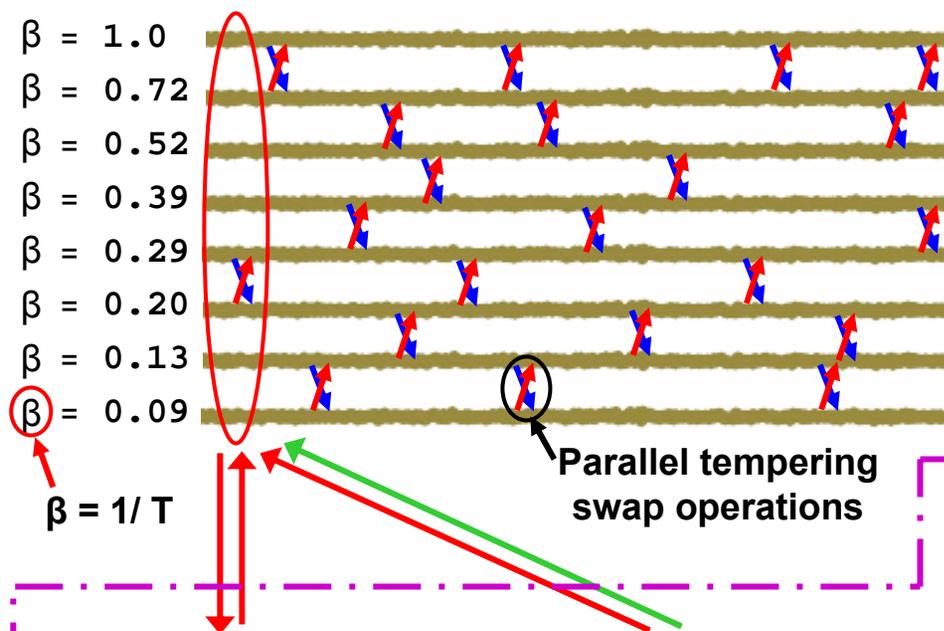
Each of these features facilitate the detection of a global minimum in χ^2 (maximum in posterior probability) in a multi-modal environment. By combining all three, the algorithm greatly increases the probability of realizing this goal.

This statistical fusion approach has been achieved through the development of a unique multi-stage adaptive control system, hence the term **controlled statistical fusion**.

Among other things the control system automates the tuning of MCMC proposal distributions for efficient exploration of the model parameter space even when the parameters are highly correlated.

Adaptive fusion MCMC

8 parallel tempering Metropolis chains



Output at each iteration

parameters, logprior + $\beta \times$ loglike, logprior + loglike
parameters, logprior + $\beta \times$ loglike, logprior + loglike
parameters, logprior + $\beta \times$ loglike, logprior + loglike
parameters, logprior + $\beta \times$ loglike, logprior + loglike
parameters, logprior + $\beta \times$ loglike, logprior + loglike
parameters, logprior + $\beta \times$ loglike, logprior + loglike
parameters, logprior + $\beta \times$ loglike, logprior + loglike
parameters, logprior + $\beta \times$ loglike, logprior + loglike

Anneal Gaussian proposal σ 's

Refine & update Gaussian proposal σ 's

2 stage proposal σ control system
error signal =
(actual joint acceptance rate - 0.25)
Effectively defines burn-in interval

Peak parameter set:
If (logprior + loglike) > previous best by a threshold then update and reset burn-in

Monitor for parameters with peak probability

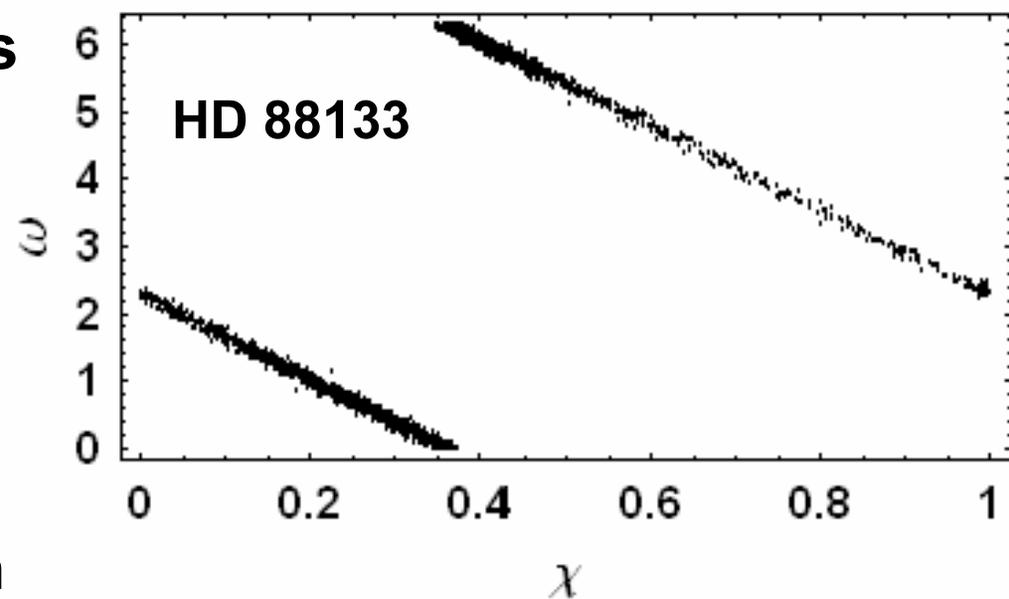
Genetic algorithm

Every 10th iteration perform gene crossover operation to breed larger (logprior + loglike) parameter set.

MCMC adaptive control system

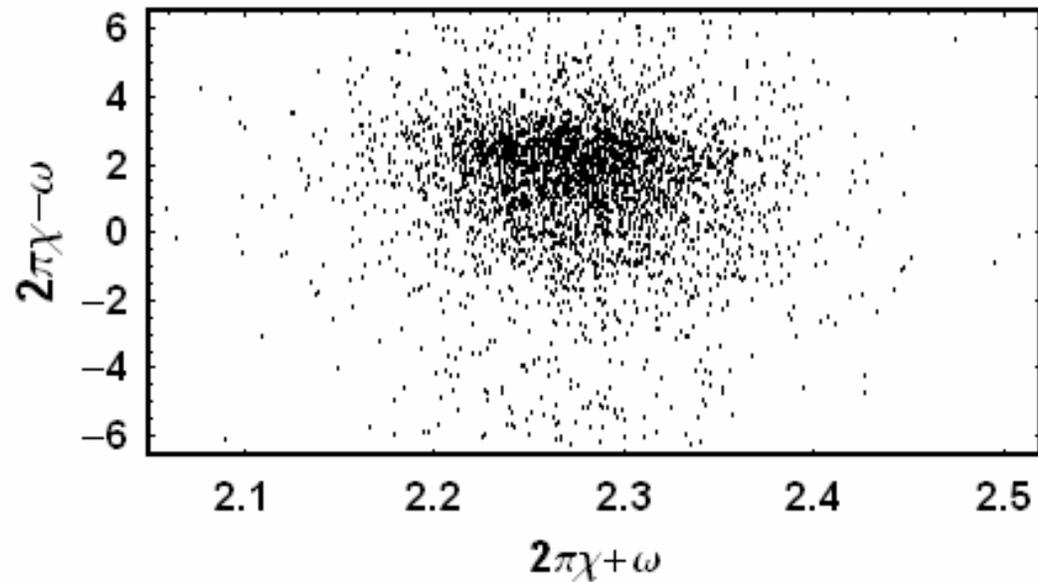
Highly correlated parameters

Top figure shows an exoplanet Example. For low eccentricity orbits the parameters ω and χ are not separately well determined. This shows up as a strong correlation between ω and χ .



One option re-parameterization

The combination $2\pi\chi+\omega$ is well determined for all eccentricities. Although $2\pi\chi-\omega$ is not well determined for low eccentricities, it is at least orthogonal to $2\pi\chi+\omega$ as shown.



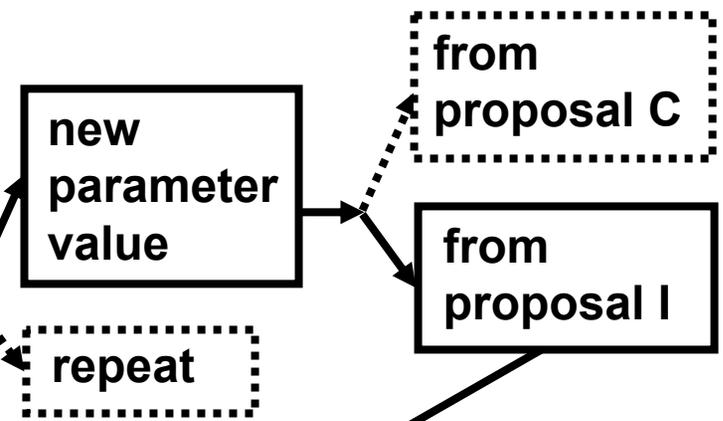
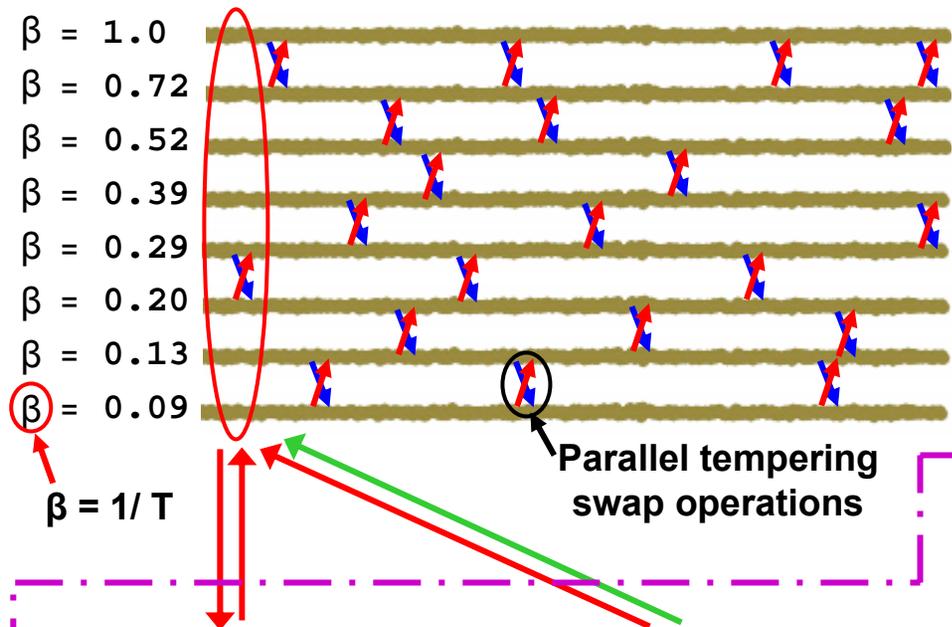
Another option

Algorithm learns about the parameter correlations during the burn-in and generates proposals with these statistical correlations.

Adaptive Fusion MCMC

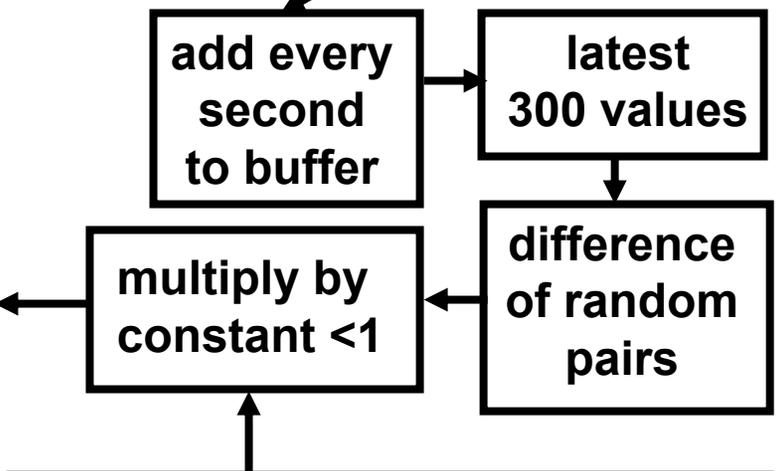
8 parallel tempering Metropolis chains

Automatic proposal scheme that learns about parameter correlations during burn-in (for each chain)



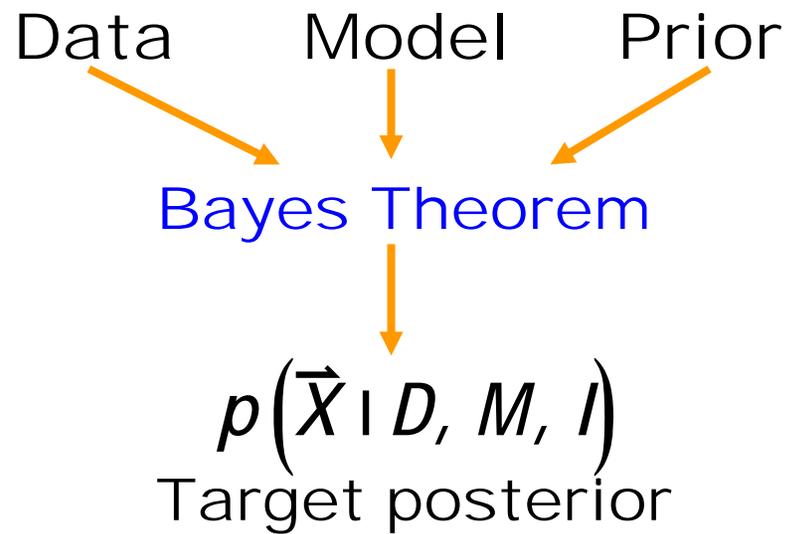
'I' proposals
Independent
Gaussian proposal
scheme employed
50% of the time

'C' proposals
Proposal distribution
with built in parameter
correlations used
50% of the time



During burn-in control system
adjusts constant so acceptance
rate from Proposal C = 25 %

MCMC adaptive control system



If you input a Kepler model the fusion MCMC becomes

A Kepler periodogram

**Optimum for finding Kepler orbits and
evaluating their probabilities.**

Capable of simultaneously fitting multiple planet models.

A multi-planet Kepler periodogram

An 8 planet model has 41 parameters.

Model space considered

Symbol	Model	# of parameters
M_{0s}	Constant velocity V + extra noise term s	$1+s$
M_{1s}	V + elliptical orbit + extra noise term s	$6+s$
M_{2s}	V + 2 elliptical orbits + extra noise term s	$11+s$
M_{3s}	V + 3 elliptical orbits + extra noise term s	$16+s$
M_{js}	V + j elliptical orbits + extra noise term s	$5j+1+s$

An extra noise term allows for an intrinsic stellar variability (“jitter”) that we model as an additional source of uncorrelated Gaussian noise with variance s^2 and add to the measurement uncertainties in quadrature. s becomes an additional parameter to marginalize over.

Note: some forms of stellar jitter (e.g., star spots) can produce Keplerian-like radial velocity variations.

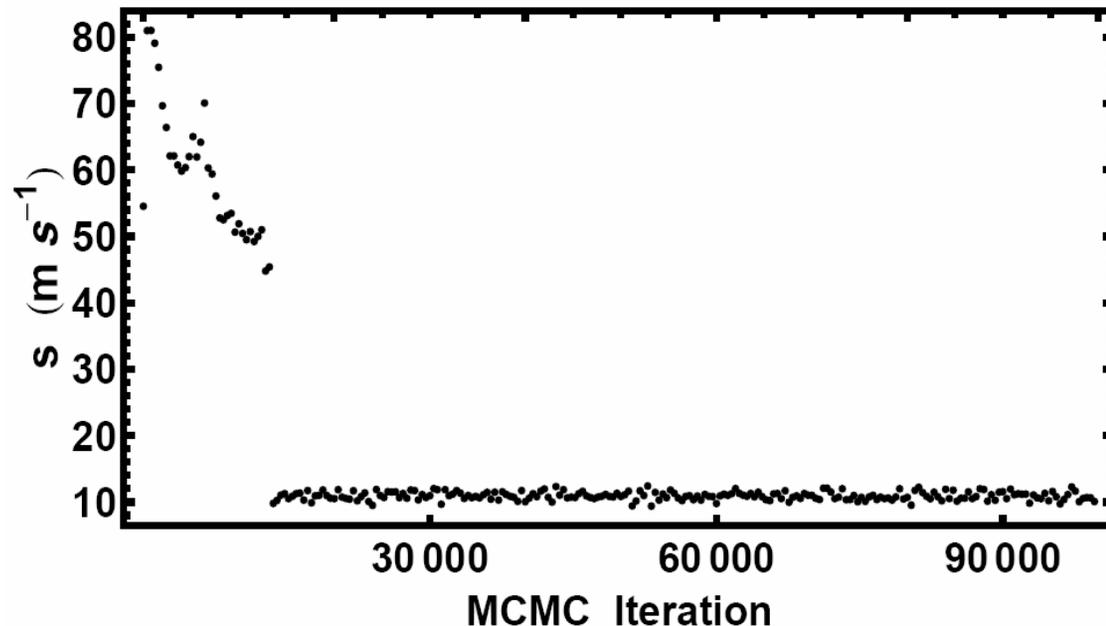
Annealing due to extra noise term, s

The inclusion of an extra noise term of unknown magnitude also gives rise to an annealing operation when the Markov chain is started far from the best-fit values.

If only known observational errors are included, the posterior probability distribution is often very “rough” with many local maxima throughout parameter space.

When s is included, Bayesian Markov chain automatically inflates s to include anything in the data that cannot be accounted for by the model with the current set of parameters and the known measurement errors.

This results in a smoothing out of the posterior surface and allows the Markov chain to explore the parameter space more quickly. The chain begins to decrease the value of s as it settles in near the best-fit parameters. **This behavior is similar to simulated annealing, but does not require choosing a cooling scheme.**



Parameter priors

Parameter	Prior	Lower bound	Upper bound
Orbital frequency	$p(\ln f_1, \ln f_2, \dots, \ln f_n M_n, I) = \frac{n!}{[\ln(f_H/f_L)]^n}$ (n =number of planets)	1/1.5 d	1/1000 yr
Velocity K_i (m s^{-1})	Modified Jeffreys ^a $\frac{(K+K_0)^{-1}}{\ln \left[1 + \frac{K_{\max}}{K_0} \left(\frac{P_{\min}}{P_i} \right)^{1/3} \frac{1}{\sqrt{1-e_i^2}} \right]}$	0 ($K_0 = 1$)	$K_{\max} \left(\frac{P_{\min}}{P_i} \right)^{1/3} \frac{1}{\sqrt{1-e_i^2}}$ $K_{\max} = 2129$ K_{\max} corresponds to a max. planet-star mass ratio = 0.01
V (m s^{-1})	Uniform	$-K_{\max}$	K_{\max}
e_i Eccentricity	a) Uniform b) Ecc. noise bias correction filter	0 0	1 0.99
ω_i Longitude of periastron	Uniform	0	2π
s Extra noise (m s^{-1})	$\frac{(s+s_0)^{-1}}{\ln \left(1 + \frac{s_{\max}}{s_0} \right)}$	0 ($s_0 = 1$)	K_{\max}

^a Since the prior lower limits for K and s include zero, we used a modified Jeffreys prior of the form

$$p(X|M, I) = \frac{1}{X + X_0} \frac{1}{\ln \left(1 + \frac{X_{\max}}{X_0} \right)} \quad (4)$$

For $X \ll X_0$, $p(X|M, I)$ behaves like a uniform prior and for $X \gg X_0$ it behaves like a Jeffreys prior. The $\ln \left(1 + \frac{X_{\max}}{X_0} \right)$ term in the denominator ensures that the prior is normalized in the interval 0 to X_{\max} .

47 Ursae Majoris is a solar twin ~46 light years away in Ursa Major

History

- 1) 1996, report of a $P = 1090$ day companion

Butler, R. P. & Marcy, G. W. 1996, ApJ, 464, L153

- 2) 2002, report of a 2nd companion,
 $P = 2594 \pm 90$ days

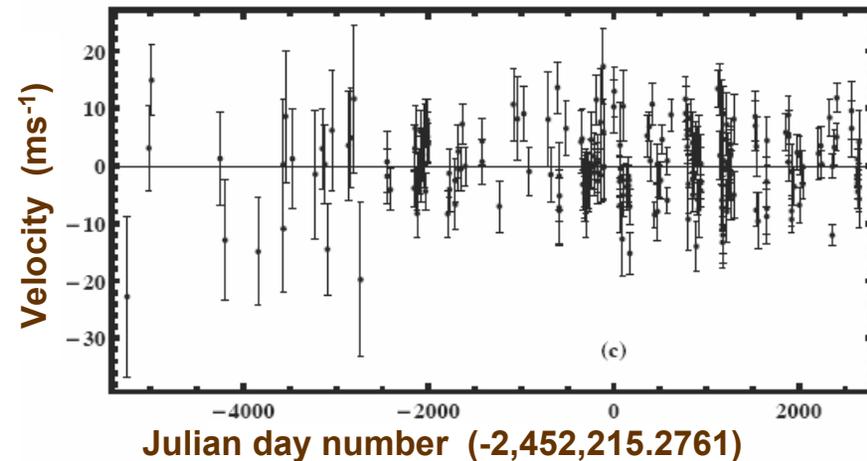
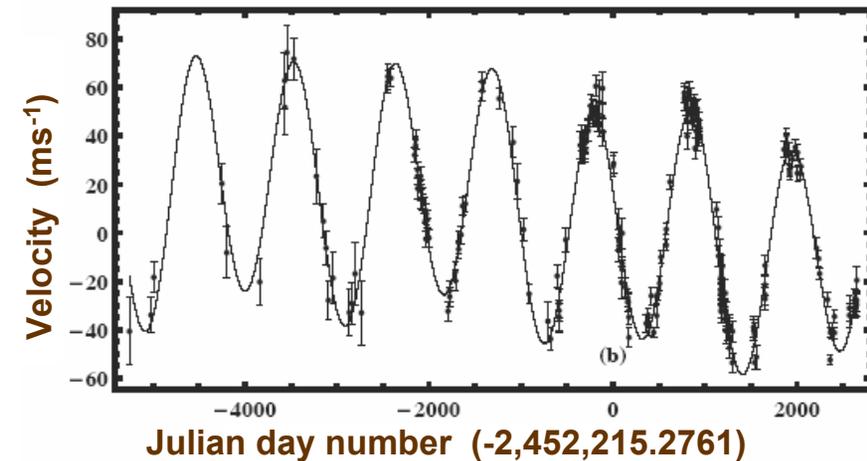
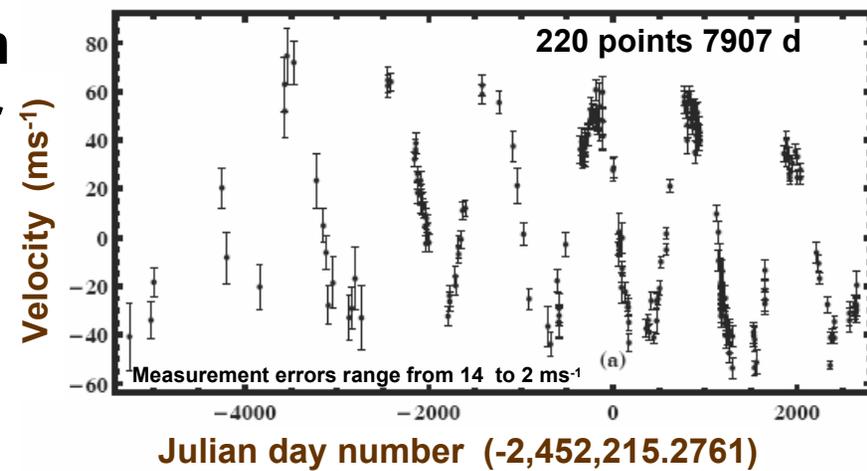
Fischer, D. A., Marcy, G. W., Butler, R. P., Laughlin, G. L., and Vogt, S. S., 2002, ApJ, 564, 1028

- 3) 2004-2009, several papers report either no 2nd planet or a 2nd planet with $P = 9660$ days when ecc. of 2nd planet set = 0.005, value found by Fischer et al., 2002.

eg., Wittenmyer, R. A., Endl, M., Cochran, W. D., Levison, H. F., Henry, G. W., 2009, ApJS, 182, 97

- 4) 2010, Gregory, P. C., & Fischer, D. A.,
MNRAS, 403, 731, 2010

FMCMC confirms 2400 d planet and finds evidence for a 3rd planet of ~10000 d.

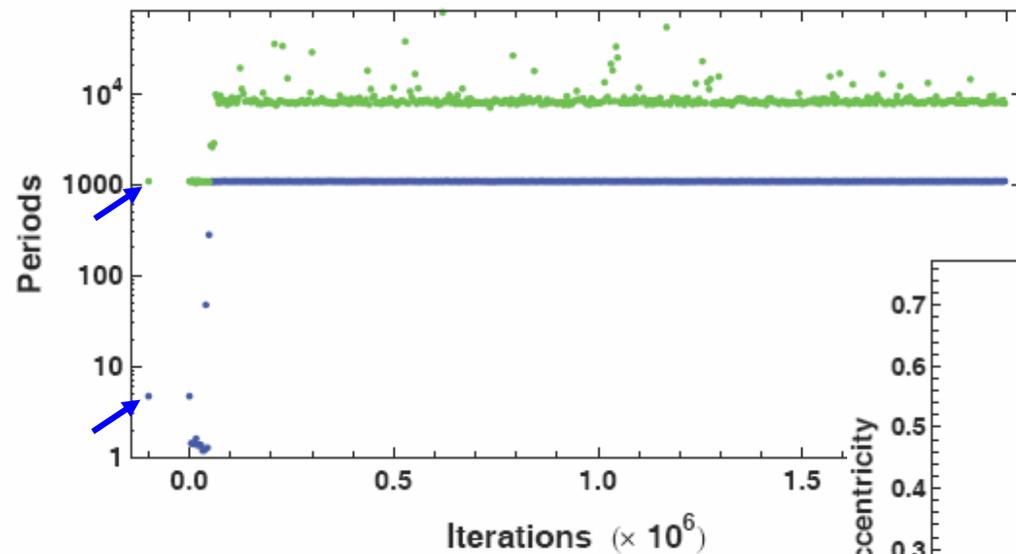
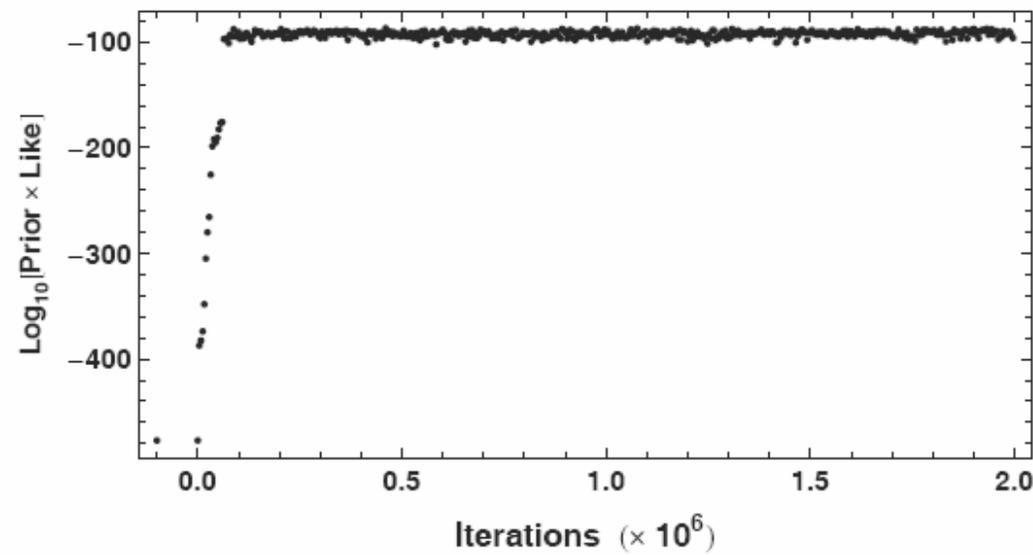


47 Ursae Majoris

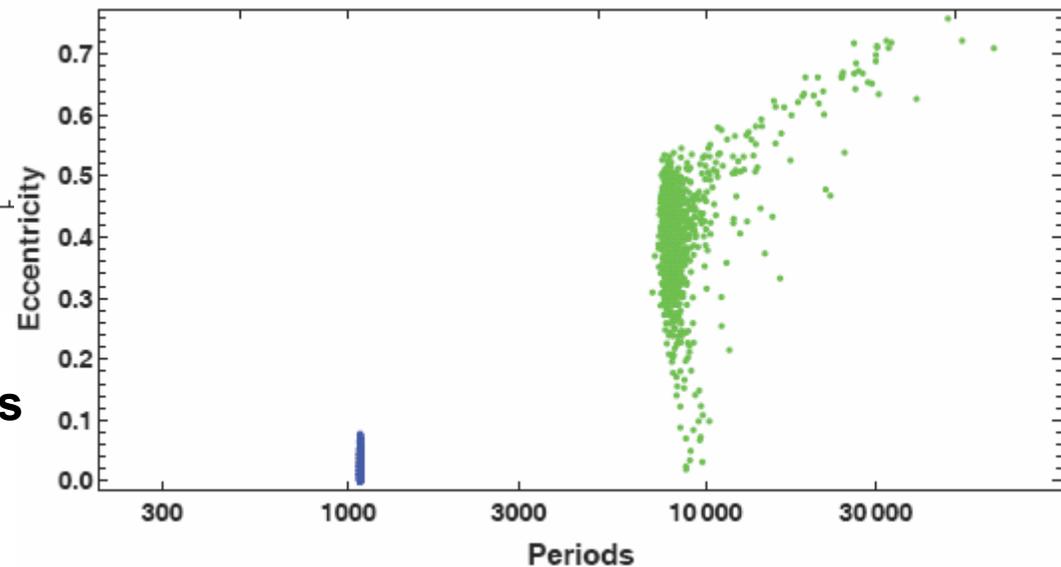
2 planet FCMC model fit

For two planet model we do not detect 2400 d period!

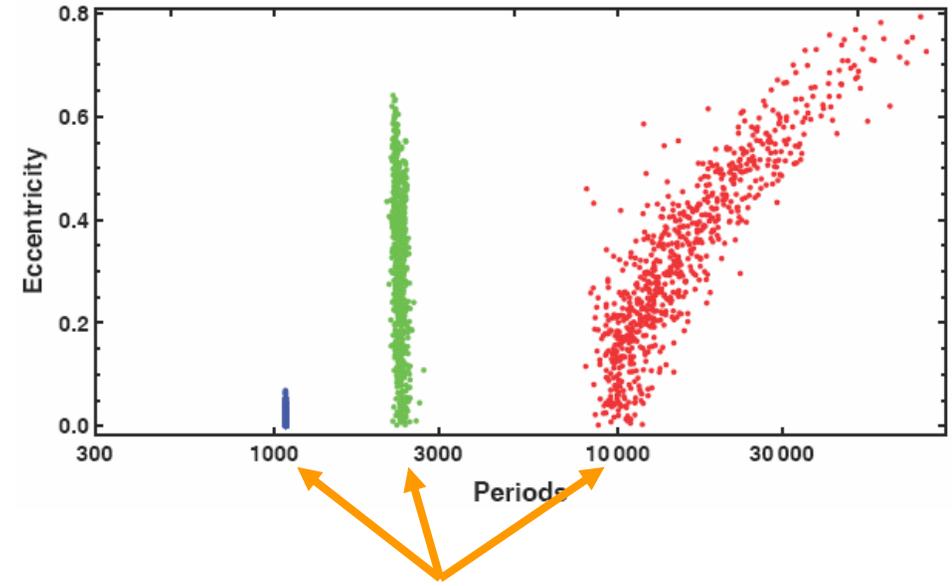
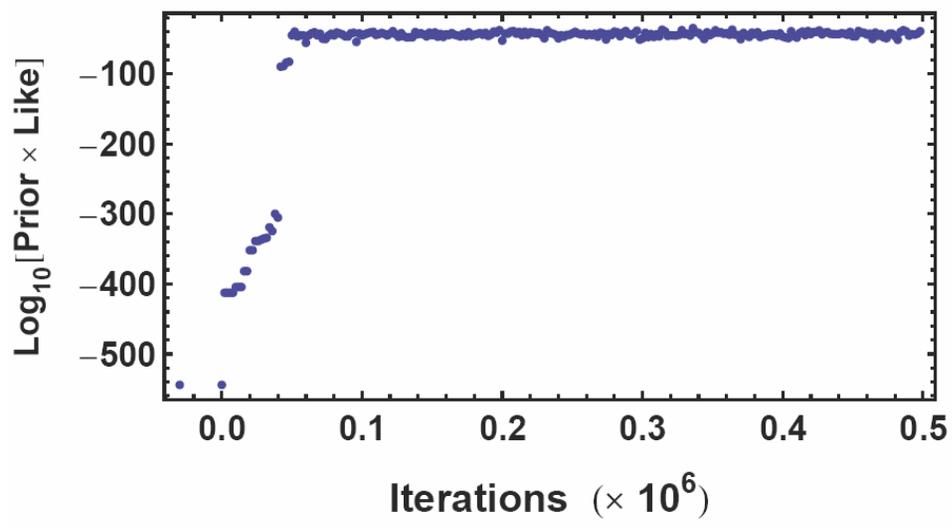
No evidence for 3rd planet in residuals.



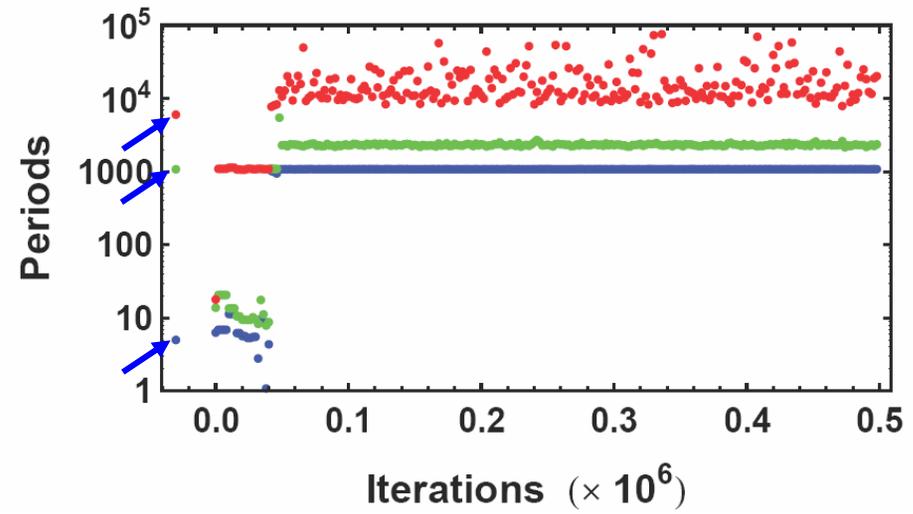
Arrow  indicates starting periods



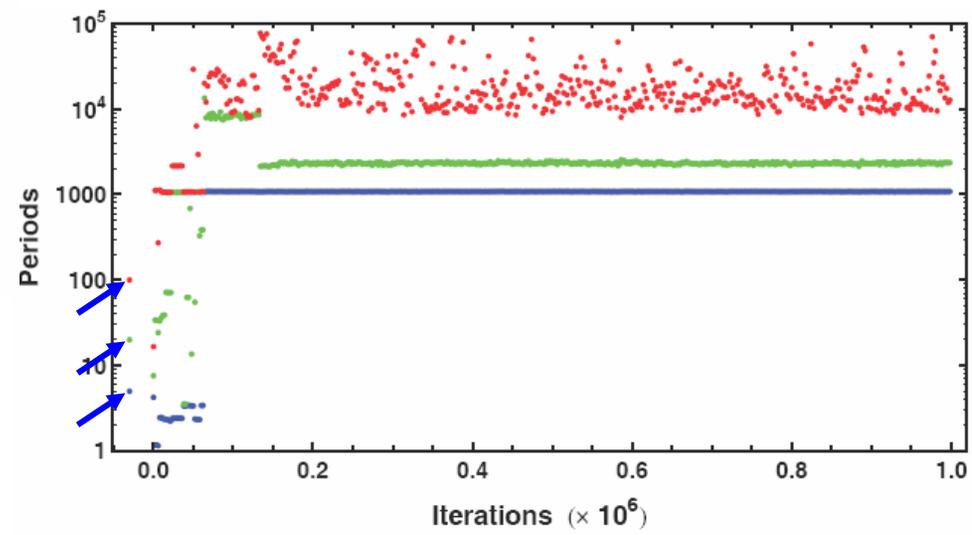
47 Ursae Majoris 3 planet FCMC model fit, Lick telescope data



Evidence for 3 planets
Longest period not well constrained
because period is longer than data span



Arrow  indicates starting periods

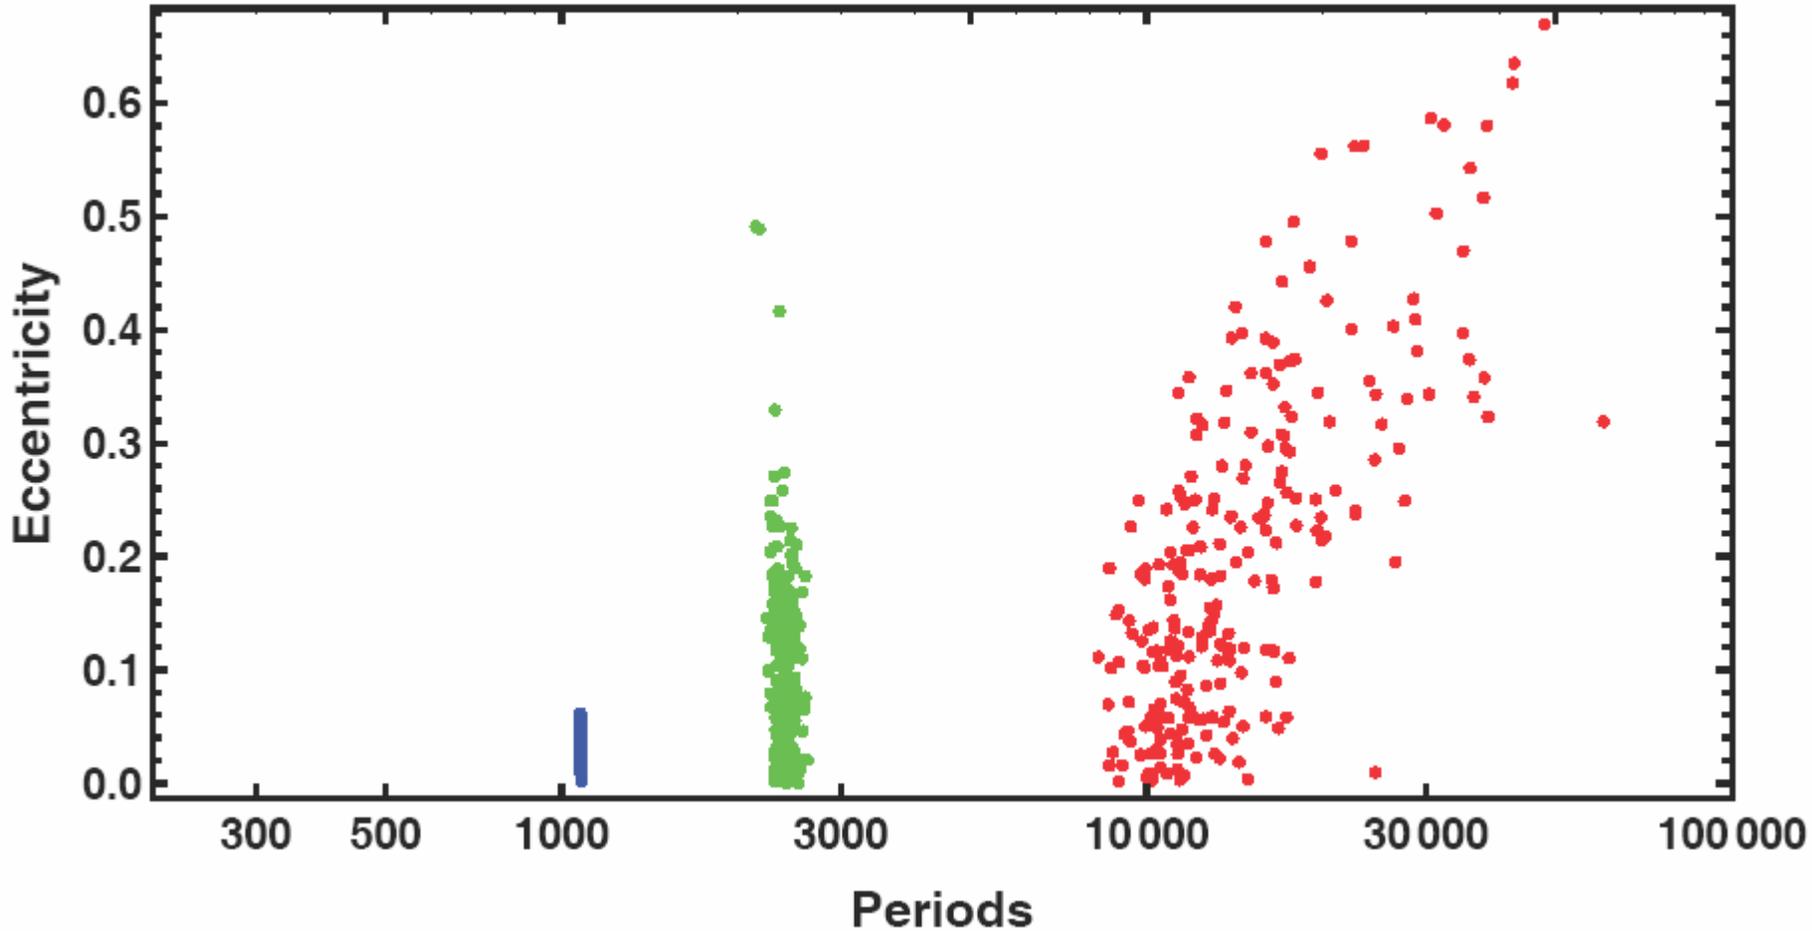


47 Ursae Majoris

3 planet FMCMC model fit

Lick + Hobby-Eberly + Harlem J Smith telescope data.
Included 2 additional residual offset velocity parameters.

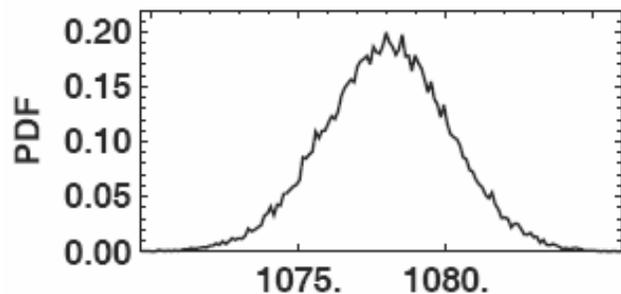
Total of 413 velocity measurements.



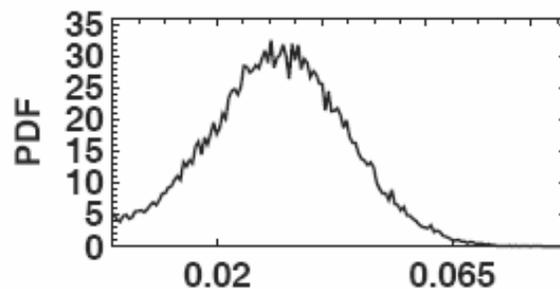
Probability of a 3 planet model is 2×10^5 > than for 2 planet model

47 Ursae Majoris

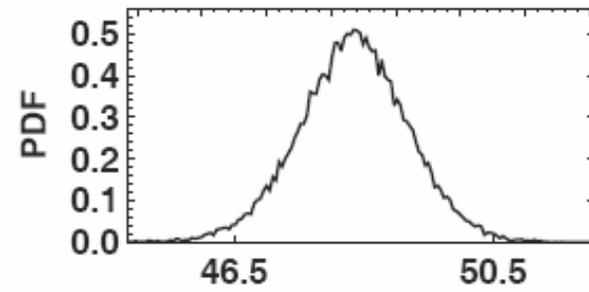
Parameter marginal distributions for a 3 planet HMCMC of the combined Lick, HET, and HJS telescope data set.



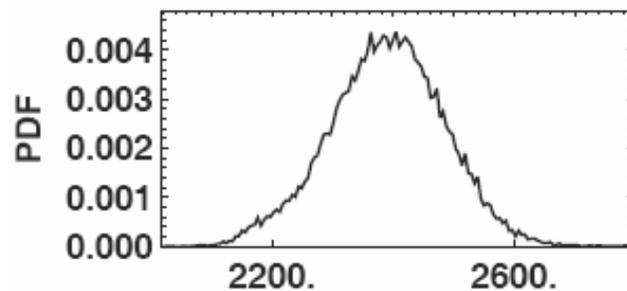
P_1 (d)



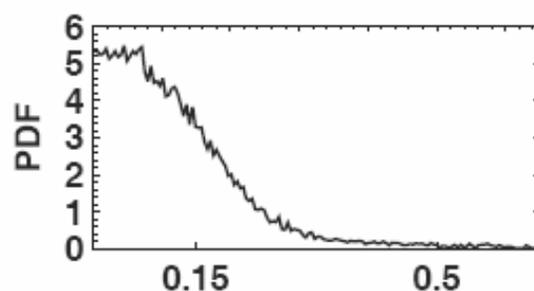
e_1



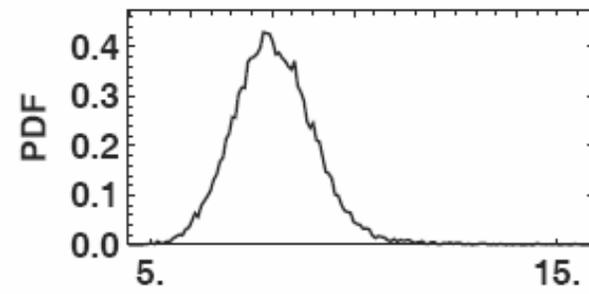
K_1 (m s^{-1})



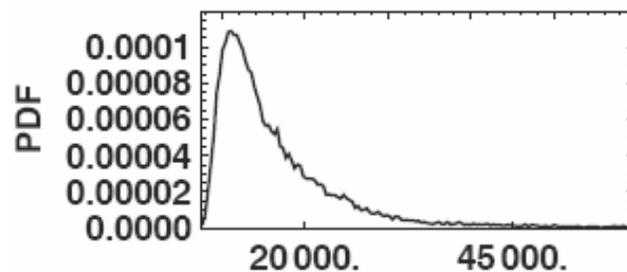
P_2 (d)



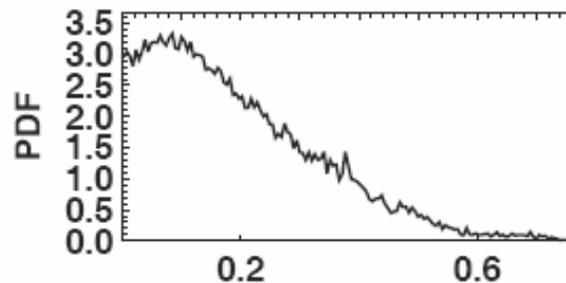
e_2



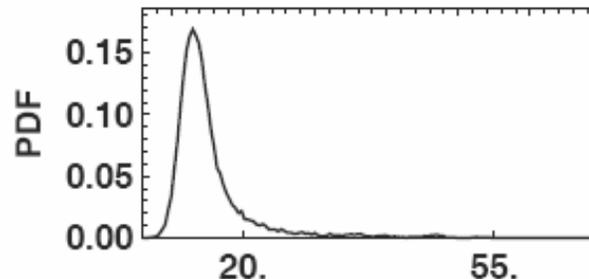
K_2 (m s^{-1})



P_3 (d)



e_3



K_3 (m s^{-1})

Gliese 581 the star with two possible habitable zone planets

History

- 1) 2005 to 2009,
Planet e is 1.9 Earth mass
Planet b is 16 Earth mass
Planet c is 5 Earth mass
Planet d is 7 Earth mass

Latest paper:
M. Mayor et al., *A&A*, 507, p. 487, Nov. 2009

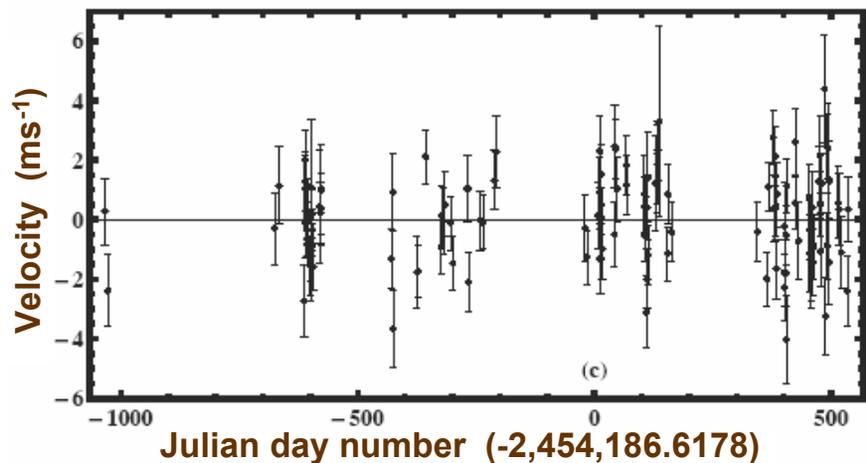
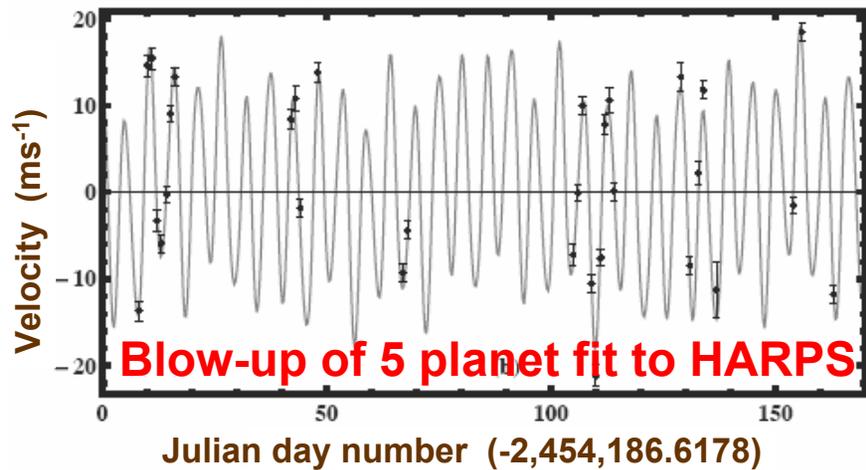
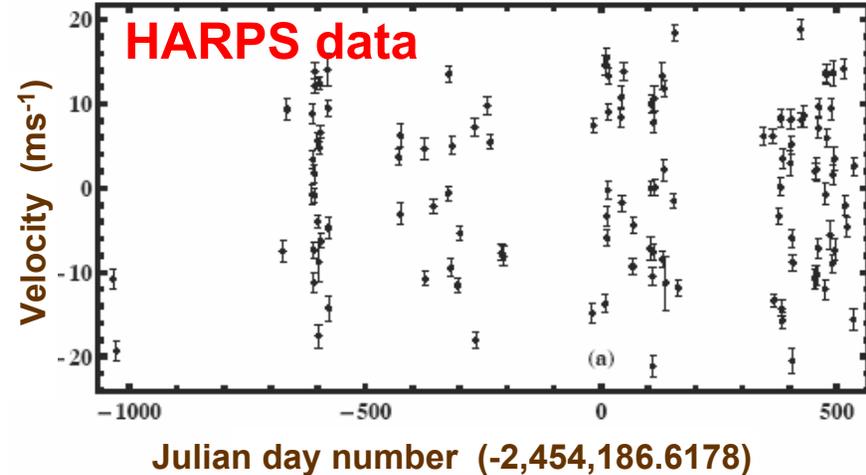
- 2) 2010,
Planet f is 7 Earth mass
Planet g is 3.1 Earth mass

Steven Vogt et al., *ApJ*, 723, p. 954, 2011,
their analysis assumes all circular orbits.

- 3) 2011, Gregory, P. C., *MNRAS*, 415, 2523

Evidence for 5 planets.
Don't support claim for planet 581g.

Find evidence that Keck HIRES
uncertainties are much larger than the
quoted values by an extra 1.8 m s^{-1}
added in quadrature.



Conclusions

1. A Bayesian analysis of the latest radial velocity data sets for 47 Ursae Majoris confirms the existence of a 2nd planet with $M \sin i = 0.54 M_J$ and $P = 2400\text{d}$, and finds evidence for a 3rd planet with a $P \sim 10000\text{ d}$ and $M \sin i \sim 1.6 M_J$.
2. The analysis yields the full marginal posterior probability density function (PDF) for each model parameter, which allows for rigorous error estimates even for a planet with $P >$ data duration.
3. This built in Occam's razor allows for a direct comparison of the probabilities of models with different numbers of planets.
4. This Bayesian approach is helping to raise the bar on how much useful information can be extracted from the radial velocity data.
5. Experience gained while working with the exoplanet data sets has led to the development of an "Adaptive fusion MCMC" nonlinear model fitting algorithm, which can be applied to many other problems.

For copies of my papers please Google Phil Gregory

The rewards of data analysis:

**‘The universe is full of magical things,
patiently waiting for our wits to grow
sharper.’**

Eden Philpotts (1862-1960)

Author and playwright