Outils pour l'analyse de données de vitesses radiales

Nathan Hara

THÈSE DE DOCTORAT

de l'Université de recherche Paris Sciences et Lettres
PSL Research University

Préparée à l'Institut de Mécanique Céleste et de Calcul des Éphémérides, Observatoire de Paris

Outils pour l'analyse de données de vitesses radiales

École doctorale n°127

ÉCOLE DOCTORALE D'ASTRONOMIE ET D'ASTROPHYSIQUE D'ÎLE-DE-FRANCE

Spécialité ASTRONOMIE & ASTROPHYSIQUE

COMPOSITION DU JURY :

M^me^ Françoise Roques
Observatoire de Paris, Présidente

M. Roman Baluev
Pulkovo Astronomical Observatory
Rapporteur

M. Jérôme Bobin
Commisariat à l'Énergie Atomique
Rapporteur

M. Gwenaël Boué
Observatoire de Paris, Membre du jury

M^me^ Magali Deleuil
Laboratoire d'Astrophysique de Marseille
Membre du jury

M. Jacques Laskar
Observatoire de Paris, Membre du jury

M. Stéphane Udry
Université de Genève, Membre du jury

Soutenue par **Nathan HARA**
le 27 octobre 2017

Dirigée par **Jacques Laskar et Gwenaël Boué**

l'Observatoire de Paris | PSL RESEARCH UNIVERSITY PARIS

Abstract

Lorqu'une étoile a des compagnons planétaires, elle décrit un mouvement quasi épicycloïdal autour du centre de masse du système. Si l'orientation du plan de l'orbite le permet, un observateur situé sur la Terre peut détecter la composante de ce mouvement sur la ligne de visée grâce à l'effet Doppler. Il mesure ainsi la "vitesse radiale de l'étoile". Si cette vitesse présente des variations périodiques suffisamment claires, la présence de planètes peut être inférée et leurs orbites contraintes.

Une des difficultés de l'analyse de telles mesures est qu'une combinaison de signaux de plusieurs planètes et de divers bruits peut être confondue avec l'effet d'une planète en réalité inexistante. Après avoir présenté les effets à prendre en compte pour analyser des données de vitesses radiales, nous abordons ce problème. Pour limiter son occurrence, nous utilisons un algorithme de poursuite de base modifié, dont on démontre l'efficacité sur des signaux réels et simulés.

Nous abordons ensuite le problème de l'estimation des paramètres orbitaux pour un système donné ainsi que leur distribution pour une population de planètes. On s'intéresse en particulier à l'excentricité, dont on montre qu'elle est d'autant plus surestimée que le modèle du signal est mauvais. Nous proposons des solutions pour une estimation robuste des paramètres orbitaux.

When a star is orbited by planetary companions, it describes a nearly epicyclic motion around the center of mass of the system. When the orientation of the orbital plane is appropriate, an observer on Earth can measure the velocity of the star along the line of sight by Doppler effect. If this "radial velocity" presents clear enough periodic variations, the presence of planets can be inferred and their orbit can be constrained.

Detection and estimation of orbits is made difficult by the photon noise, the unpredictable variations of luminosity of the star as well as instrumental faults. In particular, signals from several planets can add coherently with the noises and mimic the effect of a planet absent from the system. After listing the relevant effects to make inference on exoplanets from radial velocity data, we tackle this problem. To limit its rate of occurrence, we use a modified basis pursuit algorithm, allowing to search for several signals simultaneously. The efficiency of the method is demonstrated on real and simulated signals.

We then address the problem of orbital parameters estimation for a given system, as well as the estimation of their distribution on a planet population. We look in detail at the eccentricity, and show that its overestimation increases as the model moves away from the correct one. We suggest methods for robust inference of orbital parameters.

Remerciements

Je voudrais tout d'abord sincèrement remercier mes directeurs de thèse, Jacques Laskar et Gwenaël Boué, avec qui j'ai eu beaucoup de plaisir à travailler et au contact desquels j'ai beaucoup appris. Jacques, qui est un grand "compositeur", a aussi été un excellent "producteur" en s'assurant toujours que mes "albums" étaient prêts à arriver sur le marché et qu'ils auraient une audience. Les remarques sur les articles, sur les exposés, le voyage en Suisse, ont eu un rôle crucial dans le déroulement de cette thèse. Gwenaël, qui a toujours pris le temps d'écouter mes idées, même quand elles n'étaient pas encore très claires, et de relire mes travaux avec beaucoup d'attention et de rigueur en y apportant des idées nouvelles. J'essayerai à l'avenir de toujours garder dans un coin de ma tête cette question: "attends, mais tu es sûr que...", qui m'a évité bien des erreurs.

Je remercie Roman Baluev, Jérôme Bobin, Magali Deleuil, Françoise Roques et Stéphane Udry de m'avoir fait l'honneur de participer à mon jury de thèse.

Je remercie les membres de l'équipe ASD, au sein de laquelle il est très agréable de travailler. Alain Chenciner, pour les discussions que nous avons eues sur les principes d'incertitudes, la théorie de l'information et qui a toujours pris le temps de répondre à mes questions. Jacques Féjoz, qui lui aussi m'a aidé à résoudre certaines questions mathématiques tout comme Laurent Niedermann, qui en plus m'a cédé son bureau ! Alain Albouy, qui a dû fournir une bonne partie des centaines de litres de café nécessaires au bon déroulement de ma thèse. Mickaël Gastineau, qui m'a fait gagner beaucoup de temps en résolvant quasi instantanément tous les problèmes informatiques que j'ai rencontrés. Frédéric Dauvergne pour le travail stimulant sur NEAT et le RV Fitting Challenge. Philippe Robutel, avec qui j'ai eu beaucoup de plaisir à discuter autour d'un café ou d'un tableau, tout comme Hervé Manche (le café en moins, la guitare en plus). J'inclue Alexandre Correia aussi dans l'équipe ASD. Merci à lui pour son intérêt et pour avoir participé à l'article.

Les doctorants et stagiaires, Jessica Masseti, Alexandre Pousse, Adrien Leleu, Pierre Auclair-Desrotour, Thibault Castan, Timothée Vaillant, Léo Bernus, Antoine Petit, Santiago Barbieri et Louis Guillot. Le fait de pouvoir aller se voir de temps en temps pour réfléchir à une question, prendre un café ou un verre, aller en coupole, a rendu les journées plus créatives et plus agréables.

Je remercie Agnès Patu, Djamila Houibi, Amélie Muslewski et Rachida Ahmidez pour l'organisation des voyages et pour leur efficacité. Pour Jérusalem, Agnès, vraiment merci !

Merci à l'équipe de l'observatoire de Genève, Rodrigo Diaz, Jean-Baptiste Delisle, Xavier Dumusque, Rosemary Mardling, Damien Ségransan, et particulièrement Stéphane Udry, pour avoir pris le temps de travailler ensemble. La semaine que Jacques et moi avons passée là-bas a été très enrichissante.

Merci à mon coloc et ami de longue date Thibault de Poyferré pour les quatre années de bonne camaraderie, et pour avoir éclairci certains points de maths. Impossible de ne pas mentionner les six autres de la photo, pour qui les mots me manquent.

Merci à mes parents, mes frères et soœurs, qui m'ont toujours soutenu dans tous les sens du terme. Sans eux je n'en serai certainement pas arrivé où j'en suis. Qu'ils soient assurés de ma grande gratitude et profonde affection.

Merci à ma chère Emmanuelle, qui m'a toujours épaulé et a supporté avec patience les aléas d'une vie de chercheur. Merci au Atch.

Contents

List of Figures

List of Tables

Far better an approximate answer to the right question, which is often vague, than an exact answer to the wrong question, which can always be made precise.

John Tukey

Chapter 1

Introduction

Little details have special talents in creating big problems.
Mehmet Murat Ildan

1.1 First observations

In the middle of the XIX$^{\text{th}}$ century, Jacob (1855) interpreted abnormalities of the 70 Ophiuchi star's orbit as either non-universality of the law of gravitation, or a planet orbiting the star. Unfortunately, both claims were wrong. Nearly a century after, one, possibly two, Jupiter mass companions were reported around Barnard's star (van de Kamp 1963, 1969), also later dismissed. In 1988, a search for substellar companions lead Campbell et al. (1988) to conjecture that 7 out of 15 stars that were observed could have a companion in the 1-9 Jupiter mass range, but remained cautious about the discoveries. A bona-fide detection was finally made in 1989 with the discovery of an eccentric companion of HD114762 (Latham et al. 1989). This detection was done through radial velocity measurements, themselves possible thanks to Doppler effect. As the star moves back and forth on the line of sight because of its companion, the observer sees periodic shifts of the spectral lines of the star. Such measurements yield the mass projected on the line of sight and thus a minimum mass. In this case, $m \sin i$ was estimated to be $11 M_J$. The authors concluded that this object is likely to be a brown dwarf, but also considered the possibility that it be a giant planet.

The next major discovery was obtained by timing of the millisecond pulsar PSR B1257+12 (Wolszczan & Frail 1992). Pulsars are neutron stars that rotate rapidly with a very steady frequency. Since they emit radio waves in a fixed direction with respected to their attached reference frame, an observer on Earth receives a periodic signal. If the pulsar moves in a non uniform way, its motion translates to modifications of the reception time of the pulses. In particular, if another body orbits a pulsar, one will see its reflex motion through the periodic shifts of the signal received. With this technique two companions were detected. The precision of the measurements even allowed to see the effect of mutual interaction between the planets. The author leveraged this information to estimate the mass of the two planets, which were 4.3 ± 0.2 and 3.9 ± 0.2 terrestrial masses.

The field of exoplanetary sciences was truly kicked-off by the discovery of a giant planet around 51 Pegasi, with a minimum mass approximately half of Jupiter's and a period of

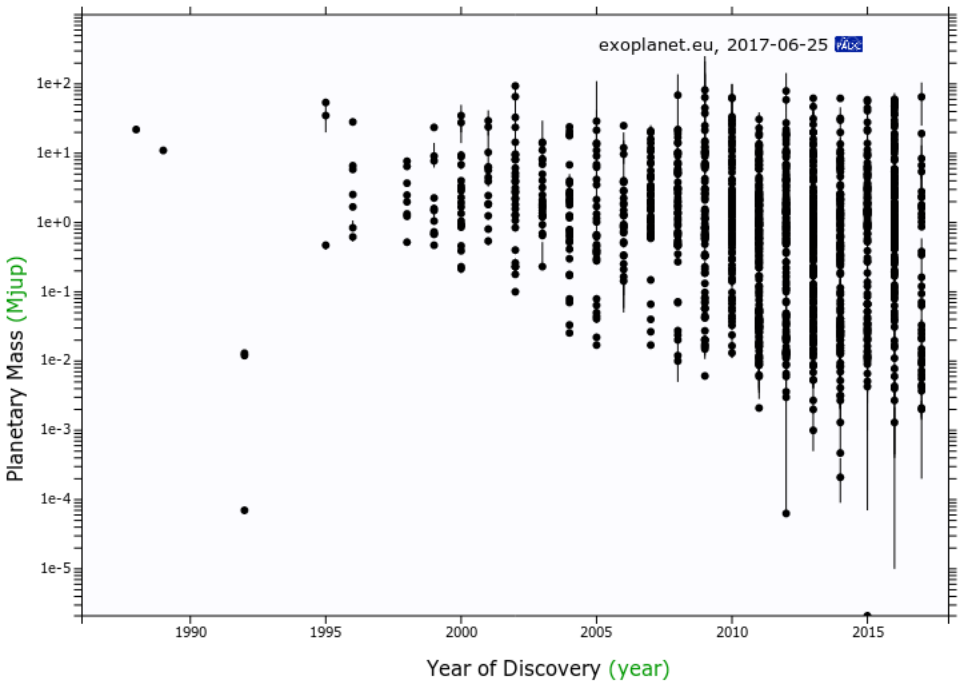

Figure 1.1: Masses of exoplanets as a function of their year of discovery (generated on exoplanet.eu).

4.2308 days (Mayor & Queloz 1995). The fact that the planet was massive and close to its host star was key in its detection, and was also very startling at the time as giant planets were expected to be beyond the snow line. Since this discovery, there have been detections of planets with smaller and smaller projected mass with the radial velocity technique, as shown in figure 1.1.

Among those, 51 Peg b is not the only puzzling object. Other giants were found close to their stars, but also systems with architectures possibly very different from the Solar System. These discoveries are used to assess formation scenarios of planetary systems. Planets in the so-called habitable zone are searched for, with the remote hope of finding life.

We would not want to make mistakes on the typology of planets in the universe in the first place. But since the measurements taken have strong consequences on our understanding of the physics of the planets, it is all the more important to analyse the measurements carefully. First not to be mislead by spurious detections, and also not to miss subtle hints in the data that could turn out to be crucial. The purpose of this thesis is to devise tools that could be useful to extract information from radial velocity measurements. We will present mainly two analyses, one for searching for candidate periodicities for planets (chapter 2) and one about the robustness of the orbital parameter estimates, focusing on eccentricity (chapter 3). In the present chapter, we will attempt to outline the aspects that should be taken into account in the process of making a claim from the data. This includes the sources of signal we are aware of, but also the inference process itself and the pros and cons of signal processing techniques.

1.2 Radial velocity signals

In this section, we present radial velocity data and the various phenomena that shape it. A recap chart summarizing the known sources of signal, typical amplitudes and time-scales is provided at the end of this section in table 1.1. The effects we present are the object of a wide literature, our presentation does not aim at entering into their details, but rather to focus on their signature in the data.

1.2.1 Motion of a star with planetary companions

The effect we wish to detect is the star wobbling due to potential companions. Let us first describe this motion in its simplest form. We consider a planetary system in its inertial frame, whose origin is its barycentre, and the z-axis pointing to the observer. In figure 1.2, we show two examples of the motion of a star, (left) and the corresponding radial velocities as function of time (right). The upper panel shows a case of star orbited by a single planet. The middle one shows the same quantities when three non coplanar planets orbit the star. We show on the bottom panel the projection of this latter motion onto the plane x, y, which is what ideal astrometric measurements would look like (without noise, star motion nor parallax).

The motion of the star represented is the sum of the motions obtained by considering the two body problems { Star, planet k } for k = 1,2,3. This approximation exactly comes down to neglecting the interactions between planets, which is valid in most cases, though there are a few detected systems where the gravitational interactions can be seen. For instance, four planets were detected in GJ 876 (Correia et al. 2010), and among those two giant planets in 2:1 mean motion resonance at 30 and 60 days which strongly interact. Remarkably, the outer planet orbits at 124 days, such that the system is in a Laplace resonance (4:2:1) (Nelson et al. 2014). Starting from best fit values of the orbital elements, the numerical integration on 100 years integration yields a complicated motion of the planets, instead of steady ellipses (see figure 1.3). The orbits of the giant planets are represented in red and blue. We will now stick to the sum of Keplerian motions, which are good approximations in general, at least over the time-scale of observations.

To be able to analyse the radial velocity measurements, we need a mathematical model of the signal, which can be obtained from Newton laws of motion. Let us first consider a planet J of mass m orbiting a star S of mass M, supposed to be known, as represented figure 1.4. Denoting by $\overrightarrow{r} = \overrightarrow{SJ}$ and $\mu = G(m + M)$, where G is the gravitational constant. The equation of motion is (Newton 1687)

$$\frac{d^2\overrightarrow{r}}{dt^2} = -\mu \frac{\overrightarrow{r}}{\|\overrightarrow{r}\|^3}, \tag{1.1}$$

which is a second order differential equation on the three-dimensional space. From Cauchy-Schwartz theorem, we need to specify a six dimensional vector (position and velocity) as initial condition to obtain a uniquely defined solution. Here, the parameter μ is not considered as an unknown given that it is well approximated by $\mu = GM$ which is only a function of the stellar mass, the latter being considered as known. However, to reproduce observational data, we do need an additional parameter because we only have access to the motion of the star relative to the inertial frame whose position is given by $\overrightarrow{GS} = -(m/(m + M))\overrightarrow{r}$. It thus

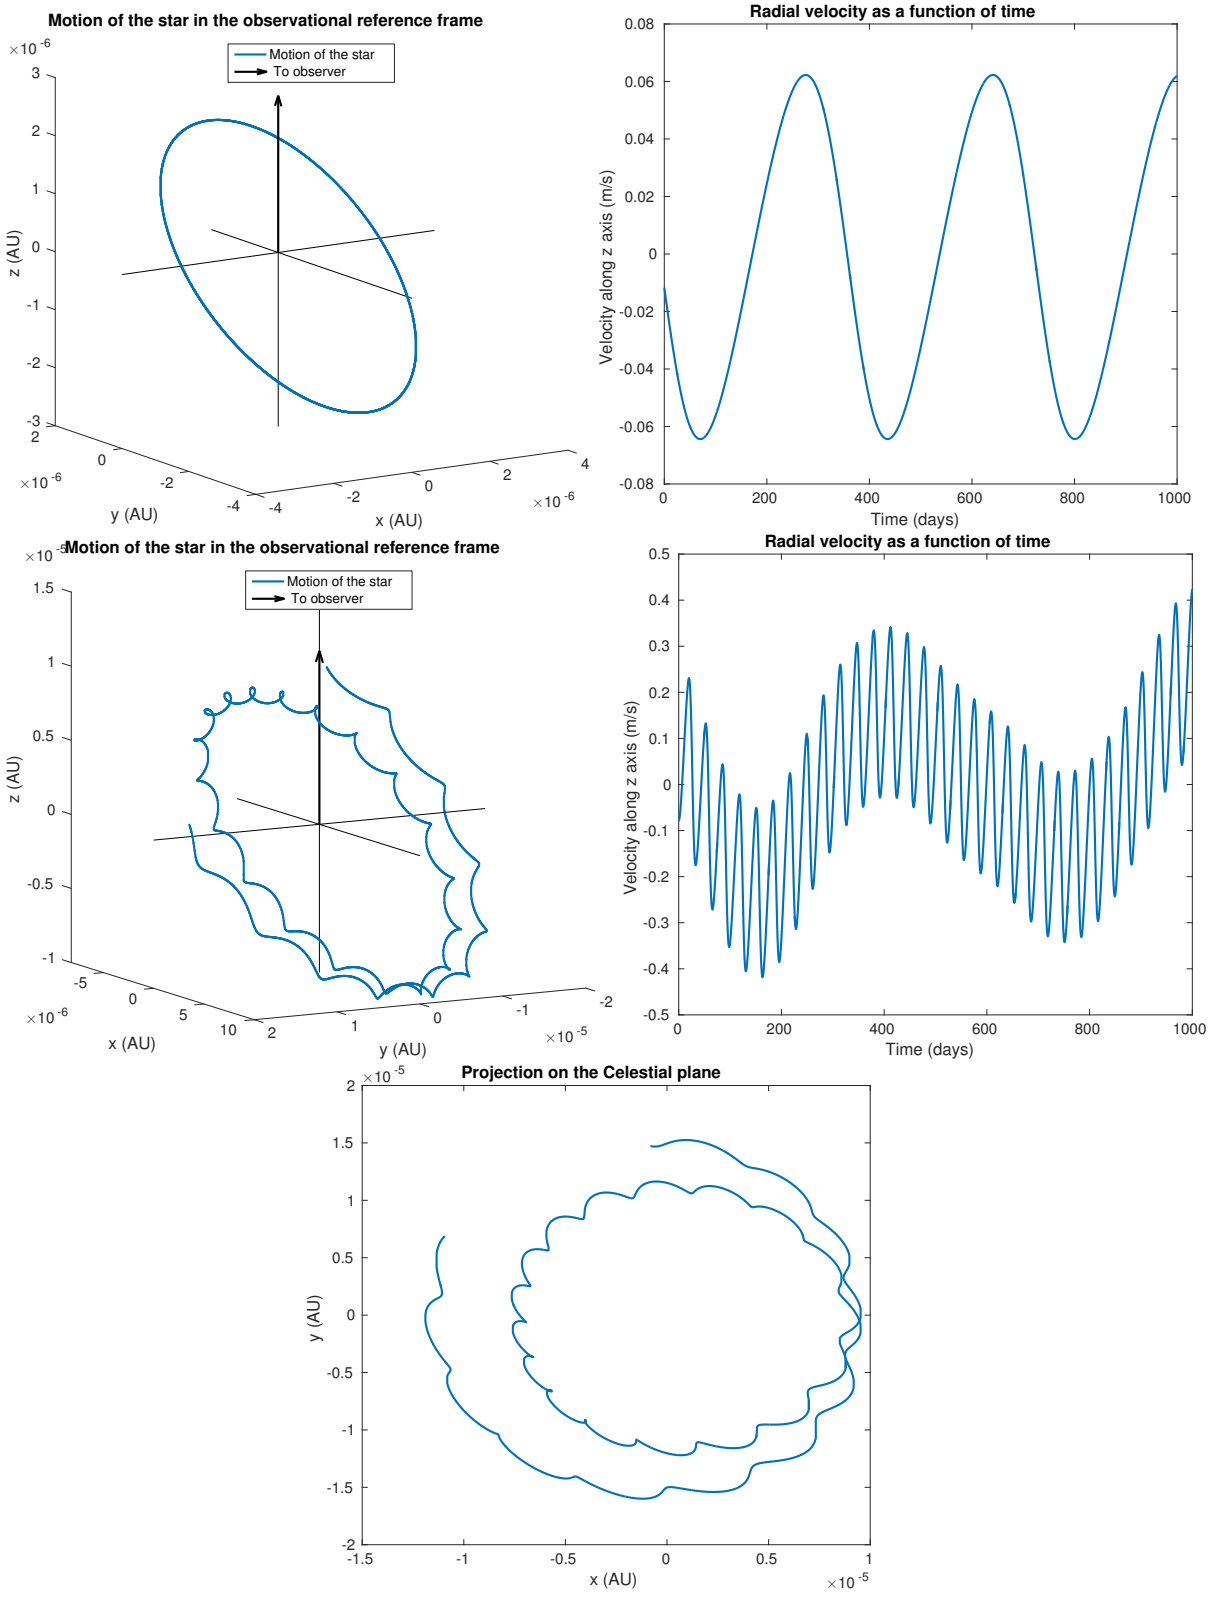

Figure 1.2: Idealized motion of the center of mass of a star orbited by planets (left) and the projection of the velocity along the line of sight (right). The upper plots represent this motion in the case of a single planet orbiting a star of 1 $M_\odot$ with mass 1 $M_\oplus$ at 1 AU and eccentricity $e = 0.1$. The middle plots represent the same quantities with additional planets of terrestrial mass at 0.2 and 1.3 AU, the bottom plot represents the projection of this motion on the x, y plane.

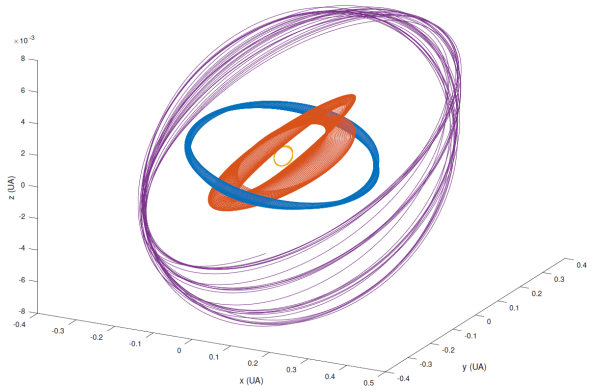

Figure 1.3: Simulated motion of planets whose orbital elements are taken initially as best fit values of GJ 876

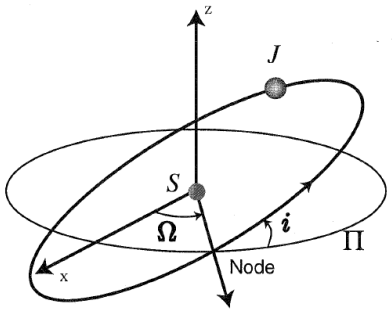

Figure 1.4: Parametrization of the motion of a planet J around a star S. Π is the plane of the sky (courtesy of J. Laskar)

depends on the planet mass m, which means that a set of seven parameters $\boldsymbol{\theta}$ completely specify the motion of the star and the position of the planet at any time of measurement t.

As will be discussed later on, when searching for planets it is convenient to first look for periodicity in the data. However, the initial conditions of the equation at an arbitrary time and the mass of the planet do not make the periodicity appear. Instead of considering those, we will use a representation of the planet motion with elliptical elements, more convenient for our purposes. These are classically defined as

- m: mass
- e: eccentricity.
- i: inclination.
- ω: argument of periastron.
- Ω: right ascension at ascending node.
- P: period or equivalently the mean motion $n = 2\pi/P$.
- M_0: mean anomaly at a reference time.

The angles i, ω and Ω are defined with respect to an inertial reference frame as shown figure 1.4, which is taken as the one defined above : the z axis points to the observer, the x and y directions are chosen so that x, y, z is a direct orthogonal basis. It is important to specify the direction of x when dealing with astrometric measurements but it is unnecessary for radial velocity analysis. The semi major axis a of the orbit is given by the third Kepler's law of motion $n^2a^3 = G(m + M)$.

It is easy to show that when the orbital elements are defined as such, the velocity of the star projected on the z axis is (Perryman 2011) (with a minus sine for consistency of the

definition of elliptical elements),

$$-\dot{z}(t, e, K, P, \omega, M_0) = K(\cos(\omega + \nu(t, e, P, \omega, M_0)) + e\cos\omega) \tag{1.2}$$

$$\cos\nu = \frac{\cos E - e}{1 - \cos E} \tag{1.3}$$

$$\sin\nu = \frac{\sqrt{1-e^2}\sin E}{1 - \cos E} \tag{1.4}$$

$$E - e\sin E = M_0 + \frac{2\pi}{P}t \tag{1.5}$$

$$K = \left(\frac{2\pi G}{P}\right)^{\frac{1}{3}} \frac{m\sin i}{(m+M)^{\frac{2}{3}}} \frac{1}{\sqrt{1-e^2}} \tag{1.6}$$

where t is the time, ν is the true anomaly. Unfortunately, the projection onto one of the axis causes the disappearance of some information: Ω does not appear anymore. Equation 1.6, from Cumming et al. (1999), shows that the inclination i and the mass of the planet m are degenerated, only the projected mass $m\sin i$ is knowable. Overall, five parameters describe the radial velocity of a star.

If several planets orbit the star, the radial velocity in the inertial reference frame attached to the system is

$$\dot{z}(t) = \sum_{k=1}^{n_p} \dot{z}_k(t) \tag{1.7}$$

where $\dot{z}_k(t)$ denotes the velocity due to planet k. In that case, $5n_p$ parameters are needed to represent the motion of the star.

Since the star moves across the galaxy, the direction pointing to the observer is not fixed. The reference frame of the target rotates with respect to the observer and is therefore not exactly inertial. When observing a circular planet of mean motion n_1 and assuming the target rotates around the Sun with mean motion n_2 (which is a local approximation), it is easily shown that the apparent mean motion of the planet is $n_1 + n_2$. Since n_2 is extremely small compared to n_1, it can be safely neglected. Therefore an observer in an inertial reference frame such as the Solar System Barycentric Reference frame observes $\dot{z}(t) + V(t)$ where $V(t)$ is the radial component of the target reference frame and $\dot{z}(t)$ is given by expression 1.2.

The main effect not taken into account so far is the motion of the Earth. The displacement of its center of mass as well as the rotation contribute to the relative velocity of the star and the observer. The removal of Earth motion components is achieved with planetary ephemeredes such as Standish (1990) or Fienga et al. (2008). The associated corrections are best explained when considering the actual measurements, which is the object of next section.

1.2.2 Doppler spectroscopy

The spectrum of a star in the visible domain contains absorption lines. These are short intervals of wavelengths such that photons whose energy lies in this range are absorbed in the upper parts of stars atmospheres. As a result, the observer does not receive light at those wavelengths. The spectrum appears as a continuum of frequencies where some are missing

(see figure 1.5). As the star moves, the wavelengths are shifted by Doppler effect and so are the absorption lines, whose spacing is approximately conserved.

The absorption lines act as a comb that can be compared to reference wavelengths to derive the shift of the spectrum. Since the instruments are not entirely stable, even if the star was perfectly motionless there would be a drift of the spectrum, thus the reference wavelengths must be calibrated. This is done with two techniques: with an iodine cell put on the path of light, that imprints its absorption lines on the spectrum (e.g. Marcy et al. 1997) or by a stable lamp, whose emission lines are used to calibrate the CCD detector (Baranne et al. 1996). For instance in the case of SOPHIE, the calibration is performed with a Thorium-Argon lamp (Perruchot et al. 2008).

In the case of a calibrating lamp, the observed spectrum is correlated to a template which is shifted in wavelength as shown figure 1.5. At each wavelength shift between the template and the spectrum, the correlation is computed (represented by the total area in pink). The correlation as a function of the wavelength shift is called the cross correlation function (CCF), an example of which is given figure 1.5, bottom right, from Mayor & Queloz (1995). The CCF is then fitted by a Gaussian function whose mean is the estimated velocity of the star.

The obtained wavelength, $\lambda_{\rm obs}$, has some contributions from the motion of the Earth. To isolate the contribution of the star, observers compute the wavelength that would have been observed in the International Celestial Reference System (ICRS) which is an approximation of inertial frame whose origin is the barycentre of the solar system (Rickman 2001). The corrected wavelength λ_B is

$$\lambda_B = \lambda_{\rm obs} \frac{1 + \frac{1}{c}\boldsymbol{k}.\boldsymbol{v}_{\rm obs}}{1 - \frac{\Phi_{\rm obs}}{c^2} - \frac{v_{\rm obs}^2}{2c^2}} \tag{1.8}$$

where $\boldsymbol{k}$ is the vector pointing from the observer to the target in the ICRS, $\boldsymbol{v}_{\rm obs}$ is the velocity of the observer, $\Phi_{\rm obs}$ is the gravitational field he experiences. As said above, the velocity of the observer is obtained from the Earth ephemeris (Standish 1990; Fienga et al. 2008). The time of observation should also be corrected (Capitaine et al. 2003). The velocity of the star along the line of sight $\boldsymbol{k}.\boldsymbol{v}_\star$ is then obtained by equating λ_B with the expression of wavelength of the light emitted and the velocity of the target in the ICRS

$$\lambda_B = \lambda_0 \frac{1 + \frac{1}{c}\boldsymbol{k}.\boldsymbol{v}_\star}{1 - \frac{\Phi_\star}{c^2} - \frac{v_\star^2}{2c^2}} \tag{1.9}$$

where $\Phi_\star$ is the gravitational field at its surface. Since we are interested in relative variations over time the $\Phi_\star$ term can in general be dropped (Lovis & Fischer 2010). Note that as these corrections do not depend on the wavelength, nothing is lost by applying them to the mean wavelength computed via the CCF rather than to each wavelength individually.

Measuring a displacement of 1 m.s^{-1}, for a spectrograph whose resolution is 100,000 requires to detect displacement as small as 1/3000 of a line width, that is circa 1/1000 of a CCD pixel (Lovis & Fischer 2010). To achieve such precision or better, which is routinely done by HARPS for instance, thousands of lines are needed to average the noise out. Building such instruments comes with technical challenges, whose effects on the signal are outlined next section.

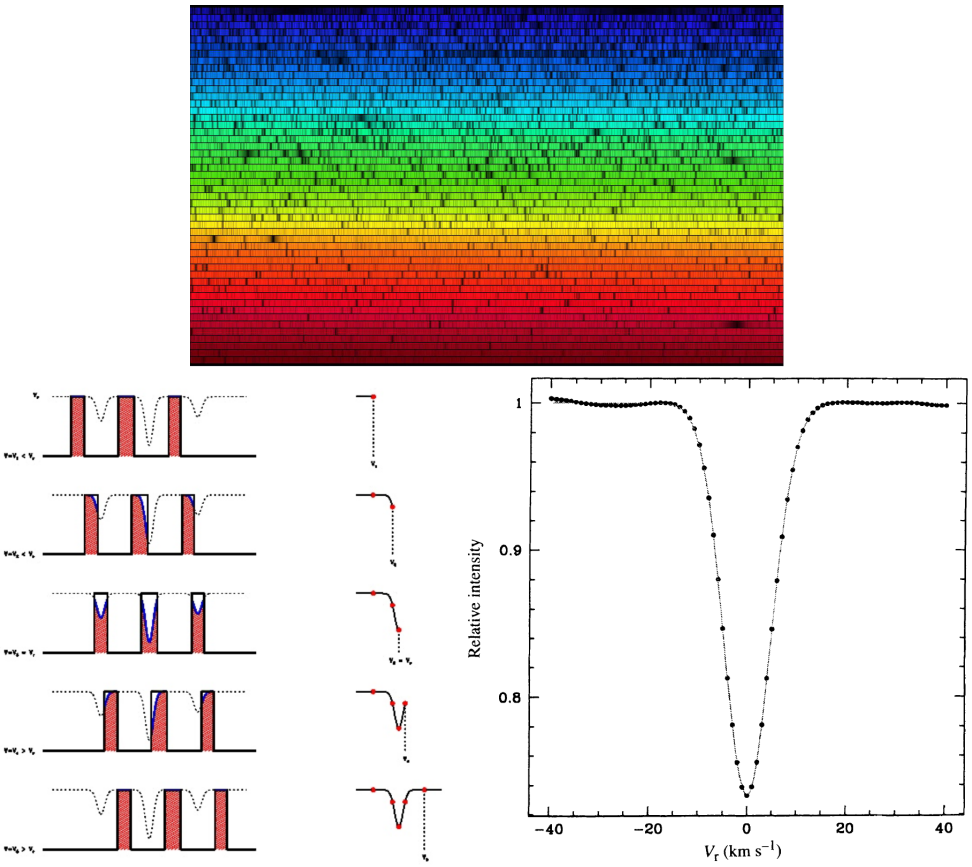

Figure 1.5: Top: Solar spectrum from 392 nm to 692 nm (https://www.cfa.harvard.edu/ssp/stars_planets/solarspectrum.html), bottom left: Correlation between a template spectrum and an actual spectrum (from Melo (2001)) and an example of cross correlation function (CCF) (Mayor & Queloz 1995)

1.2.3 Instrumental effects

A thorough presentation of the spectrograph used (échelle grating) is out of the scope of this work. The presentation of instrumental considerations will focus on the impact on the signal of the instrument properties. In other words we shall see to which extent the measured spectra differ from ideal ones.

Like in any optical device, the photon noise is a source of uncertainty. Its effect can be quantified using a Poisson process model for the arrival times of the photons. With this modelling, one can express the average and standard deviation of the number of photons received per wavelength (Pelat 2013). Let us denote by $n_t(\lambda)$ the true average number of photon per unit time and per wavelength, and by Δt the time during which the shutter is open (the integration time). The number of photons received at wavelength λ will be on average $n_t(\lambda)\Delta t$ and its standard deviation $\sqrt{n_t(\lambda)\Delta t}$. This value is assimilable to the uncertainty on the true light flux.

On top of this issue, if the spectrograph is not in a vacuum chamber, the variations of pressure and temperature lead to a change of refraction index of the air which change the light

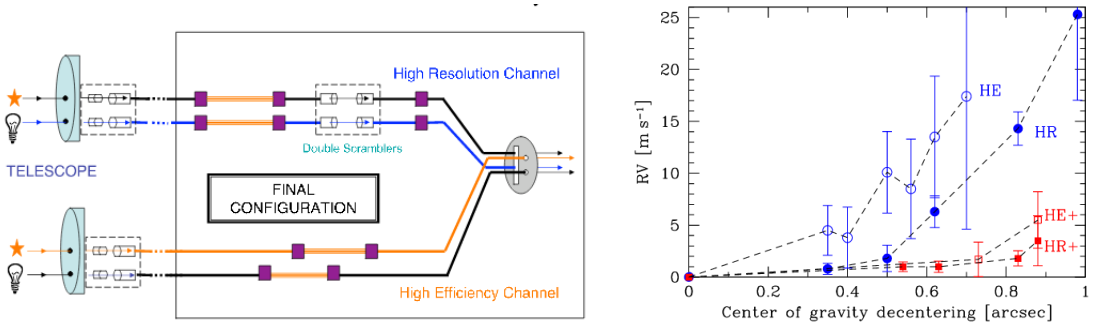

Figure 1.6: Schematic of the SOPHIE optical fibers feeding the slit (left). Variation of the estimation of radial velocity before and after introducing octogonal fibers. Figures from Perruchot et al. (2011) (left) and Bouchy et al. (2013) (right)

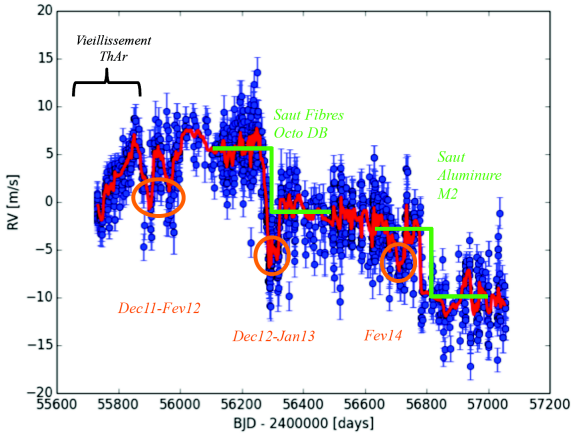

Figure 1.7: Radial velocity of known stable targets as a function of time obtained with SOPHIE (courtesy of F. Bouchy)

paths and induce some large deviations of the estimated velocity (100 m.s^{-1}/K or m.s^{-1}/mbar). Changes of temperature lead to modifications of the relative positions of the optical devices that might induce changes of optical properties such as the PSF and cause variations of the estimated radial velocity up to hundred of meters. High precision instruments such as HARPS are in a void chamber and in pressure and thermally regulated rooms. The estimation of radial velocity depends on the position of the photocenter (that is the barycenter of the entering light weighted by intensity) at the entrance of the slit. Such changes might result from modifications of the seeing, atmospheric condition, telescope focus or guiding etc. As the photocenter moves the spectral lines are shifted and the estimation of the radial velocity varies. In the case of spectrographs whose slit is fed by an optical fibre, the amplitude of this shift can be reduced by placing fibers with a non circular cross section. This is what has been done for the spectrograph SOPHIE (Perryman 2011), where fibers with octagonal sections have been introduced on the light path, as shown figure 1.6. Figure 1.6 (right) shows the variation of the radial velocity as a function of the offset of the photocenter before and after the change for both channels of the instrument (respectively in blue and red, the channels

being HE and HR). The new device produces radial velocity (RV) which are sensitive to the range of offset up to 5 $m.s^{-1}$, as opposed to 25 before change.

The precision of the instrument must be stable on timescales greater than the orbital periods to allow a detection. The calibration by iodine cells or Thorium-Argon mentioned in the last sections cannot be perfect. Figure 1.7 shows the evolution through time of the estimated radial velocity of targets that are known to be particularly stable. The period from 55700 to 55800 encompassed by the black accolade displays a drift as the Thorium-Argon lamp gets older.

This figure shows other several interesting feature, especially two jumps highlighted in green. Those happened respectively after the introduction of octagonal optical fibres and after the mirror of the telescope has been re-aluminised. The monitoring of stable stars allows to reduce the effects of such offsets which are hard to predict. The averaged measures (red curve) is subtracted from other radial velocity measurements to improve the stability of their distance to a stable reference.

Finally, let us mention that the variation of size and sensitivity of CCD pixels. Variations of sizes between pixels distort the wavelength solution from calibration . Also, the read out noise increases as the signal to noise ratio decreases, introducing some systematic effect.

Contamination

The spectrum measured by the spectrograph can contain information from a star in the field of view indistinguishable from the target star (e.g. Cunha et al. 2013, and reference therein). The measured spectrum can also contain signatures from the sunlight reflected by the Moon (Bonomo et al. 2010), this is especially a problem when the star is close to the Moon and the seeing is poor. Finally, the absorption lines of the atmosphere are not subject to the Doppler effect that shifts the stellar spectrum and might vary in depth. This problem is especially important for spectrographs operating in the infra-red region such as CARMENES or Spirou where the atmospheric absorption spectral lines are more numerous.

In HARPS, the spectral regions where the atmosphere has the deepest lines are simply neglected in the calculations (Mayor et al. 2003). As the performance of the instruments improves, the errors induced by shallower lines becomes more important and other, more sophisticated options are considered. Some are based on forward modelling of the absorption (e.g. Bailey et al. 2007; Seifahrt et al. 2010; Cotton et al. 2014; Cunha et al. 2014). In another spirit, Artigau et al. (2014) has used a non parametric method to distinguish stellar and atmospheric lines (Principal Component Analysis).

As an example, figure 1.8 shows the impact on simulated radial velocity time series of ignoring the spectral content in certain regions when the spectral lines are deep or shallow and the ignored regions are determined by a threshold. When telluric lines are well identifiable, thresholding the spectrum seems to be an acceptable technique, while this type of correction leaves systematic effects on the RV estimation for shallower lines that have amplitudes comparable to stellar lines.

1.2.4 Stellar types

The precision of the radial velocity estimate depends on the morphology of the spectrum. Since the estimation is obtained by averaging the displacements of the spectral lines, spectra

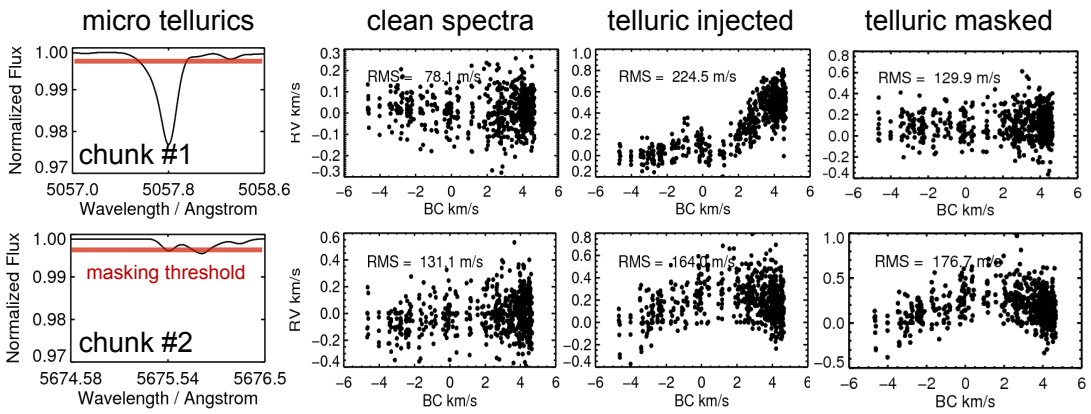

Figure 1.8: Spectra and extracted RVs for two example $2\mathring{A}$ spectral chunks, plots from left to right: micro telluric lines within the chunk (note depths), extracted RVs for > 700 observations with no telluric lines injected, with telluric injected, and with telluric injected but masking out the telluric-contaminated pixels when extracting RVs. x-axis is the barycentric correction velocity of the star. Figure and caption from Sharon Wang's poster at the Extreme Precision Radial Velocity meeting II, 2015, Yale University.

with numerous and stable lines provide more accurate estimates. More precisely, Connes (1985) and Bouchy et al. (2001) compute the best error achievable on the radial velocity taking into account photon noise and a Gaussian error on the flux of pixel i due to the faults of the CCD defaults of standard deviation σ_D [1]

$$\sigma_{\text{RV}} = \frac{c}{Q\sqrt{F\Delta t}} \;\; ; \;\; Q = \sqrt{\sum_{i=1}^{n} \frac{\lambda_i^2 \left(\frac{\partial A_0}{\partial \lambda}(\lambda_i)\right)^2}{A_0(\lambda_i) + \sigma_D^2(i)/\Delta t}} \tag{1.10}$$

where F is the total incident flux (total number of received photons by time unit), Δt is the integration time and Q is the quality factor. $A_0(\lambda)$ is the theoretical average number of photons per unit time received at wavelength λ normalized by the total flux, which therefore depends only on the spectral type and λ_i is the wavelength corresponding to pixel i of the CCD detector (or the center of pixel i). The quality factor contains the squared derivative of the spectrum's derivative. It means that the higher the slope at λ, the more information is carried by the spectrum at this wavelength.

In practice, one wants to avoid too hot stars ($T_{\text{eff}} \gtrsim 10,000K$), whose photosphere elements are mostly ionized and therefore do not have lines in the visible spectrum. A fast rotation tends to broaden lines, decreasing the value of the spectrum derivative and are not suitable targets for highest precision measurements. Finally coldest stars ($T_{\text{eff}} \lesssim 3500K$) tend to have overlapping lines and lower flux, therefore lower signal to noise ratio. Precise measurements are available for $F5$ to $M5$ stars, as well as the red giant and clump region of the HR diagram (Lovis & Fischer 2010).

[1] To obtain the result, Connes (1985) invokes calculations from Connes 1984, which is not publicly available. Appendix A discusses this calculation and makes it more explicit.

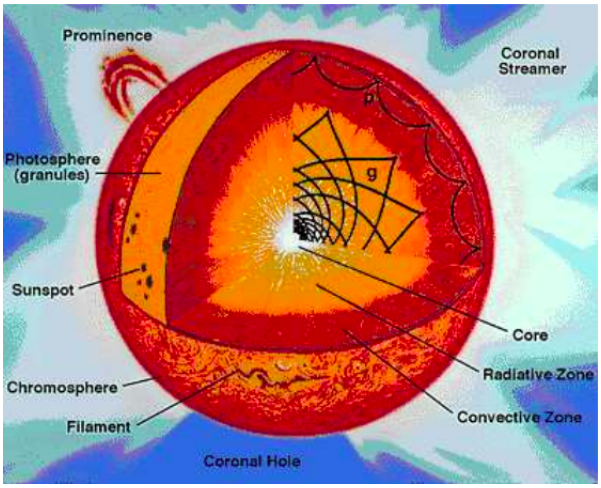

Figure 1.9: Structure of a Sun-like star. Credit: NASA

1.2.5 Stellar noise

The HARPS and HARPS-N spectrographs (resp. Mayor et al. 2003; Cosentino et al. 2012) can achieve a precision on radial velocity of 1 - 0.5 $m.s^{-1}$. This value is expected to be improved to 0.1 $m.s^{-1}$ by ESPRESSO in the near future. At this level of precision, the main obstacle to detection and characterization of exoplanets is the stellar noise. Since the subject is crucial, it will be treated in a little bit more detail.

(Extremely) simplified structure of a star

Giving a fair description of stellar structure is way beyond of the scope of this work. Interested readers might refer to the classical textbooks (Phillips 1999; Prialnik 2000) for an introduction to stellar physics. We will only define terms that are essential to understand the RV signatures of stellar features.

A star might be roughly divided in concentric layers, as shown figure 1.9, that have different properties. The core of the star, that concentrates most of its mass, is where hydrogen fusion takes place and generates thermal energy. This one is transmitted outwards by radiative transfer in the so-called radiative zone. After a random walk of ≈ 10 Myr, the photons reach the upper part of the radiative zone. At this point, the transfer of heat is done through convective transfer, hot gas rising up and colder gas rising down due to buoyancy. On top of the convective zone is the photosphere, where light is emitted outwards and which can be seen as the surface of the star. The uppermost layers are the chromosphere and finally the corona, which is a plasma envelope that surrounds stars . In the case of the Sun, the ejection of Coronal mass is known as Solar wind. Note that this description is invalid for coolest stars ($< 0.35\ M_{\odot}$), where the convective zone can extend all the way to the center (Hansen & Kawaler 1994).

The presence of a magnetic field in sunspots had been noted already by Hale (1908), and its origin remains a subject of active research. In the case of the Sun, the physical process at the origin of the magnetic field (the dynamo) is likely to be shearing forces at the interface of the convective and radiative zone (the tachocline, Spiegel & Zahn 1992). The shearing forces

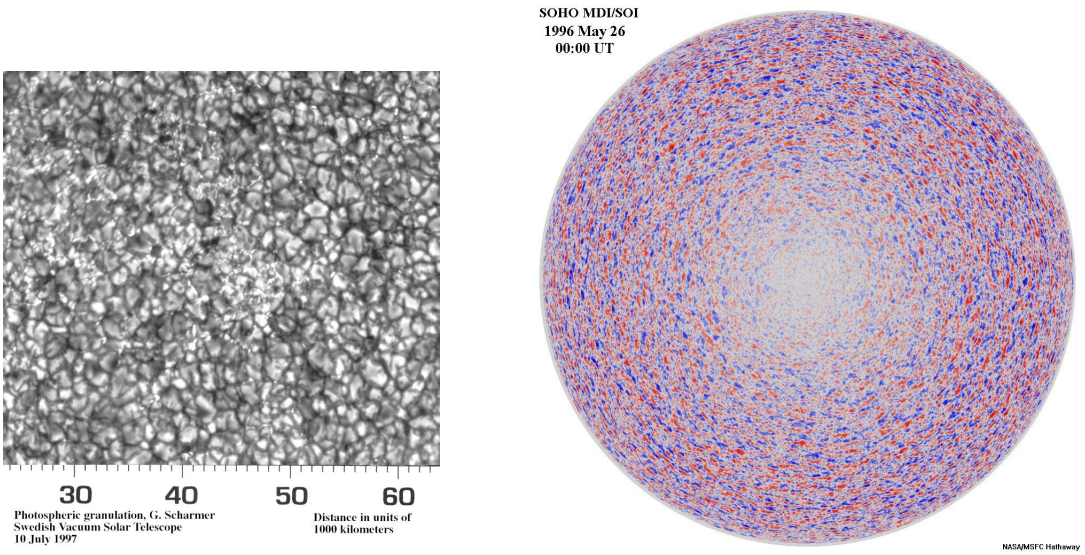

Figure 1.10: Granules (left) and super granules (right) on the surface of the Sun (NASA, https://solarscience.msfc.nasa.gov/feature1.shtml)

would be resulting from the interactions between the solid rotation of the core and radiative zone, as well as the variable rotation rate as a function of radius in the convective zone. This leads to expect different structures of magnetic fields in lower mass stars (Chabrier & Küker 2006; Walder et al. 2012). For a discussion of the physical origin of the solar dynamo see Charbonneau (2010). We will restrict ourselves to list the consequences of the magnetic field on the observation.

Acoustic waves

Stars are crossed by pressure waves that arise from inhomogeneities in the convective zone, whose study is termed Asteroseismology (e.g. Christensen-Dalsgaard 2014). Some of these waves cause radial oscillations of the star that have an effect on the measurements, as shown by Dumusque et al. (2011b). The typical time-scale of the p-modes oscillation is 5 to 15 min and they produce RV variations up to a few m.s^{-1}. Dumusque et al. (2011b) also shows that this effect can be mitigated by choosing an integration time $\gtrsim$ 10 min in order to average it out.

Granulation

In the convective zone, the hot gas rises up until it reaches the photosphere, cools down and goes backwards in the direction of the star centre. This means that the gas whose velocity is pointing to the observer is on average hotter than the descending one and results in an overall blueshift of the surface. Figure 1.10 (left) shows a picture of the Sun's surface taken by the Swedish Vacuum Solar Telescope, where the color level scales the level of the light flux. The surface exhibits cells of $\approx$ 1000 km that are brighter in the center, where the gas emerges, and darker on the rim where gas goes down. These have a lifetime of approximately 20 min, after which other granulation cells appear.

Convection cells of greater scales are present on the Sun. These phenomena are termed supergranulation for cells of ≈ 40,000 - 50,000 km, which last for ≈ 1 day (Kueveler 1983; Del Moro 2004) and mesogranulation for cells whose size and lifetime is in between (Deubner & Fleck 1989; Roudier et al. 1998). The blueshift and redshift induced by supergranules is represented figure 1.10 (right).

Before the characteristics of granulation were precisely known, Harvey (1985) suggested a simple statistical model for granulation. The surface of the Sun is supposed to be a sum of a great number of cells whose velocity behave like independent processes with exponential autocovariance. These ones have a characteristic time τ_s and standard deviation σ_s that depends on the scale of the granulation s. Averaging these processes, he obtains (without providing calculation details) that the noise induced at scale s is modelled by a stochastic process of power spectrum

$$P_s(\nu) = \frac{4\sigma_s^2 \tau_s}{1 + 2\pi\nu\tau_s}. \tag{1.11}$$

Harvey also suggests the activity-induced noise can be modelled by such a process. This model has been used, along with a model of p-modes frequency to fit the spectrum of five stars densely sampled (Dumusque et al. 2011b).

Spots, faculae and plages

Sunspots are known to arise from magnetic fields at the surface of the sun that blocks convective flux. The corresponding regions of the photosphere are 600 to 1800 K cooler than their neighbourhood, thus appear darker. By extension, any region on the surface of a star that is cooler due to magnetic activity is called a starspot (Schrijver 2002). Strassmeier (2009) (and reference therein) shows that, alike their solar analogs, the lifetime of starspots grows with their initial size (Petrovay & van Driel-Gesztelyi 1997). But unlike sunspots, they might form in stellar poles (Hussain 2002) and last longer (on average a year, up to 4.5 years (Strassmeier & Hall 1994)). Furthermore, spots are associated with magnetic fields that inhibit partly the convection, therefore the convective blueshift, resulting in a net redshift (Cavallini et al. 1985).

Sunspots are surrounded by bright points called faculae which also come from magnetic fields. These are ≈ 100 K hotter than the average on the photosphere, but this effect on radial velocity is not dominant (Dumusque et al. 2014). Alike Spots, the associated magnetic field limits convection, which reduces the blueshift (Lagrange et al. 2010). Those are tubes tangent to the stellar surface and therefore they appear brighter on the limbs. Faculae are associated with other structures situated in the chromosphere called plages, through processes that are not completely understood (e.g Schrijver 2002). These have a spectral signature on the Ca II K line, which allows to detect their presence, hence magnetic activity (Bruzek 1977). Furthermore, plages introduce a small difference in flux and appear brighter and brighter in contrast with their neighborhood as they approach limbs (Chapman et al. 2001).

The main effect of spots and faculae on radial velocity estimate is the following. When the spot is on the half of the star moving towards the observer, there is a deficit of blue-shifted light, therefore a global red shift. As it passes on the receding half the situation is inverted, there is a global blueshift of the velocity. This periodic motion can mimic a planet signature at the stellar rotation period or an harmonic and lead to false detections as noted by Queloz

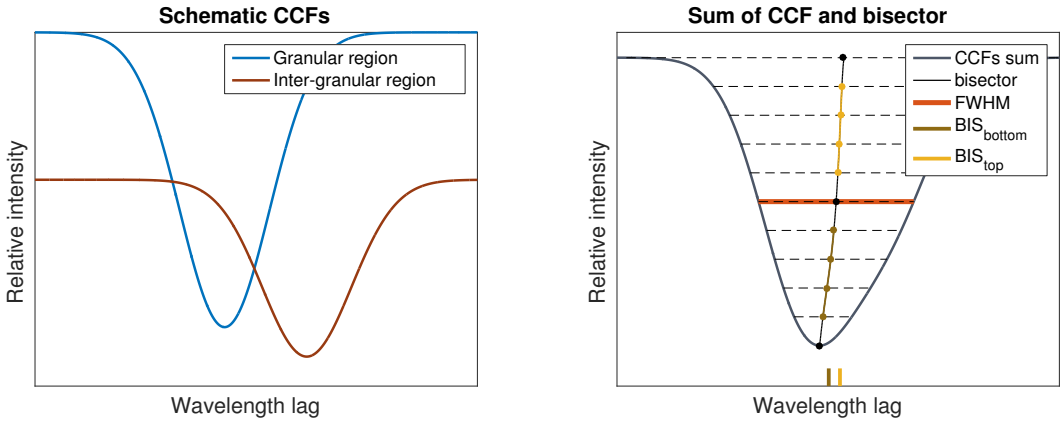

The CCF is not only useful to estimate the radial velocity. The various phenomena at the surface of the star also affect its shape. For instance, the contribution of the center of the convection cells, blue shifted, will be superimposed with the contribution from their edges, red shifted. These contributions integrated on the stars are schematically represented in the figure above (right). The observer measures the sum of the two, which is broadened and asymmetric. The broadening and its variations are seen by computing the full width at half maximum (in red, right figure). Queloz et al. (2001) suggests to compute the asymmetry as follows. First, for each level of relative intensity I , one computes the mid-point of the line joining the two CCF points at I (round dot) . Then one computes the average of midpoints on the 10 - 40 and 60 - 90 % upper parts of the intensity range, which are represented by dark and light yellow check-marks on the wavelength axis. The bisector span is the difference between these two values, $\mathrm{BIS} = \mathrm{BIS}_{\mathrm{top}} - \mathrm{BIS}_{\mathrm{bottom}}$.
Other indicators can be derived from the CCF whose merits are investigated in Santerne et al. (2015). They find that the so-called BiGauss (Nardetto et al. 2006) and W_{span} (Figueira et al. 2013) are most sensitive to line profile asymmetry. Note that thes are most useful when using a box-shaped template to compute the CCD (Eggenberger & Udry 2010).

Figure 1.11: Indicators derived from a spectrum

et al. (2001). Fortunately, differences between active and inactive stars can be seen on the cross-correlation function (CCF) defined section 1.2.2.

The signatures of spots, faculae and plages has been investigated with numerical simulations by Boisse et al. (2012) (SOAP software) and Dumusque et al. (2014) (SOAP 2.0). The star surface is subdivided in cells, each of which have a CCF that depends on whether it is occupied by a quiet photosphere, a spot or a plage (no mention is made of faculae), and the velocity of the cell. The CCF is obtained as a weighted sum of the cells CCFs whose weights are given by a limb-darkening model. The CCF models are obtained by using the HARPS reduction pipeline with a solar spectrum of a spot and of a quiet photosphere. The CCF

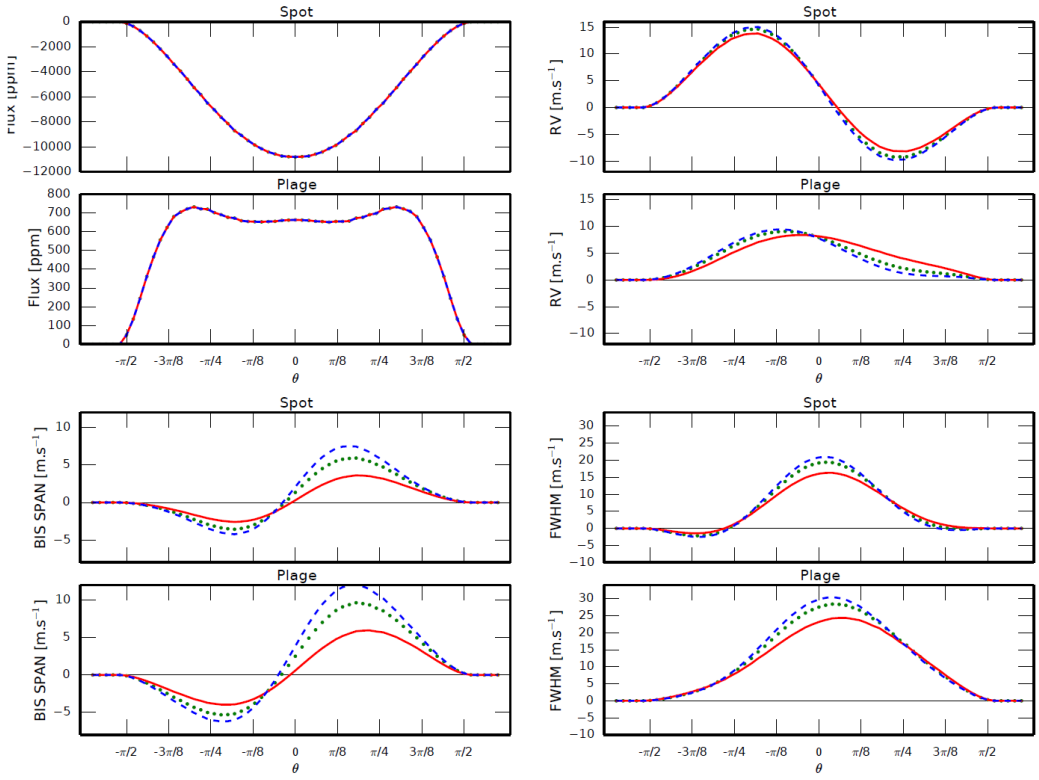

Figure 1.12: Simulation of SOAP 2.0 (from Dumusque et al. (2014)) Flux, radial velocity, FWHM and bisector span as an equatorial spot or a plage moves across the stellar disk. The different colors correspond to different instrument resolutions. R > 700,000 (blue, ≈ resolution of the Fourier Transform Spectrograph, FTS), R = 115,000 (green, ≈ resolution of HARPS), R = 55,000 (red, ≈ resolution of CORALIE)

of active regions serve as template for spots and plages, which are identical in that respect. The difference lies in their temperature models, which affects their weight in the sum. They compute also other indicators from the CCF (FWHM and BIS), whose definition is given in figure 1.11. We reproduce a sample of their results figure 1.12, which is the variation of the indicators mentioned as well as photometric flux for a single spot or plage on the equator as it moves across the stellar disk. The good news is that indicators are correlated with the RV variations induced by the active regions. Since RV variations should not have signatures in activity indicators, activity indicators allow in principle to disentangle planets from activity. This is not totally true as those are noisy and more complex than idealized cases. Secondly, the spot pattern varies in time in a stochastic fashion as we shall see next section.

Long term activity

If the spots were motionless with respect to the surface of the star and of constant flux, their signature RV would be a periodic signal at the star rotation period. This one would be exactly defined by its Fourier coefficients. These are useful characterization of the spot effects, but because of stochastic apparition and disappearance of spots the global effect is

more complicated.

Figure 1.13 shows measurements of the area covered by spot as a function of time and latitude (Hathaway 2010), as well as the total area covered by spots as a function of time. The upper plot exhibits a clear pattern termed "butterfly". The spots appear at $\approx 30 - 35°$ at minimum activity and migrate downwards with a periodicity of eleven years. Magnetic cycles are not observed only on the Sun. The monitoring of the CA II H & K lines of 1296 Sun-like stars (Th HK survey) revealed that 60% show periodic variations of 7 to 30 years (Baliunas et al. 1998), 25% showed variability without clear periodicity and 15% seemed quiet. More recently Sanchis-Ojeda & Winn (2011) and Sanchis-Ojeda et al. (2011) showed that one can recover spots location from their occultation by a planetary companion. In the same spirit Llama et al. (2012) could reconstruct some butterfly-like diagrams from Kepler observations. Furthermore, Lockwood et al. (2007) showed that on old stars, Ca II H & K are correlated positively with chromospheric activity, which is a sign of an activity dominated by faculae, while on younger star the correlation is negative, indicating a spot-dominated activity. The line between those two behaviour is situated at $\log R'_{HK} = -4.96$.

Lovis et al. (2011a) showed that the variation of activity on these long time scales could lead to radial velocity variations of 25 $\mathrm{m.s^{-1}}$ on the time scale of the cycle. However, the very fact that these discoveries were possible shows that magnetic variability is a traceable quantity. Furthermore, it was shown that individually, spots are correlated with FWHM and BIS, and long-term magnetic activity is known to impact RV measurements insofar the number of spots, faculae and plages vary in number, size and location across the star. Still, if the variations of the number of spots is close enough to a planet signature we could be mislead. In order to correct RV time series, we would like an indicator that informs us closely on the contribution of activity to the radial velocity estimate. Dumusque et al. (2012) further noted that the variation of $\log R'_{HK}$ are correlated with radial velocity variation. More recently, Haywood et al. (2016) observed the surface of the Sun, onto which spots can be resolved simultaneously with HARPS observations of its light scattered on Vesta. It shows in particular that the full-disc magnetic flux density might be a better indicator of activity.

Flares

Sudden ejection of coronal mass due to breaking of magnetic fields line, known as flares, can produce variations of estimated radial velocity up to tens of meters per second Reiners (2009). These are rare events, accompanied with an increase of brightness and have a clear signature on the H_α line.

Gravitational redshift

As seen section 1.2.2, the wavelength of the light is red-shifted by the gravitational field of the star. A change of radius of the star means a change of gravitational field on the photosphere, therefore a fluctuation of the estimated radial velocity. Cegla et al. (2012) explored with simulations the possibility of such variations. They found that the magnetic fields act as an inhibitor of convection, by trapping and cooling down the gas going upwards. This results in a decreased effective radius, whose fluctuations can cause a RV jitter of 0.3 to 4 cm/s. Secondly starspots are known to be depleted compared to the average surface of the star (known as the Wilson effect, Bray et al. 1965). The overall modification of the radius due to the appearance

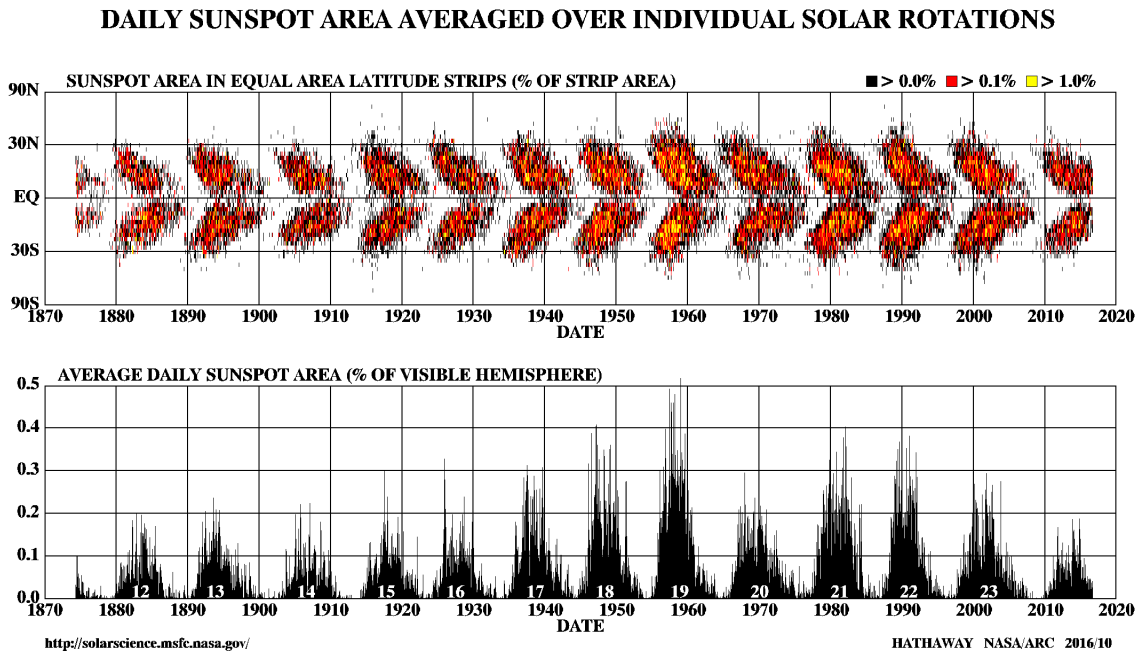

Figure 1.13: Area covered by Sunspots by latitudes of equal area as a function of time (top). Area covered by sunspots averaged per day, as a function of time. From Hathaway (2010)

and disappearance of starspots should amount to ≈ 2 cm/s.

1.2.6 Pipeline

A measurement of radial velocity is the result of a computation starting with the number of photons measured by CCD detectors and numbers from the calibration step of the instrument information. Algorithms that automate this computation are called pipelines, and are critical parts of the reliability of radial velocity measurements. It is beyond the scope of this work to enter in the details of pipelines, we simply mention a few works for further reading. For example, Zechmeister et al. (2014) proposes an optimal extraction method to obtain a one-dimensional spectrum from a two-dimensional CCD array, as encountered in echelle-grating spectrographs, see reference therein for other techniques. Bauer et al. (2015) outlines a wavelength calibration technique for Fabry-Perot etalons. The HARPS pipeline is described in Anglada-Escudé & Butler (2012).

Table 1.1: Summary of known signal sources. Numbers for instrumental effects and barycentric correction are taken from Lovis & Fischer (2010). Others are from Cunha et al. (2013) (blending stars), Cegla et al. (2012) (Gravitational redshift), Reiners (2009) (flares), Dumusque et al. (2011b) (p-modes, granulation), Dumusque et al. (2011a) and Dumusque et al. (2014) (spots, faculae, plage), Baliunas et al. (1998); Lovis et al. (2011a) for activity cycles.

Effect	**Amplitude and timescale**	**Correction (if available, or modelling)**
	Signal	
Planet motion generate a reflex motion of the star	1 cm/s to hundreds of m.s^{-1} $<$ 1 day to hundreds of years	$v(t) = V + \sum_{i=1}^{n_p} K_i(\cos(\omega_i + \nu_i(t)) + e_i \cos\omega_i)$ (see equation 1.2)
	Uncertainties	
Barycentric correction: wavelength an observer at the center of ICRS would receive	0.01 (BEPOP) - 0.5 m.s^{-1} 1 year and harmonics	Correction from planetary ephemeris
Star blends: A background star contaminates the spectrum	Depends on the difference in magnitude and spectral type. On average $\leqslant$ 10 cm/s but can reach 1 m.s^{-1} for $\Delta m \leqslant 8$ $\approx$ Planet orbital period	Activity diagnostics (Santerne et al. 2015)
Lunar & Telluric Contamination: the spectrum might be contaminated by lunar or telluric lines as well as faint companions that add up with the stellar spectrum and deform the CCF	Telluric lines: a year Moon: bad sky transparency, bad weather	Subtracting CCF of the channel pointing at the sky from the CCF of the observations (Bonomo et al. 2010)
Photon noise	$\frac{c}{Q\sqrt{F\Delta t}}$ with Q quality factor depending on the stellar type, F photon flux, Δt integration time	The number of photons received is counted in real time to check achievement of required SNR.
Variation of the air refraction index: change in temperatures and pressure induce spectral line shifts	Temperature change: 100 m.s^{-1}/K Pressure change 100 m.s^{-1}/mbar observation night	Instruments in vacuum chambers, put in thermal and pressure regulated room.
Thermal and mechanical effects: variations of temperatures induce changes of instrument properties	up to hundreds of m.s^{-1}	

Slit illumination variation: variations of the photocentre change the response of the instrument	Photocenter shift induce variations of 30 km/s/slit width for a resolution of 100,000	Fibers of non circular sections that scramble the incident light.
Calibration imprecision: associating a wavelength to a pixel is done by calibration several times per night but not perfectly	Variations of 1-2 $m.s^{-1}$, down to 30 cm/s, in the long term (several months) depending on techniques	Improving stability of calibration instrument (easier said than done)
Detector faults: the variation of properties pixel per pixel and read-out noise create systematics	Increased relative noise at low SNR Possible periodic effect if a deficient detector is periodically reached	Improving CCD detector regularity and stability, check for systematic yearly or half-yearly effect.
Instruments upgrades: change of components might induce improved precision but also offsets that are hard to predict	up to several $m.s^{-1}$, but can be corrected	approximation from stable star monitoring.
Granulation: convective motion in the star's atmosphere cause local variation of the atmosphere radial velocity	$\approx$ 0.1 - 4 $m.s^{-1}$ τ_1 = 20 min, $\tau_2 \approx$ 40 min, $\tau_3 \approx$ 1 day	Stationary stochastic process of spectral density $P(\nu) = \sum_{k=1}^{3} \frac{\sigma_k^2 \tau_k}{1+(2\pi\nu\tau_k)^2}$ (discussed model). Observation strategy: taking several observations per night
p-modes: The propagation of acoustic waves in the star mantel cause it to oscillate	$\approx$ 0.1 - 4 $m.s^{-1}$ 5 - 15 minutes	$\gtrsim$ 10 min observations to average out the effect.
Relativistic terms: gravitational field $\Phi_\star$ and star velocity $v_\star$ effects on the received wavelength	$\approx$ 0.06 cm/s observation timespan	See eq 1.9
Spots: darker region of stellar surface, inhibit convection, introduce asymmetry of the light flux between the approaching and receding parts of the stellar disk	$\approx$ 0.1 - 10 $m.s^{-1}$ Stellar rotation period and harmonics	Understanding via numerical models and observations
Faculae/Plages: brighter region of stellar surface, inhibit convection, introduce asymmetry of the light flux between the approaching and receding parts of the stellar disk	$\approx$ 0.1 - 10 $m.s^{-1}$ Stellar rotation period and harmonics	Partial correction via FWHM, bisector span, $\log R'_{HK}$

Activity cycles: number of starspots, plages and there location vary stochastically during the activity cycle	Up to 25 m.s^{-1} $\approx$ 7 - 30 years	Modelled as stochastic processes (so far Gaussian processes)
Flares: Sudden ejection of coronal mass	Several 10 m.s^{-1} $\approx$ 20 mins flashes	Seen in e.g. H_α line

1.3 Inference

I learned from Popper what for me is the essence of scientific investigation — how to be speculative and imaginative in the creation of hypotheses, and then to challenge them with the utmost rigor, both by utilizing all existing knowledge and by mounting the most searching experimental attacks. In fact, I learned from him even to rejoice in the refutation of a cherished hypothesis, because that, too, is a scientific achievement and because much has been learned by the refutation.

John Eccles

In the work already mentioned which I was working on at the time, I referred to passages from this book more or less at random, not in order to refute them — what business is it of mine to refute! — but, as befits a positive mind, to replace the improbable with the more probable and in some circumstances to replace one error with another.

Friedrich Nietzsche, On the Genealogy of Morality

1.3.1 Context

As seen last section, a considerable amount of work has been done to build outstanding instruments as well as to identify, model and possibly correct the components of the spectroscopic signals. Furthermore, observation nights are expensive and time consuming. It is therefore important to exploit as much as possible the data available and to try to optimize observation strategies. Before presenting actual data analysis techniques, that rely on probability theory, let us take a step back and think of what we are trying to do.

In this thesis, the main concern is to gain information from observations, which is often referred to as inference. The difficulties of such a purpose are to a certain extent independent of the type of observations and of the information to be extracted. In this chapter we present some general considerations about inference that apply to the study of extrasolar planets.

We will take as a practical definition of inference the fact of characterizing the properties of an object about which some information is missing. On the contrary, in the context of deductive reasoning, one can specify hypotheses from which propositions are derived. When dealing with experimental data, we only have partial and imprecise information. For instance,

a noise model can be seen as partial information: we obviously do not know the exact noise value at the moment of the measurement, but we know that it is going to behave roughly like a Gaussian variable that has a certain variance σ. Furthermore, this model is imprecise, since we know that the noise does not follow exactly this Gaussian law. Unlike a mathematician who can choose the hypotheses he starts from, the experimenter cannot assume a certain law for the noise and be perfectly right simultaneously.

Inferring properties of the real world has been a subject of intense debate among philosophers that might now seem outdated thanks to the powerful statistical tools that have arisen in the 20^{th} century. However, it is not always clear whether the statistic we compute answers the question that is being asked. A blind use of statistics can lead to spurious claims, as has been the case in medicine (Ioannidis 2005) because of an unreflected and incorrect use of the so-called p-values.

We will not dwell on the philosophical debate, but state what is the epistemological choice made here — and underlying in most scientific research, — that is the Popperian response to the problem of induction. Since it is not possible to build general cases from data, the position adopted is to make clear hypotheses, then derive consequences on what should be observed in the framework of deductive reasoning, and finally compare to experiments. If the results disagree, at least one of the hypotheses should be rejected. If not, it will only be said that data corroborates a hypothesis, keeping in mind that no refutation or corroboration is definitive.

Trying to find hypotheses that agree with a data set can be compared with the resolution of an equation $f(x) = y$, y being the data, f the model and x a parameter of the model. This one can be seen as a vector equation, which encompasses agreement with data and with other knowledge we might have. The set of solutions of the equation is similar to a set of non-rejected hypotheses, those that reproduce with "good" accuracy the data and other prior knowledge.

If we want to make a strong case for a particular hypothesis, this might be seen as proving uniqueness of the solution. Also, when the space of solution/hypotheses is too hard to explore, only an approximation of a subset of solutions will be available. Finally, it might turn out that the space of possible solutions/hypotheses was initially too restrictive and the solution lies outside, which is similar to situations that can be encountered in equation solving. With this viewpoint, trying to select a model among those compatible with the data amounts to introduce exterior sources of information, additional "equations" to constrain explanations, for instance preference for simple models, as discussed in section 1.3.2. In the context of exoplanet detection, to evaluate the plausibility of a planetary systems whose parameters are estimated via a statistical analysis, one might consider other source of information such as the type of the star or the stability of the estimated system.

Furthermore, let us stress that the "equations" putting constraints on our hypotheses depend on the sequence of event that leads to a scientific claim. A hypothesis formulated after seeing the data should not be treated as if it had been formulated beforehand. The more hypotheses we explore, the more likely it is to find one that reproduces the data. If the hypothesis found to be most consistent is assessed as if there was no prior search, its likelihood will be greatly overestimated. Funny instances of this problems are available on http://tylervigen.com/discover where one can discover correlations by exploring large data sets. Figure 1.14 we show an example of a correlation between the number of doctorate

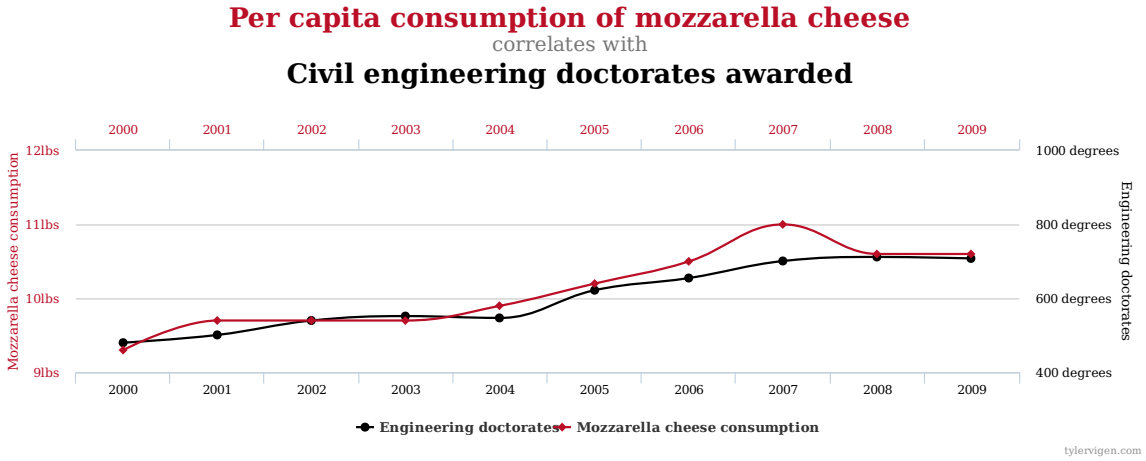

Figure 1.14: From http://www.tylervigen.com/spurious-correlations.

in civil engineering and the consumption of mozarella, both in the United states. The 95.86% correlation coefficient is obviously due to the fact that a correlation is searched in a wide database of curves. In RV data analysis, one can always reproduce the data more and more closely as the number of planets in the model increases. To claim a detection, it must be ensured that a good fit with a planetary model does not only arise from the fact that many orbital configurations were considered.

A major problem remains: what criteria should be used to express a disagreement between data and a model? The tools available rely heavily on probabilistic tools, presented next section.

1.3.2 Uncertainty and probability

To put in practice the Popper method in a quantified way we need a mathematical concept that embodies uncertainty. There are two traditions to model uncertainties: starting from the frequency of an event and starting from requirements on inference rules compatible with common sense. The following sections outline these two approaches.

The difference can be approached from a historiographical perspective. Our goal is not to try to unveil what authors have meant by frequentist and Bayesian and to comment these choices. In the following discussion we define what we intend by frequentist and Bayesian based on what seems to be their common use and then discuss these notions so defined. This discussion should be seen as a starting point from existing literature to discuss the more general question: what is a useful measure of uncertainty? In simpler terms: how sure am I?

The point of view we ultimately defend is that the interpretation of probability (frequency, subjective measure of certainty...) and what are the relevant numbers to make inference (shall we use Bayes formula?) are two different matters. Secondly, that using Bayes formula has the advantage of stating more clearly assumptions but that the computational burden makes the computation slow to re-do with different assumptions. Methods not relying on the Bayes formula are blind to the fact that the data might have been generated from a population, but since these are faster, they are more practical to see the dependence of the result on a model change.

Frequency of an experimental result

Random variables In the first case, one considers the outcome of an experiment as a random variable. That is a real valued function on a space Ω to $\mathbb{R}^N$, where N is assimilable to the number of observations,

$$\begin{aligned} \boldsymbol{Y} : \Omega &\to \mathbb{R}^N \\ \omega &\to \boldsymbol{y}. \end{aligned}$$

For $\boldsymbol{Y}$ to be a random variable, one must add that Ω is a measurable space of total measure equal to one, and $\boldsymbol{Y}$ is a measurable application (see for instance Rényi 2007). Let us first assume that $\boldsymbol{Y}$ takes a finite number of values $\boldsymbol{Y}_1 ....\boldsymbol{Y}_n$. A classical interpretation of saying that "$\boldsymbol{Y}$ follows the probability distribution $p_1, ..., p_n$ is that if we were to do the exact same experiment m times, the outcomes $\boldsymbol{y}_1, ...\boldsymbol{y}_n$ would be such that the number of outcomes $\#\{\boldsymbol{y} = \boldsymbol{Y}_i\}/m$ tends to p_i as m tends to infinity. So when the number of outcomes is finite, the outcome of the experiment is $\boldsymbol{Y}_i$ is considered as an event that happens on average every $1/p_i$ times. This interpretation is termed frequentist, and was promoted for instance by Ronald A. Fisher.

Let us note that so far we do not need a notion of randomness. We can think of a machine that selects the outcome of the experiment so that if it was to be done m times, the frequency of each $\boldsymbol{Y}_i$ would tend to p_i. The "machine" can proceed in a deterministic way — just as a computer generates pseudo-random sequences — as long as the average properties are satisfied and the experimenter does not know the selection process. The random variable is then a convenient way to express a partial knowledge of the experiment outcome: we know how often a situation will happen on average, but not exactly when.

In general, $\boldsymbol{Y}$ will depend on assumed physical properties gathered in a vector of parameters $\boldsymbol{\theta}$ through a function $\boldsymbol{f}(\boldsymbol{\theta})$ and the uncertain part comes from a so-called additive noise, random variable $\boldsymbol{\epsilon}$ so that $\boldsymbol{Y} = \boldsymbol{f}(\boldsymbol{\theta}) + \boldsymbol{\epsilon}$. The goal is to put constraints on the value of $\boldsymbol{\theta}$, which are the physical parameters of interest, based on observations $\boldsymbol{y}$. To do so, one can consider that $\boldsymbol{\theta}, \boldsymbol{f}$ and the distribution followed by $\boldsymbol{\epsilon}$, denoted by $f_{\boldsymbol{\epsilon}}$, are fixed, therefore the distribution of $\boldsymbol{Y}$ is known. Consequently, it is possible to compute the probability of obtaining the outcome $\boldsymbol{y}$ if it was generated via $\boldsymbol{Y}$. The lower is this probability, the less likely is the model. It means that $\boldsymbol{\theta}$, $\boldsymbol{f}$ or the distribution of $\boldsymbol{\epsilon}$ is unlikely to be correct. If no parameter $\boldsymbol{\theta}$ is compatible with the observations, then probably $\boldsymbol{f}$ or $f_{\boldsymbol{\epsilon}}$ are incorrect. In this method, we are considering what are the random variables that could have generated $\boldsymbol{y}$ not too rarely.

Non parametric interpretation The reference to parametrized models $\boldsymbol{f}(\boldsymbol{\theta})$ tends to mask an interesting interpretation of inference from random variables. What do we expect to unveil with such a procedure? We are implicitly assuming that the data have been generated according to a distribution $f_{\boldsymbol{Y}_t}$ which plays the role of the unknown of our equation. We are looking for probability distributions $f_{\boldsymbol{Y}}$ that could have generated the data and hope $f_{\boldsymbol{Y}_t}$ is among them. Since we cannot search among all the possible distributions, the search is restricted to a certain space, which constitutes another assumption.

Considering that the data have been generated according to a probability distribution means, in the frequentist interpretation, that the data are a particular realization of a more

general class of events. It is however not always clear if it is possible express uncertainty this way. For instance, should we allow ourselves to consider "the chance of meeting my colleague Randy in the elevator this Monday morning" as the frequency of an event among a set of possibilities? To do so, we need to see it as a particular occurrence of a wider class that could be repeatable (meeting a colleague a Monday morning) or thinking of our world as a particular one among others — which is not so intuitive.

Uncertainty on the model If we know the data we see or we expect to see is generated via a certain probability distribution, if we delimit the space of probability distributions we look for, then the uncertainty is quantified by the frequency of an event. However, we have an uncertainty on whether the right framework was chosen to conceptualize the experiment, and this uncertainty cannot be expressed in terms of frequency of an event[2]. This means in practice that computed probabilities are not an absolute, perfect measure of uncertainty, and we should not jump too fast to conclusions even from compelling numbers. There is also a very fundamental uncertainty on whether it is correct to consider that the data were generated by a random variable.

Section 1.3.1, it was said that it is different to make inferences before and after the data has been taken. This discussion allows to make this claim precise. A scientist can conceive an experiment whose outcomes are delimited and which can be repeated. In this case, the uncertainty can be effectively seen as partial knowledge on the rate of occurrence of an outcome. The computed probabilities will be helpful to make inferences. On the contrary, when we are given data, we don't know how it was produced. It is not clear what is the class of events it belongs to and what repetitions of the experiment would mean. The uncertainty is great, but difficult to quantify. Another way of looking at it is that as we acquire knowledge on the data, we are in a different state of information, we add other constraints to possible explanations. If we ignore that information, the risk of spurious claims might increase.

What to think about exoplanets then? In the context of astronomy, the interpretation of probability as a frequency does not seem in general to raise problems. We could not estimate better the distribution of a certain property of stars than by knowing how many stars have this property divided by the number of stars in the observable universe. Also, as more measurements are taken, the number of times the error will lie in a certain interval tends in general to a Gaussian probability distribution.

As said above, in the context of inference, assumptions have to be made to come back to the framework of deductive reasoning. As long as the subtended population of events is made clear and assumptions are stated, the analysis has a meaning. What remains unclear is whether the computed probability helps us making correct claims. This problem has been tackled in different ways, in particular by using another interpretation of probabilities presented in the next section. We discuss the differences with the frequentist approach paragraph "Bayesian vs Frequentist".

[2]Unless it is itself seen as a frequency in a more general framework, itself embodied in an even more general one. However to avoid infinite recursion an assumption has to be made at some point. The use of the Bayes formula appears when embedding models within each other.

Uncertain claims

The formalism of probability theory can be justified from another point of view, which traces back to Laplace and also Keynes (1921). This one consists in extending the laws of Aristotelian logic, formalised by Boolean algebra, to situations where we do not have certain knowledge. To be clear, we are considering propositions that are either true or false, but on which we have incomplete information, and we evaluate our degree of belief that a statement is true. This problem is different from evaluating the degree of truth of an imprecise proposition. For instance, we want to know whether "a planet of mass superior to Jupiter orbits star HDx" is correct or not, and we evaluate our belief this is the case. We are not concerned with evaluating in which extent "planet HDxb orbiting star HDx is big" is true.

The major development of this approach came from Cox's theorem (Cox 1946, 1961), later popularised by Jaynes & Bretthorst (2003). Cox's starting point is the question: what do we expect from a quantity that evaluates the plausibility of a proposition A holding another proposition B as correct? For our purpose, Cox (1946) argues that it seems reasonable to impose several requirements to a quantified measure of plausibility. For the mathematical definition of these requirements, we adopt (mostly) the conventions and presentation order of Halpern (1999b).

1. The degree of plausibility of a proposition holding another proposition as true should be represented by a real number. Implicitly, this means that we are considering a function P and a set W such that for any two subsets U, V of W with $U \neq \emptyset$, $P(V|U)$ is a real number.

2. The plausibility of the negation of a any proposition V, denoted by $\bar{V}$, given $U \neq \emptyset$, depends only on the plausibility of V given U. In other words, there exists a function $S : \mathbb{R} \to \mathbb{R}$ such that $P(\bar{V}|U) = S(P(V|U))$.

3. The plausibility of a conjunction of any events V and V', $V \cap V'$ such that $U \neq \emptyset$ depends only on the plausibility of $V|U$ and of $V'|V \cap U$. In other words, there exists a function $F : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ that verifies $P(V \cap V'|U) = F(P(V'|V \cap U), P(V|U))$.

We are considering propositions are sets belonging to a wider set of all acceptable propositions W, as is classic in Boolean algebra or probability theory. The propositions are identified with the "set of situations" they cover. The condition $U \neq \emptyset$ means simply that the proposition U we assume is not unrealisable. It would not make much sense to search for the plausibility of V knowing U and $\bar{U}$, or to be more concrete, the plausibility that we are going to arrive on time given the train is late and the train is not late. The assumption that the plausibility is a real number does not mean, at least at this point, that plausibilities are ordered on an absolute scale. We could decide that a plausibility equal to two indicates a higher degree of confidence than a plausibility equal to three.

The second hypothesis states that if we know the plausibility of an event, then plausibility of its negation is fully determined. This means in particular that the plausibility of $\bar{A}|C$ does not depend on A nor C but only on $P(A|C)$. If the plausibility of arriving on time knowing the train is late, and the plausibility of receiving a call from a cousin knowing she is in town have the same value, then the plausibility of their respective negations will be the same. The third proposition can be understood as compatible with usual logic. To decide whether V and V' are true knowing U, we can start by proving that $V|U$ is true, and with that knowledge,

proving that V' is true or conversely. The fact that these two operations commute yields that F should be associative, $F(x, F(y, z)) = F(F(x, y), z)$ which is key in obtaining Cox's theorem from these assumptions, namely: P has to be isomorphic to a conditional probability measure on W, i.e. there exists a bijective continuous function $g : \mathbb{R} \to \mathbb{R}$ such that $g \circ P$ verifies

$$g \circ P(W|W) = 1 \tag{1.12}$$

$$g \circ P(V' \cup V|U) = g \circ P(V'|U) + g \circ P(V|U) \text{ for disjoints } V, V' \tag{1.13}$$

$$g \circ P(V|U \cap U') \times g \circ P(U|U') = g \circ P(V \cap U|U') \tag{1.14}$$

as soon as $U', U \neq 0$. Unfortunately, as it is Cox's proofs (Cox 1946, 1961) are incorrect. With only these assumpions, Halpern (1999a) exhibits a counterexample when the set of possible propositions W is finite. However, Paris (1994) showed that with additional hypotheses, the proof can be made rigorous. These are, as stated in Halpern (1999b),

1. The range of P is $[0, 1]$ (or any closed interval).
2. $P(\emptyset|U) = 0$ and $P(U|U) = 1$ for $U \neq \emptyset$.
3. S is decreasing (see condition 2 above).
4. F is strictly increasing on both coordinates in $(0, 1]$ and is continuous (see condition 3 above).
5. For all $0 \leqslant \alpha, \beta, \gamma \leqslant 1$ and $\epsilon > 0$, there exists sets $U_4 \subset U_3 \subset U_2 \subset U_1$ such that $U_3 \neq 0$ and each $|P(U_4|U_3) - \alpha|$, $|P(U_3|U_2) - \beta|$ and $|P(U_2|U_1) - \gamma|$ is less than ϵ.

The first four assumptions do not seem so stringent. The assumptions 1 and 2 jointly mean that the range of P cannot be infinite, there is a closed scale of certainty, starting from 0 for an impossible event and reaching one for a certain event. The assumptions 3 and 4 state respectively that the more plausible A, the less plausible its negation $\bar{A}$ and the more plausible $A|B \cap C$ or $B|C$, the more plausible $A \cap B|C$. The assumption number five is not very appealing and hard to interpret as such, but one of its consequences is easier to grasp: W must be infinite. Indeed, 5 implies that for all $U \neq \emptyset$, the values taken by $P(V|U)$ should be dense in $[0, 1]$.

These results have cast some doubt on how compelling Cox's arguments were regarding the use of probability as privileged system of inference. To cut short to an abstract discussion on the rules of inference (see for instance Horn 2003), let us point out that in astronomy we are interested in physical quantities that vary continuously, therefore the infiniteness of W is verified. We shall neither consider theories in which the degree of belief is represented by two numbers (Shafer 1976; Dubois & Prade 2012), which are too far from common practice. Let us conclude from this discussion that the formalism of probability theory seems attractive even when considering rules of inference from an abstract point of view.

Bayesian vs Frequentist

Bayes formula In the following, we will encounter data analysis techniques that are stamped "Bayesian" or "frequentist" and a choice will have to be made. To define the differences

between the two approaches, let us first write Bayes theorem. This one gives the probability of a hypothesis H knowing data y, that is respectively in the discrete and continuous cases,

$$\underbrace{\Pr\{H|\boldsymbol{y}\}}_{\text{Posterior (probability)}} = \frac{\overbrace{\Pr\{\boldsymbol{y}|H\}}^{\text{Likelihood}} \quad \overbrace{\Pr\{H\}}^{\text{Prior (probability)}}}{\underbrace{\Pr\{\boldsymbol{y}\}}_{\text{Normalizing constant}}} \quad \text{and} \tag{1.15}$$

$$\underbrace{p(\boldsymbol{\theta}|\boldsymbol{y})}_{\text{Posterior (distribution)}} = \frac{\overbrace{p(\boldsymbol{y}|\boldsymbol{\theta})}^{\text{Likelihood}} \quad \overbrace{p(\boldsymbol{\theta})}^{\text{Prior (distribution)}}}{\underbrace{p(\boldsymbol{y})}_{\text{Normalizing constant}}}. \tag{1.16}$$

the braces indicate the name usually given to each term of the formula. We use the notation p for all distributions, which is common use, but is an abuse of notation since these are different functions. In general, one refers to "the prior" or "the posterior" to talk about prior and posterior distributions or probability. If we have n hypotheses $H_0, ...H_{n-1}$ that are disjoint and such that $\sum \Pr\{H_k\} = 1$, or the continuous parameters $\boldsymbol{\theta}$ describe a set Θ such that $\int_\Theta p(\boldsymbol{\theta})\,\mathrm{d}\boldsymbol{\theta} = 1$, the formulae become

$$\Pr\{H_i|\boldsymbol{y}\} = \frac{\Pr\{\boldsymbol{y}|H_i\}\Pr\{H_i\}}{\sum_{k=0}^{n-1} \Pr\{\boldsymbol{y}|H_k\}\Pr\{H_k\}} \tag{1.17}$$

$$p(\boldsymbol{\theta}|\boldsymbol{y}) = \frac{p(\boldsymbol{y}|\boldsymbol{\theta})p(\boldsymbol{\theta})}{\int_\Theta p(\boldsymbol{y}|\boldsymbol{\theta})p(\boldsymbol{\theta})d\boldsymbol{\theta}}. \tag{1.18}$$

It is not always clear what is meant by frequentist and Bayesian. When these two terms are opposed, it is — as far as we know — regarding one of the two following issues: should we interpret probabilities as limit frequencies of events or measures of uncertainties in general? Secondly, should we use the Bayes formula even if the prior is unknown? The methods that are called frequentist (p-values, maximum likelihood,...) never make use of the prior. We will thus call frequentist any analysis that does not use formula 1.17 nor 1.18, and restrict itself to using $\Pr\{\boldsymbol{y}|H\}$ or $p(\boldsymbol{y}|\boldsymbol{\theta})$.

An example Let us leave aside the problem of the interpretation of probabilities for a moment and focus on the use of the Bayes formula. To illustrate the differences between using the Bayes formula or not, we consider a classical example of the Bayesian literature. Let us suppose that a patient is tested for an illness, and denote by I and $\bar{I}$ the events "the patient is ill" and "the patient is not ill", respectively. We denote by T^+ a positive test (which says the patient is ill) and by T^- a negative one. This test is not perfect, and the uncertainty on its outcome is such that $\Pr\{T^+|I\} = 0.95$. When a patient is ill, the test will be positive in 95 cases out of 100. We further assume that the probability of the test to be positive when the patient is not ill is $\Pr\{T^-|I\} = 0.04$. Note that there is no reason why we should have $\Pr\{T^+|I\} + \Pr\{T^-|I\} = 1$. For instance the device for the test could be broken and always output T^+ with probability 0.95 regardless of the state of the patient.

From the physician and patient point of views, what matters is the probability of being ill if the test is positive or negative. In the first case, formula 1.17 gives

$$\Pr\{I|T^+\} = \frac{\Pr\{T^+|I\}\Pr\{I\}}{\Pr\{T^+|I\}\Pr\{I\} + \Pr\{T^+|\bar{I}\}\Pr\{\bar{I}\}} = \frac{1}{1 + \frac{\Pr\{T^+|\bar{I}\}\Pr\{\bar{I}\}}{\Pr\{T^+|I\}\Pr\{I\}}} \tag{1.19}$$

To compute this probability, we need another information, which is the probability of being ill $\Pr\{\bar{I}\}$. If the disease is relatively rare, for instance $\Pr\{\bar{I}\} = 0.01$, we obtain that $\Pr\{I|T^+\} \approx 0.19$. Even though the test seemed reliable because of $\Pr\{T^+|I\} = 0.95$, due to the rarity of the disease when the test is positive there is only one chance out of five that the patient is actually ill.

Unkown prior This example shows that the underlying population ($\Pr\{I\}$) might induce a large difference between $\Pr\{T^+|I\}$ and $\Pr\{I|T^+\}$. Now the problem is that in general, the probability $\Pr\{I\}$ is unknown. A common practice in case of uncertainty would be to assign $\Pr\{I\} = 1/2$. In that case,

$$\Pr\{I|T^+\} = \frac{\Pr\{T^+|I\}0.5}{\Pr\{T^+|I\}0.5 + \Pr\{T^+|\bar{I}\}0.5} = \frac{1}{1 + \frac{\Pr\{T^+|\bar{I}\}}{\Pr\{T^+|I\}}} \approx 0.96. \tag{1.20}$$

Then it seems likely again that the patient is ill if the test is positive. Note that

$$\frac{\Pr\{T^+|\bar{I}\}}{\Pr\{T^+|I\}} \approx 1/23 \tag{1.21}$$

which also seems to favour to a convincing extent the hypothesis that the patient is ill.

Since the result depends on the prior, it might be tempting to just reject its use, and stick to ratios of the type $\Pr\{T^+|\bar{I}\}/\Pr\{T^+|I\}$. However, we just saw that this one is a particular case of a Bayesian ratio, when the priors of alternative hypotheses are equal to each other. More importantly, if the disease was indeed rare and concerning 1% of the population, starting a treatment would be a wrong choice roughly four times out of five on average. We believe that a good interpretation of the test example is that the test is simply not precise enough to make a reliable inference. Secondly, that ideally, inference should not be based on a single prior, but one should evaluate the sensitivity of the result to a change of prior information. This indeed reflects better our state of knowledge when we simply do not have a clear idea of the prior distribution. Just like the fact that the actual probability of the patient is ill or not does not depend on the ignorance of the doctor, the fact that a planet orbits a star or not does not depend on the prior choice. A way to tend to robustness is to impose $\Pr\{T^+|\bar{I}\}/\Pr\{T^+|I\} \approx 1/N$ where N is "very large". It means that even if $\Pr\{I\}$ is small the ratio $(\Pr\{T^+|\bar{I}\}\Pr\{\bar{I}\})/(\Pr\{T^+|I\}\Pr\{I\})$ is still in favour of I. We will encounter similar ratios for the selection of model, more precisely equation (3.18), where $N = 150$ is a common choice. The only way to obtain a better ratio is to take more measurements. From a more theoretical perspective, Cox's theorem tells how to update a state of belief when data is available, but does not give guarantee on the direct usability of the posterior as an absolute measure of uncertainty. In the notations of the previous sections, it states what to do when a plausibility measure P is chosen, but not how to choose it.

Probability interpretation Regarding the question of the probability interpretation, let us note that in this example, we used the Bayes formula and still thought of the probability as rate of occurrence of a certain event. This situation is similar to information theory (Shannon 1948), where the frequency of the letters of the English language is known, and the exact posterior probability of the received signal is used to design the communication devices. The posterior probability, defined as the probability that "b was the sent message knowing a is the received message" is exactly the rate of occurrence of this event. One can loosely think of the orbit parameters of a given system as a message that we are trying to decode with measurements, and that has a certain distribution we wish to determine.

Now what if there are reliable records that the occurrence of the disease is actually 1%? The physician might want to use an additional information, which is that if the patient was sent to be tested, it is probably because other symptoms lead to a suspicion of illness I. If one can think of the probability of a randomly chosen person in a population as a frequency, a given person is ill or is not. Expressing a prior belief cannot be easily expressed as a frequency, though the way priors are set might be inspired from frequencies, for instance "80% of the patients the physicians saw that displayed a certain symptom S had I, this patient has S, therefore one can put a 0.8 probability prior on its illness".

Furthermore, if one was to think of a way to evaluate a method of diagnostic, the notion of frequency comes naturally. For instance, a hospital director knows that each physician will use several elements to take decisions such as its intuition, which cannot be expressed in terms of frequencies. However, the director can assess whether ultimately, the decision process as it is yields a satisfactory rate of healing. This is also true with a data analysis method for the discovery of exoplanets. If a data analyst was to use a certain method with priors set in a heuristic way, one could try that method on simulated data and evaluate how reliable it is. That is exactly the principle of blind tests such as the radial velocity challenge Dumusque (2016). To evaluate the performance of the method one counts how many planets were correctly detected. The uncertainty on a claim are assimilable to stating "how frequently" the method employed yielded a correct result. This applies to observation strategies, the fact of stopping observations if a star displays no clear hints of signals after, say, 50 measurements biases the population estimate.

We are not able to demonstrate that every reasonable measure of uncertainty comes down to a measure of a frequency, however be it for the justification of a prior probability or the evaluation of a method, it seems difficult to avoid the notion.s Intuitively, it seems like one needs to consider a particular event as belonging to a wider class to consider that it might not always happen. An analogy with Heisenberg's uncertainty principle could be helpful: when an event is completely specified (say, the probability of meeting my colleague Randy this Monday morning at 8:01:36 etc.), the uncertainty on its frequency is infinite, since it cannot be reproduced. On the other hand, when the class of events is very wide (the probability of meeting someone you know without appointment in the next ten years) one can have a good resolution on their rates of occurrence, but a poor one on the definition of these events.

Computations For hypothesis testing, the ideal solution would be to use Bayes formula with all possible priors and all possible models. In that case, the tested hypotheses are not any more the H_k as defined (1.17), but the family of H_k, extended to other models and their prior probabilities $\Pr\{H_k\}$. This method is fully compatible with the Popperian view

outlined section 1.3.1. Obviously such a test is not possible because it would require an infinite power of computation. Even for a single prior, when the parameters are continuous the posterior (1.18) might be very lengthy to estimate (up to several days as we shall see later), so changing the assumption on the likelihood might slow, yet very important to assess the robustness of a result, since the likelihood contains the physical model of the process. In the case of exoplanets, the models for stellar noise are known to be heuristic. Robustness to a model change seems desirable to guarantee the planetary nature of a signal.

Conclusion From this discussion, we take the following points: the true prior does have an impact on the probability that a hypothesis is true given the data. It is sound to test several priors, the problem being that it might be infeasible computationally. In some cases, methods that do not use priors could be useful because of their speed and they allow to quickly change assumptions on the likelihood, which contains the physical model of the observations. In any case, testing the whole process of inference seems sound.

Occam's razor

> *This model fits better the data therefore it is true, even though it could fit anything anyway.*
>
> *Nobody*

The search for acceptable hypotheses was compared to the resolution of an equation. Instead of searching variables $\boldsymbol{x}$ that verify $\boldsymbol{f}(\boldsymbol{x}) = \boldsymbol{y}$ we search $\boldsymbol{x}$ such that the agreement between the model and the data y is small, according to some measure of distance, $d(\boldsymbol{f}(\boldsymbol{x}), \boldsymbol{y}) \approx 0$. More precisely, the equality is transformed to degrees of preferences: the better the data are reproduced the better seems the explanation. We also use knowledge that comes from past research, which is comparable to introduce other equations or degrees of preference. For instance, the values of the orbital parameters estimated in multiplanetary systems are checked to give stable systems (Hébrard et al. 2016).

Now among the hypotheses that are compatible with data, a common practise is to select the simplest in some sense. This principle is often referred to as "Occam's razor", that cuts any bits of non essential information. This preference for simplicity is embedded more or less explicitly in most model selection methods: Solomonoff's theory of inductive inference (Solomonoff 1964a,b), Akaike's information criterion (AIC, Akaike 1974), Schwarz' Bayesian Information Criterion (BIC, Schwarz 1978), Rissanen's Minimum Length Description (MDL, Rissanen 1978), Bayes factor (Kass & Raftery 1995)... Our purpose is to make correct claims, why should we bother about the model complexity? We use the analogy with equation resolution to discuss the use of the principle.

Let us consider different models $\boldsymbol{f}_k$ with $\boldsymbol{f}_k : \mathbb{R}^k \to \mathbb{R}^n, 0 < k \leqslant n$. In other words, the functions $\boldsymbol{f}_k$ have an increasing number of parameters. We further suppose that $\boldsymbol{f}_{k+1} = \boldsymbol{f}_k + \boldsymbol{g}_k$ where $\boldsymbol{g}_k$ is a function of one variable which vanishes at some discrete set of values, and that all $\boldsymbol{f}_k$ are continuous. Now suppose the "data" is $\boldsymbol{y} = \boldsymbol{f}_{k_0}(\theta_{k_0})$ for some k_0, which can be seen as a noiseless measurement. According to our definition of the $\boldsymbol{f}_k$, for any $k \geqslant k_0$, there will exist a $\boldsymbol{\theta}_k$ such that $\boldsymbol{y} = \boldsymbol{f}_k(\boldsymbol{\theta}_k)$. However, for any $k < k_0$, there exists a θ_k such that $\boldsymbol{y} = \boldsymbol{f}_k(\boldsymbol{\theta}_k)$ only if $\boldsymbol{\theta}_{k_0}$ is such that $\boldsymbol{g}_{k_0}(\boldsymbol{\theta}_{k_0}) = 0$, which happens on a discrete set of

points among a continuous space. Now if $\boldsymbol{\theta}_{k_0}$ is selected via any continuous random law, $\boldsymbol{g}_{k_0}(\boldsymbol{\theta}_{k_0}) = 0$ almost surely never happens. In other words: the probability of selecting a spuriously simple model is zero.

This example is best interpreted in geometrical terms. One can think of the models not as functions but as manifolds, defined as the set of values taken by $\boldsymbol{f}_k(\boldsymbol{\theta}_k)$ as $\boldsymbol{\theta}_k$ describes the domain of $\boldsymbol{f}_k$. Provided the functions are well behaved, the dimension of the manifold grows with the number of parameters. If $\boldsymbol{y}_{k_0} = \boldsymbol{f}_{k_0}(\boldsymbol{\theta}_{k_0})$ is selected randomly on a manifold of dimension k_0, $\mathcal{M}_{k_0}$, its number of parameters will be underestimated only if it falls on the intersection of $\mathcal{M}_{k_0}$ and a manifold of lower dimensions, which has a null measure, thus the event has a probability zero.

In the case of nested models (the next model contains the previous one as a particular case), the choice of the simplest model guarantees that the parameters selected at least are in the model. Furthermore, provided more complex models have non null probability, if a more complex model was the true one and ours was spuriously simple it means that we would be in a situation that has zero probability. Let us note that the reasoning holds for non nested models, as long as they have increasing dimensions and null measure intersections. In the real case, the measurements are noisy so that this statement is not rigorous anymore. The probability of selecting a wrong model increases with the noise amplitude. What we showed is the asymptotic behaviour as the noise tends to zero.

Let us note that for the purpose of this discussion we assumed that models are manifolds, and that the data are not certainly in one of the manifolds. With these assumptions, Occam's razor is not a principle, but a consequence.

Implications for exoplanets

What should be concluded from the previous analysis for the study of exoplanets? We highlight these points

- The situation of exoplanet detection is similar to information theory. The emitter has a language whose distribution is the actual rate of occurrence of planets with certain characteristics. Even though the receiver selects himself the star, his measurements are assimilable to receiving a noisy signal from a subclass of possible symbols (from a particular type of star). Therefore there is not too much concern about the interpretation of probabilities.

- In any case we have to make an assumption to estimate uncertainty: the ensemble of distributions that are considered as candidates to reproduce the data. More fundamentally, it is assumed that it makes sense to quantify uncertainty by a frequency, or by the hypotheses of Cox's theorem. The uncertainty on whether these assumptions are correct cannot be expressed in terms of probability, be it considered as a frequency or as a measure of plausibility. We have to keep in mind that confidence or credible sets of values rely on these assumptions. It is sound to compute these set with other assumptions (other models, other priors) in order to evaluate the robustness of the result.

- In general, we should not ignore an information we have. One should keep in mind that the whole process of arriving to conclusions impacts the validity of inferences.

In particular, the trust in conclusions should depend on whether the hypotheses were formulated before or after seeing the data. The more we search, the more likely we are to find something. There has to be some sort of increasingly stringent requirements on detection claims as we try more hypotheses.

- As far as possible the methods should be tested on simulated cases. Data are generated according to a probability law, then the process of target selection, model acquisition up to final claim is modelled. This allows to quantify uncertainty as "if the data are generated with probability distribution f and the analysis is done with a process $\mathcal{P}$ then I will obtain a pattern of right and false claims C". This is still not perfect, because it will not take into account all the particular details of a given situation.

- Complete analysis are often impossible, most of the time we will be content to find a model for the data that seems reasonable or which is better than another one. Assumptions can simplify the problems, the danger is then to "see what we expect to see".

1.3.3 Challenges of radial velocity data analysis

After these theoretical considerations on inference methods, let us come back to radial velocity signals. In this section, we will show what are the key challenges of their analysis in cases where everything is under control, that is all the signals are simulated. This will provide useful guidelines to tackle real signals, where the uncertainty on the model will make the analysis more complicated. If we neglect the gravitational interactions between planets, the ideal observations at a time t look like

$$y(t,\boldsymbol{\theta}) = v_\star + \sum_{i=1}^{n_p} K_i(\cos\left(\omega_i + \nu(t, e_i, P_i, \omega_i, M_{0,i})\right) + e_i \cos\omega_i) \tag{1.22}$$

where the subscript i designates planet i and n_p is the total number of planets. $v_\star$ is the velocity of the system barycenter. Now let us consider that we have taken N noisy measurements at times $t_1, ...t_N$, that we consider as a vector with N components, $\boldsymbol{y} = (y_k)_{k=1..N}$. We further assume that the noise on our observations is Gaussian, independent and identically distributed of variance σ^2. It means that the probability of observing $\boldsymbol{y}$ knowing $\boldsymbol{\theta}$ (the likelihood) is

$$f(\boldsymbol{y}|\boldsymbol{\theta}) = \frac{1}{(2\pi)^{\frac{N}{2}}\sigma^N}\, \mathrm{e}^{-\frac{1}{2\sigma^2}\sum\limits_{k=1}^{N}(y_k - y(t_k,\boldsymbol{\theta}))^2} \tag{1.23}$$

This model means, in the frequentist interpretation, that the result of an experiment would give $\boldsymbol{y} \in S$, where S is a measurable set of $\mathbb{R}^N$ (a set of possible observations) one in a $\int_S f(\boldsymbol{y}|\boldsymbol{\theta})d\boldsymbol{y}$ times. Note that if the noise had a covariance matrix $\mathbf{V}$, equation (1.23) would become,

$$f(\boldsymbol{y}|\boldsymbol{\theta}) = \frac{1}{(2\pi)^{\frac{N}{2}}|\mathbf{V}|}\, \mathrm{e}^{-\frac{1}{2}(\boldsymbol{y}-\boldsymbol{y}(\boldsymbol{t},\boldsymbol{\theta}))^T\mathbf{V}^{-1}(\boldsymbol{y}-\boldsymbol{y}(\boldsymbol{t},\boldsymbol{\theta}))}\,. \tag{1.24}$$

Let us remind that the covariance matrix $\mathbf{V}$ of a random vector $\boldsymbol{\epsilon}$ whose mean is zero is

$$\mathbf{V}_{kl} = \mathbb{E}\{\epsilon_k \epsilon_l\}. \tag{1.25}$$

Be the analysis Bayesian or frequentist, we will need to compute $f(\boldsymbol{y}|\boldsymbol{\theta})$ in as many $\boldsymbol{\theta}$ as possible. When the likelihood is given by equation (1.23), it is equivalent to compute $2\sigma^2(-\log f(\boldsymbol{y}|\boldsymbol{\theta}) + N \log \sqrt{2\pi}\sigma)$, which is a strictly decreasing function of the likelihood. This one is equal to $\sum_{k=1}^{N}(y_k - y(t_k, \boldsymbol{\theta}))^2$, more compactly written as $\|\boldsymbol{y} - \boldsymbol{y}(\boldsymbol{t}, \boldsymbol{\theta})\|^2$, where $\|\cdot\|$ designates the Euclidian norm.

Basically, it means we want to know how far a given model is from the observation in terms of the usual distance, but in N dimensions. In particular, the maximum likelihood corresponds to the minimum distance. Is this one hard to find? First we can simplify the problem by noticing that the parameters we use are not optimal. We can write the velocity induced by a planet

$$K(\cos(\omega + \nu(t, e, P, \omega, M_0)) = A\dot{U}(t, k, h, P) + B\dot{V}(t, k, h, P) \tag{1.26}$$

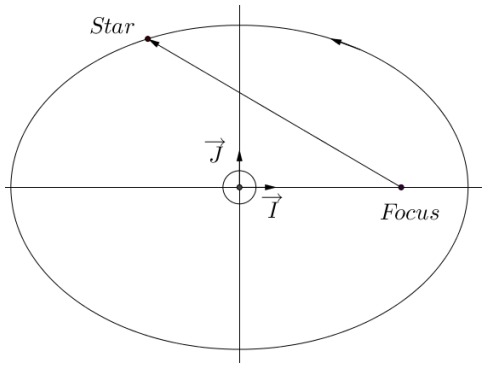

where $\dot{U}$ and $\dot{V}$ are the components of the velocity in the orbital plane, which are the projections respectively onto $\vec{I}$ and $\vec{J}$ (see figure on the left). The variables k and h are defined as $k = e\cos\omega$ and $h = e\sin\omega$. This parametrization has the advantage of having two parameters on which the model depends linearly, A and B. As a consequence, it is easy to find $A^\star$ and $B^\star$ that minimize $\|\boldsymbol{y} - \boldsymbol{y}(\boldsymbol{t}, \boldsymbol{\theta})\|$ via a linear least square when k, h, P are fixed.

As V is also a linear parameter, this leaves three non linear parameters per planet. Exploring a parameter space of 3, 6 or 9 parameters might not seem like too problematic, but unfortunately when searching frequencies, this task is lengthy due to the great number of local minima.

Let us consider a single planet and make an experiment. We take the 28 measurement dates of CoRoT-9 (Bonomo et al. 2017b) and generate a sinusoidal signal plus a Gaussian noise of standard deviation 1 m.s^{-1}. This is done in two cases: $K = 1$ m.s^{-1} and $K = 3$ m.s^{-1}. To visualise the minimization of χ^2, we proceed as follows.

1.3.4 Single planet

We first consider only circular orbits, that is $e = 0$. For each period P, we fit a model depending on linear coefficients a, b, c

$$y_P(t, a, b, c) = a\cos\frac{2\pi}{P}t + b\sin\frac{2\pi}{P}t + c. \tag{1.27}$$

on the data y, which is a linear least square problem. Let us denote by $\boldsymbol{y}_P^\star$ the best fitting sine model at period P. Instead of representing $\|\boldsymbol{y}_P^\star - \boldsymbol{y}\|$, we plot

$$\mathrm{GLS}(P) = \frac{\|\boldsymbol{y}\|^2 - \|\boldsymbol{y}_P^\star - \boldsymbol{y}\|^2}{\|\boldsymbol{y}\|^2}, \tag{1.28}$$

which carries the same information but should be more familiar to the reader, since it is exactly the definition of the Generalized Lomb-Scargle periodogram (GLS), or floating-mean periodogram (Ferraz-Mello 1981; Cumming et al. 1999; Reegen 2007; Zechmeister & Kürster 2009a). Any local maximum of this figure corresponds to a local minimum of the distance between the model and the data.

Figure 1.15 (top) displays several features that are common to noisy RV signals. First of all, the tallest peak of the periodogram does not correspond to the true period (101.35, in red), but to a signal near 1 day, spotted by the gray line on the left of the graph. Furthermore, there is a peak near the true period, centred on 76.62 days, that is almost as tall as the one corresponding to the true period.

Let us now move to Keplerian models. We denote by $\boldsymbol{\theta}'$ the orbital parameters that are not eccentricity and periastron, these are A, B and P as defined in equation (1.26) plus a constant offset $v_\star$ as defined in equation (1.22). We compute $\chi^2(e,\omega) = \min_{\boldsymbol{\theta}'} \|\boldsymbol{y}(\boldsymbol{\theta}) - \boldsymbol{y}\|^2$ for a tight grid of values of e and ω, where the minimization is performed with a non-linear Levenberg-Marquardt (Levenberg 1944; Marquardt 1963) algorithm. Such a grid of χ^2 is computed twice, with the period initialized at 101.35 and 76.62 days. The results are shown in color code in figure 1.15, respectively bottom left and bottom right. A dark blue region corresponds to lowest χ^2 (best fits), and the light yellow corresponds to high χ^2 (poorest fits). Let us make two remarks: from $e \gtrsim 0.6$, local minima start to appear. Furthermore, the model giving the best fit is obtained for $P = 76.62$ with $\chi^2 = 13.8$. This corresponds to a reduced χ^2 of 0.62, which means there is a good agreement with that model. How can we prevent this model from being validated?

One could argue that observations of a 1 m.s^{-1} signal with a 1 m.s^{-1} noise is already a limit case. Indeed, a more careful analysis shows that the highest peak has a level of significance of only 15 % (a peak at least as high is observed only in 15 % of the cases assuming the signal is pure Gaussian noise). This level is not considered high enough to claim a confident detection. The point was rather to express concerns on how good the fit is with a spurious period and a spurious eccentricity. The experiment is now done with a signal with $K = 3$ m.s^{-1}. Strong peaks are still present around one day and 77 days. This is best explained with a concept called the spectral window. This one will be discussed in the next chapter so we will not enter into the details to avoid redundancy. The difference between the 1.00072436[3] day peak and the correct one (in red) is hard to spot on the figure but the 100 days peak is indeed the true period. Note that in both cases, the orbit was circular but the eccentricity was clearly over-estimated. One of the main contributions of this thesis is to study the reason behind this bias in detail (see chapter 3).

Very eccentric planets

Now what if the planet has the same semi-amplitude, but is very eccentric? A Keplerian signal with $e = 0.9$ is simulated and the results are displayed figure 1.17. We see three signals at short periods, whose normalized reduction in sum of square (equation (1.28)), is above 0.55 and is therefore significant with a false alarm probability of 5 %. On the contrary, the true signal does not appear on the periodogram. This one is found when computing the χ^2

[3]Such a high number of digits might be puzzling, but it originates from the fact that periodograms have a constant resolution in frequency, not in periods. The smaller periods, the more digits are needed to obtain a certain precision on the frequency.

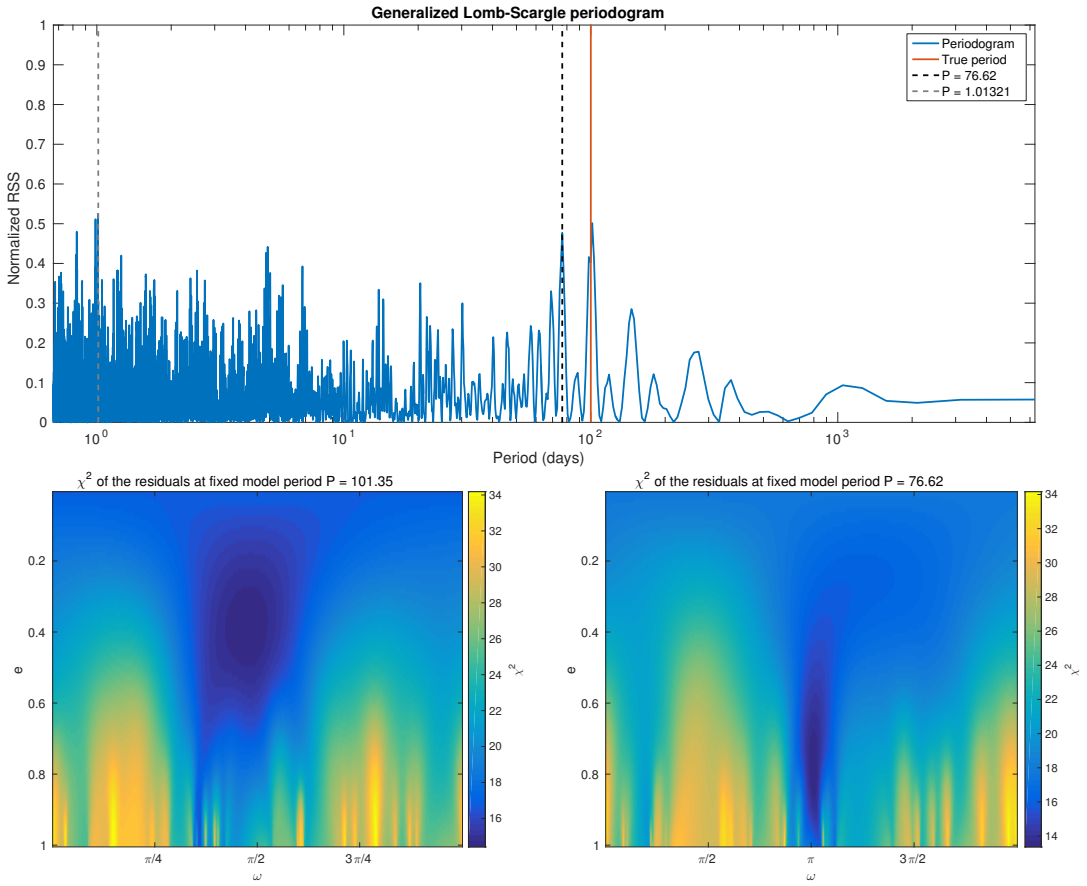

Figure 1.15: Top: periodogram of a signal of period 100 days, of amplitude $K = 1$ m.s^{-1}, evaluated on the 28 measurement times of CoRoT-9 Bonomo et al. (2017b). Bottom: maps of the χ^2 of the residuals for a Keplerian fit with a grid of (e, ω). For each value of (e, ω), the minimization is performed with a local minimization algorithm (Levenberg-Marquardt type, coded in matlab) over the orbital elements A, B and P as defined in equation (1.26) plus a constant offset $v_\star$ as defined in equation (1.22). The non linear fit is initialized with a period at 101.35 days (bottom left) and 76.62 days (bottom right).

map at $P = 100$, where a minimum χ^2 island appears at high eccentricity with $\chi^2_{\rm min} \approx 35$. The examination of this map also shows that this island is centered on a local minimum. If a non-linear least-square fit of a Keplerian model was initialized at the correct period, it would converge to the true eccentricity. Unfortunately there was no obvious indication that it was worth looking for a signal at this period.

False alarm probability (FAP)

We used above the expression "false alarm probability of 5 %" without defining it. Here it means that if the data is pure white, independent noise with $\sigma = 1$ m.s^{-1} (our model), the value of the maximum peak of the periodogram will be at least as high in 5 % of the cases. False alarm probabilities can be defined in every experiment. It requires to choose a so called

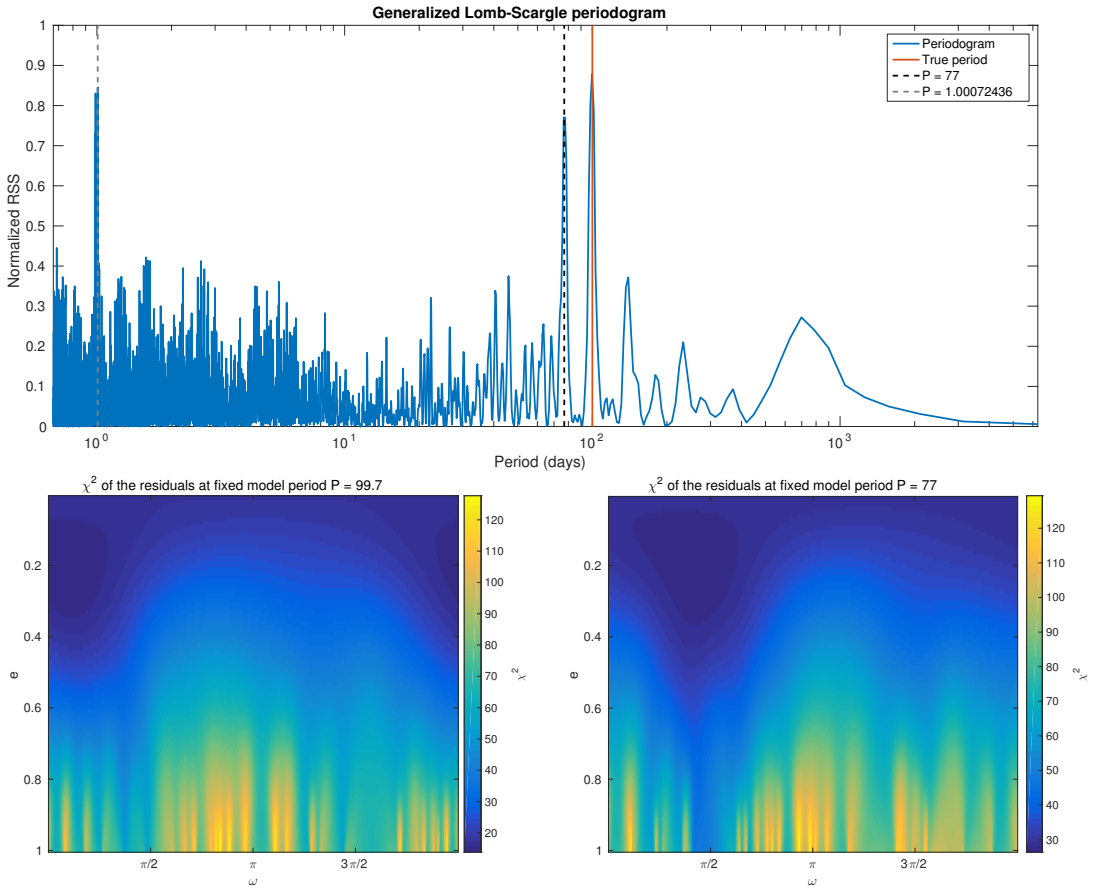

Figure 1.16: Periodogram of a signal of period 100 days, of amplitude $K = 3$ m.s^{-1}, evaluated on the 28 measurement times of CoRoT-9 Bonomo et al. (2017b). Bottom: maps of the χ^2 of the residuals for a Keplerian fit with a grid of (e, ω). For each value of (e, ω), the minimization is performed with a local minimization algorithm (Levenberg-Marquardt type, coded in matlab) over the orbital elements A, B and P as defined in equation (1.26) plus a constant offset $v_\star$ as defined in equation (1.22). The non linear fit is initialized with a period at 100 days.

"test statistic" T, which is a function of the experiment outcome $\boldsymbol{y}$. Now if we suppose that the experiment outcome is a random variable $\boldsymbol{Y}$ that follows a certain law (for instance is pure Gaussian noise), and denote that hypothesis H_0, the distribution of $T(\boldsymbol{Y}|H_0)$ can be computed. For a given experiment outcome, the false alarm probability, or here p-value of the statistic T evaluated in $\boldsymbol{y}$, is

$$p(\boldsymbol{y}, T) = \Pr\{T(\boldsymbol{Y}|H_0) \geqslant T(\boldsymbol{y})|H_0\}. \tag{1.29}$$

In other words, how frequently should the test give T at least as high as the one measured, under the hypothesis that Y follows the distribution H_0? Intuitively we expect that if the value is below a certain value α, H_0 can be rejected. One can show that the p-value is uniformly distributed under H_0, which implies $\Pr\{p(\boldsymbol{y}, T) \leqslant \alpha|H_0\} = \alpha$. In simpler terms, the probability of rejecting H_0 *knowing* H_0 is true is α. It implies that the probability of rejecting

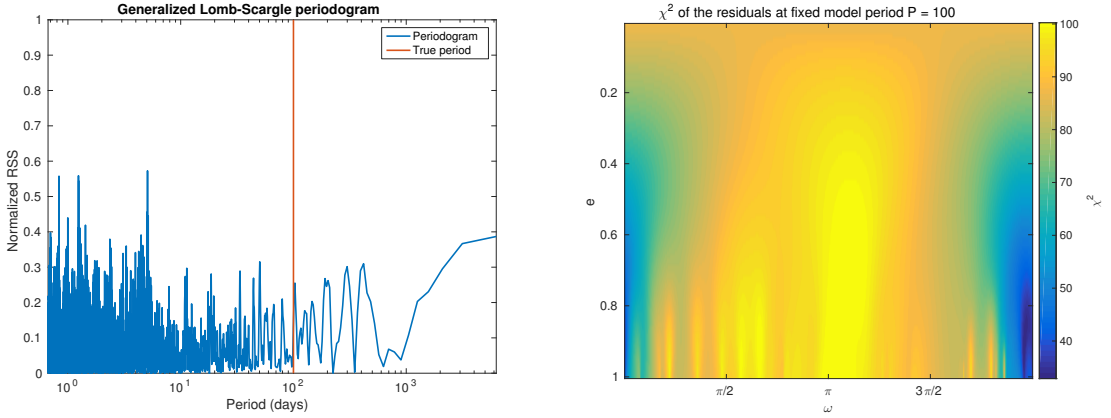

Figure 1.17: Periodogram of a signal of period 100 days, of amplitude $K = 3$ m.s^{-1} and eccentricity $e = 0.9$, evaluated on the 28 measurement times of CoRoT-9 Bonomo et al. (2017b).

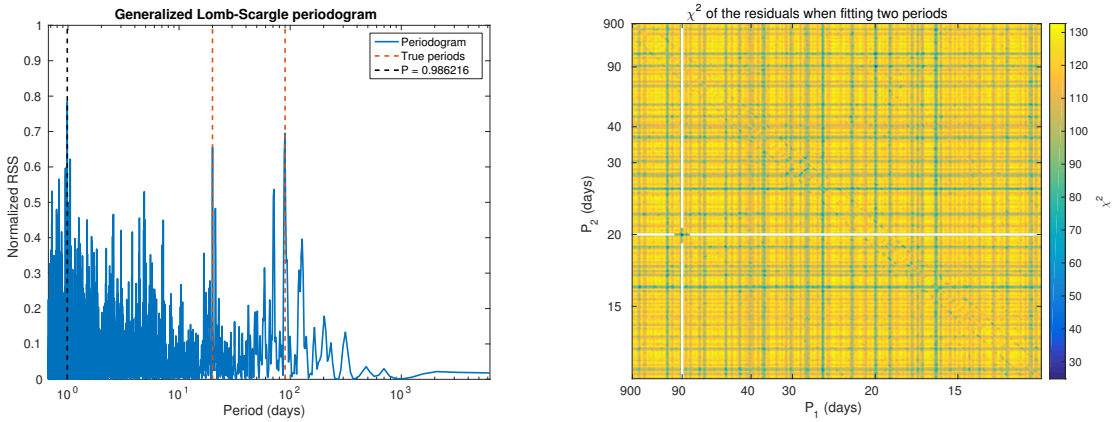

Figure 1.18: left: Periodogram of two signals with respectively $K = 5, 3, 2$ m.s^{-1}, $P = 90, 20, 5$ days and eccentricity $e = 0$ evaluated on the 28 measurement times of CoRoT-9 Bonomo et al. (2017b). Noise is Gaussian and independent with $\sigma = 1$ m.s^{-1}. Right: squared norm of the residuals of the fit of two sines plus a constant.

H_0 and H_0 is true verifies $\Pr\{(\text{reject } H_0) \cap (H_0 \text{ is true})\} = \Pr\{\text{reject } H_0 | H_0\} \Pr\{H_0\} \leqslant \Pr\{\text{reject } H_0 | H_0\} = \alpha$. In exoplanet searches, $\alpha = 0.1\%$ is considered to be a strong evidence.

The efficiency of p-values as planet detection indicators is debatable, since they only allow to control the rate of false rejection of knowing the signal is given by a certain model. The fact that white noise can be rejected does not automatically mean that one can claim a planet detection. The high periodogram peak might arise from another type of variation such as under-estimated noise level, unseen companion at another period, non-Gaussian noise, etc. As the confidence level improves these issue become less and less of a problem. When using small values of α, p-values have proved to be useful for analyses of radial velocity signals.

Another problem of the p-value is that it only controls the rate of false positives, and not the one of false negatives. Let us suppose there is only one alternative to H_0, denoted by H_1. If H_1 is rare (its probability is less than the p-value), then it means that on average, we will

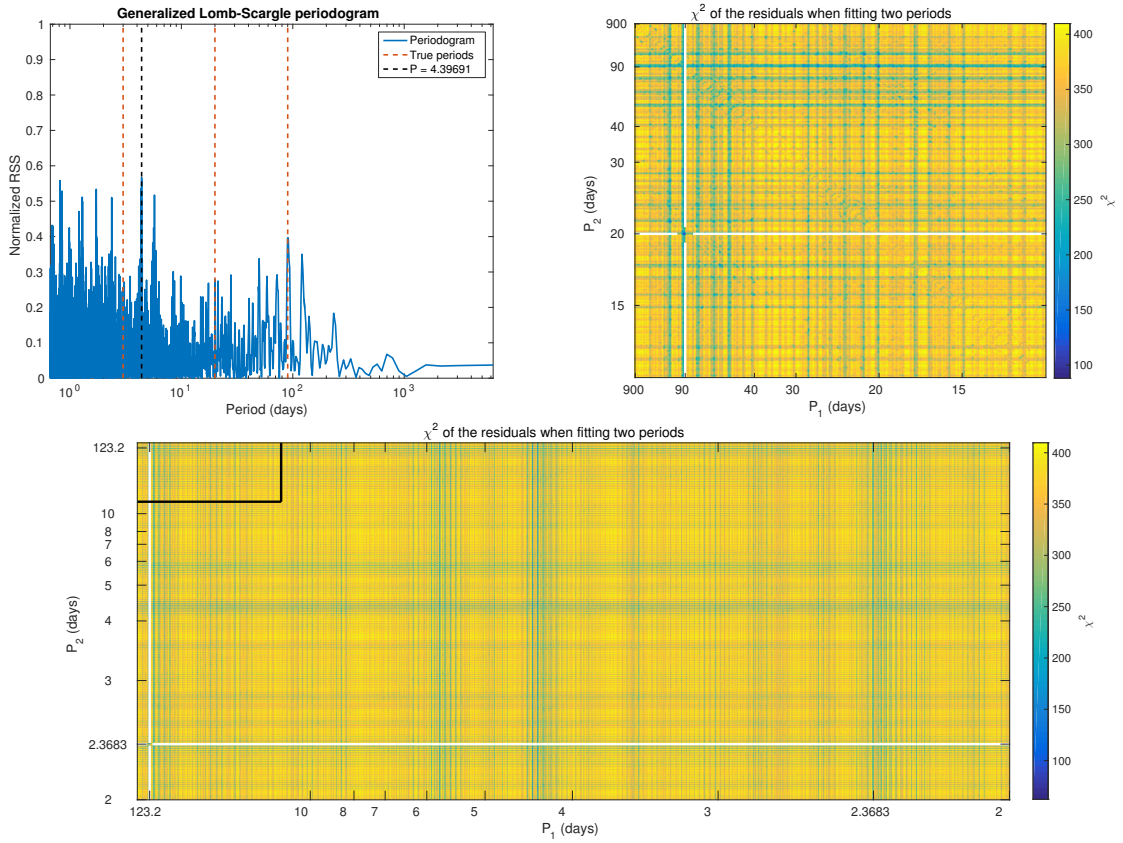

Figure 1.19: Periodogram of three signals with $K = 3$ m.s^{-1} and eccentricity $e = 0$, with periods 90 and 20 days evaluated on the 28 measurement times of CoRoT-9 Bonomo et al. (2017b). Noise is Gaussian and independent with $\sigma = 1$ m.s^{-1}. Right: squared norm of the residuals of the fit of two sines plus a constant. Bottom: also represents the squared norm of the residuals of the fit of two sines plus a constant for a wider range of periods. The χ^2 map represented on the upper right figure is delimited by the plain black lines on the north-west corner of the plot.

have more false discoveries than true discoveries. In the limit case where H_1 never happens, all discoveries are spurious. This problem was pointed out by Soric (1989) for medical studies, and was already encountered in section 1.3.2 in the illness test example. Here also, if we ignore the fact that the data belong to a certain population spurious inferences might be expected. For a purely frequentist treatment of the specific issue of false discovery rate with p-value, see Benjamini & Hochberg (1995). The presence of planets does not seem so rare that this should be too major a cause for concern.

From a practical standpoint, the distribution of the test statistic — here the maximum of the periodogram — is not easy to compute analytically. To estimate FAPs, it is common practice to generate thousands of noise realization, either by generating them from a theoretical law or by shuffling residuals of the best fit model. For each realizations y_k, the test statistic $T(\boldsymbol{y}_k)$ is computed. Then one simply counts the number of occurrences of $T(\boldsymbol{y}_k)$ above $T(\boldsymbol{y})$, which divided by the number of samples gives an estimate of the FAP. This

process, known as bootstrap (Efron 1979) is accurate but requires numerous computations, and is therefore slow.

1.3.5 Multiplanetary case

These difficulties are severely worsened by the possibility of observing multiplanetary systems. As we shall see next chapter, there is a non negligible risk of taking a combination of signals from several planets as a single planet. To illustrate it, we generate two signals of circular orbits with $K = 3$ m.s^{-1} at 90 and 20 days, still on the same array of time measurements with random phases and independent Gaussian noise.

The periodogram of the data is displayed in figure 1.18 (top left). The two periods appear relatively clearly and are significant. There are still unpleasant aliases around one day, and the maximum of the periodogram is at $P = 0.986216$. We now fit two sines at a time plus a constant and compute the χ^2, that is the squared norm of the residuals. The frequencies of the sines are taken on a regularly spaced grid. We find that the minimum is attained at the correct periods, 89.5 and 20.01 days, localized by white bands in figure 1.18 (top right).

Now we do the same procedure on another signal. This one is simulated in the same way, except that the amplitude of the 90 days period is increased to 5 m.s^{-1}, and we add another circular orbit with semi-amplitude $K = 2$ m.s^{-1} at $P = 3$ days. Figure 1.19 shows the periodogram of the signal and reveals that the maximum is on neither of the true periods. However, when searching for two signals simultaneously the minimum is again at correct periods (87.7 and 20.01 days, $\chi^2 = 80$). Unfortunately, when doing the search on a wider scale the global minimum arises at 123.2 and 2.3683 days, $\chi^2 = 62$, which is localized by white bands in figure 1.19.

The black box represents the space that is delimited by figure 1.19 (top right). This is a tiny fraction of the space of parameters explored to find the new global minimum. This highlights an important issue: the more parameters are tried, the more likely we are to find a spurious best fit. In that case, the measurements are not closely spaced enough to hope for detections at ≈ 2 days. Since the RMS of the signal is too high for it to be pure noise, what would have likely happened is that more measurements would have been taken before claiming anything. Note that all examples taken here were not specifically designed to be particularly pathological. The periods were drawn at random in $\log P$ then fixed, and the problems shown in the chosen examples appeared after generating one to five realizations of the noise.

Consequences on error estimations

Figure 1.19 (top right and bottom) exhibit numerous stripes, which all indicate local minima. With additional assumptions discussed in chapter 2, the situation might be better. But even in this case, it is not completely ensured that there is an exotic combination of frequencies that give the best fit.

Besides, when calculating the error via posterior distributions, one has to sample parameters from the posterior distribution. In MCMC methods Metropolis et al. (1953); Hastings (1970), it will be difficult for walkers to get out of a local minimum in period. In practice the walkers do stay confined in the local minimum. This basically means that the computation of the error via MCMC can be though of a conditional probability "knowing that the period

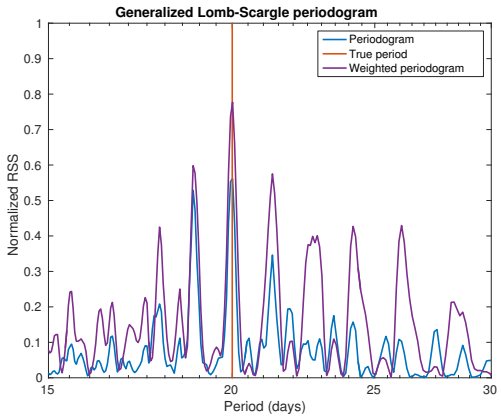

Figure 1.20: Periodogram of a system with $K = 3$m/s on the measurement times of CoRoT-9 (Bonomo et al. 2017b) with a matrix of elements $\sigma_{ij} = \mathrm{e}^{-\frac{|t_i - t_j|}{\tau}}$, $\tau = 25$ days.

is between P_1 and P_2", where P_1 and P_2 are values that are too improbable to be crossed. This effect results in underestimation of error bars.

Noises

On top of these issues, there are complicated noises which can bury signals or mimic them. If the noise is Gaussian multivariate, such that its distribution is given by equation (1.24), the periodogram we computed is not equivalent to the maximum likelihood over the linear parameters of the model. A straightforward derivation shows that we should be taking

$$\mathrm{GLS}_W(P) = \frac{\|\mathbf{W}\boldsymbol{y}\|^2 - \|\mathbf{W}(\boldsymbol{y}_P^\star - \boldsymbol{y})\|^2}{\|\mathbf{W}\boldsymbol{y}\|^2}, \tag{1.30}$$

instead of equation (1.28) where $\mathbf{W} = \mathbf{V}^{-\frac{1}{2}}$. Figure 1.20 shows an example of what happens when the noise is not taken into account (blue curve): the signal might be underestimated and pass below the significance threshold. In the example presented, the difference is not spectacular but assuming the residuals of the fit follow the same distribution as in the white noise case (which is only approximately true, as shown in Appendix E) a 0.2 difference of value of the periodogram here means a factor 10 difference of false alarm probability (0.21 to 0.02), which is non negligible.

In section 1.2.5, particular attention was given to various sources of stellar noise. How does it look like? Figure 1.21 shows a signal simulated for the RV Fitting Challenge (Dumusque 2016; Dumusque et al. 2016), which is based on SOAP 2.0 (Dumusque et al. 2014) and our current understanding of long term evolution of magnetic activity. Since we will enter more into the details of this challenge later on, let us just record that they look extremely different.

Overall

In summary, we believe the key aspects of radial velocity data seem to be the following:

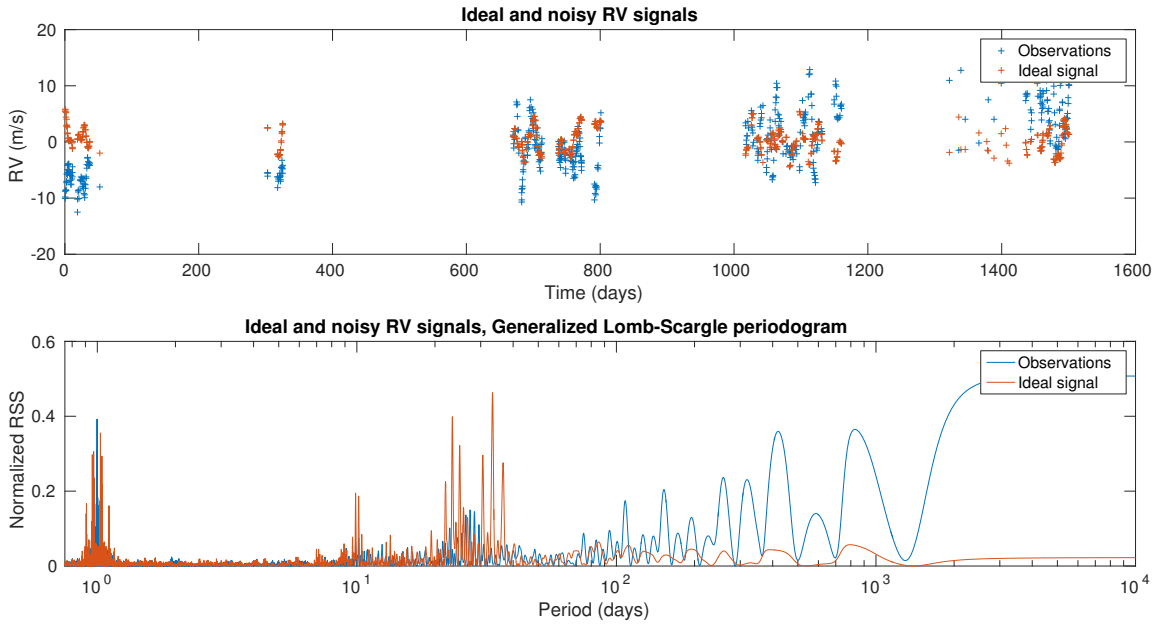

Figure 1.21: System one of the RV Fitting Challenge (Dumusque 2016). Top: RV data with and without simulated stellar noise. Bottom: Generalized Lomb-Scargle periodogram of both time series.

- Exploring complicated phase spaces with many local minima, which are first due to the search of several frequencies and complicated structures at high eccentricities. This is increasingly difficult with the number of planets and the complexity of the noise model.
- Correctly taking into account the (incomplete) exploration of models in the error bars. Also, how to decide how many planets should be reported, since different configurations will fit acceptably the data set. This is related to the fundamental problem outlined in section 1.3: the more we search, the more likely we are to find something that fits the data spuriously.
- Identifying apt models for stellar noise and instrumental noises.
- Taking into account our uncertainty on models themselves to avoid spurious detections.
- Improving the mathematical understanding of the signal, in particular obtaining analytical formulae, to fasten up data fits and computations of error bars.

Note that this concerns individual radial velocity systems, the problem of correction of planets population has not been discussed yet. Some aspects of the problem are discussed in chapter 3.

1.3.6 Existing analysis techniques

Overview

The processing of a given spectrum to obtain the time series of radial velocity or other indicators, and the analysis of the resulting time series are in general treated separately. Articles

addressing the problem of inference of planetary companions directly from the spectra are starting to emerge, but are not yet applied to real cases. For instance, Davis et al. (2017) applies principal component analysis to the spectra to unveil directions of variability, in the same spirit as Artigau et al. (2014), who used it to separate telluric lines from stellar absorption lines. The online tool AGATHA (Feng et al. 2017) allows to compute the periodogram on certain ranges of wavelength to unveil possible variations. Full exploitation of the spectral content is also being explored via manifold learning (Stenning et al, in prep.). The idea behind these techniques is that we expect the correlated noise due to the stellar variability, the atmosphere or the instrument to have signatures that depend on the wavelength, while the signature of the RV should not.

We will now focus on the time series analysis, which represents the vast majority of the RV data analysis literature. Let us note that time series analysis techniques already use the difference between the spectral signatures of stellar variability and planetary companions via FWHM, BIS (see figure 1.11) or specific spectral lines such as the calcium II H & K (Vaughan et al. 1978) or H_α (Pasquini & Pallavicini 1991; Montes et al. 1995). The radial velocity signal induced by a companion has another property that can be leveraged to validate a planet: it is stable over time. We expect the stellar variability not to keep the same frequency over long periods of time.

Using the stability over wavelengths and time of the companion-induced RV variation seems promising, since it is a step towards guaranteeing the "uniqueness" of the companion hypothesis. It is likely that these idea will build upon previous analyses on radial velocity signals, which we present in the following sections. We distinguish three types of works: periodograms, Bayesian analysis and Gaussian processes. This distinction is made based on what is put forward in the published analyses, but as we shall see the frontier is not so clear.

Periodograms

We already defined the generalised Lomb-Scargle periodogram in equation (1.28) and used it in section 1.3.3 to show some of the challenges of radial velocity data analysis. Since it is a widely used tool and is key in the claims of planet detection, we will spend a little more time introducing it from a historical perspective. The term "periodogram" was initially introduced by Schuster (1898). In his own words:

> *It is the object of this paper to introduce a little more scientific precision into the treatment of problems which involve hidden periodicities, and to apply the theory of probability in such a way that we may be able to assign a definite number for the probability that the effects found by means of the usual methods are real, and not due to accident.*

Schuster was not the first author to try to unveil periodicities in time series. He was however the first one to use probability theory to assess whether a periodicity was due to non-periodic fluctuations. He gives the following definition of the periodogram. For a function y of time and a frequency ω,

$$r(\omega) = \frac{2}{T}\left\{\left(\int_{t_1}^{t_1+T} y(t)\cos\omega t\ dt\right)^2 + \left(\int_{t_1}^{t_1+T} y(t)\sin\omega t\ dt\right)^2\right\}^{\frac{1}{2}} \tag{1.31}$$

Which is followed by

> *for different adjacent values of ω, the quantities r will fluctuate about some mean value r' so that the probability of r being greater than $\lambda r'$ is $\mathrm{e}^{-\frac{\pi\lambda^2}{4}}$.*

In more modern terms, Schuster projects the time series $\boldsymbol{y} = (y(j\Delta t))_{j=1..N}$ (discretized with evenly spaced sample, the integrals become sums) on the vectors $\boldsymbol{u}(\omega) = (\cos(\omega j\Delta t))_{j=1..N}$ and $\boldsymbol{v}(\omega) = (\sin(\omega j\Delta t))_{j=1..N}$ where N is the number of samples. Assuming $\boldsymbol{y}(t)$ is pure white Gaussian noise, the scalar products of $\boldsymbol{f}$ and $\boldsymbol{u}(\omega)$ and $\boldsymbol{v}(\omega)$ both follow a Gaussian distribution, therefore their modulus (1.31) follows a Rayleigh distribution.

After Schuster (1898), techniques have been refined to account for irregularly spaced measurement. Lomb (1976) and Scargle (1982) extend the formula of the periodogram to non evenly spaced measurements but the spirit is the same: the data is projected onto the linear space generated by the vector $\boldsymbol{u}(\omega), \boldsymbol{v}(\omega)$ to obtain a sinusoidal model $\boldsymbol{y}_{\text{fit}}(\omega) = a(\omega)\boldsymbol{u}(\omega) + b(\omega)\boldsymbol{v}(\omega)$. The distribution of $P(\omega) = \|\boldsymbol{y}\|^2 - \|\boldsymbol{y} - \boldsymbol{y}_{\text{fit}}(\omega)\|^2$ is computed in Lomb (1976). The higher this quantity is, the closer to zero is the norm of the residuals, which indicates a good fit between a certain period and a periodicity at ω. The major improvement lies in the fact that not only the distribution of the $P(\omega)$ is computed under the assumption that $\boldsymbol{y}$ is pure Gaussian noise, but also the distribution of the maximum of $(P(\omega_k))_{k=1..N}$ where $P(\omega_k)$ are independent random variables. Indeed, when one searches for periodicity, a range of periods is explored. As argued section 1.3.1, the fact that one has searched in a wide space must be taken into account. Even if the signal is pure noise, if all sinusoidal models with periods ranging from zero to infinity are considered, there will always be models reproducing the observations to arbitrary accuracy (the set $\{\cos\omega\boldsymbol{t}/\|\cos\omega\boldsymbol{t}\|, \omega \in \mathbb{R}_+\}$ is dense on the sphere for almost all $\boldsymbol{t}$).

When the period is not uniformly sampled, it is not clear how many independent periods in a certain range are in the signal, therefore one can use the bootstrap technique described in section 1.3.4 to compute the false alarm probability. One generates M realizations of a Gaussian noise numbered by k (or M random shufflings of the best fit residuals), computes the periodogram for each of them and its maximum $P_{\max,k}$, and uses the empirical distribution $(P_{\max,k})_{k=1..M}$ as a proxy for the theoretical one. It is possible to fit a model to that empirical distribution, especially to consider an "effective" number of independent frequencies calibrated on test data sets as a function of the frequency span and the number of measurements (e.g. Horne & Baliunas 1986; Cumming 2004).

These principles can be extended to comparisons between non purely sinusoidal models. For instance Ferraz-Mello (1981) adds a constant to the sine model and Cumming et al. (1999) computes the distribution of the amplitude of a periodogram peak under the Gaussian independent noise hypothesis. For the same model, Zechmeister & Kürster (2009a) compiles the approximations of the false alarm probability and analytical formula for this type of periodogram. Baluev (2013b) considers the more general face of periodic function non-linearly fitted. Cumming (2004), Gregory (2007a), Zechmeister & Kürster (2009b), O'Toole et al. (2009) and Baluev (2015b) go further and use as a candidate function not a sine wave, but a full Keplerian function, thus defining a "Keplerian periodogram". The difference between the approaches lies in the way the parameter space is explored, and the computation of the statistical significance of a detection. Note that Gregory (2007b) has "Keplerian periodogram" in the title, but it is a full Bayesian exploration of the parameter space, presented next paragraph.

A Keplerian periodogram is computationally much more demanding than a classical one,

since the model depends non-linearly not only on the period, but also on the eccentricity and the argument of periastron. This computational burden becomes especially problematic to compute false alarm probabilities. Analytical approximations are difficult to obtain and performing bootstrap requires to recompute the periodogram on several hundreds or more of randomly generated noise.

The analytical estimation of the false alarm probability as well as computational issues have been tackled in a series of paper with increasingly complex models Baluev (2008, 2009, 2013,a,b, 2015b). These ones present two features: the periodogram always compares a model H with a model H plus a candidate signal of parameters $\boldsymbol{\theta}$ with p components which always includes the frequency, and that we denote by $K(\theta)$. The periodogram can take the form

$$z_0(\boldsymbol{\theta}) = \frac{1}{2}(\chi^2(K(\boldsymbol{\theta})) - \chi^2(H)) \tag{1.32}$$

$$z_1(\boldsymbol{\theta}) = (N-p)\frac{\chi^2(K(\boldsymbol{\theta})) - \chi^2(H)}{\chi^2(H)} \tag{1.33}$$

and also others, that we do not present for the sake of brevity. The notation $\chi^2(M)$ denotes the χ^2 of the residuals when model M is fitted on the data set. In the Keplerian model, $\boldsymbol{\theta} = (\omega, e, \varpi)$ where e, ϖ designate the eccentricity and longitude of periastron.

The analytical approximation of the FAP relies on considering the functions $z_i : \omega \to P(\omega)$ or $z_i : \omega, e, \varpi \to P(\omega, e, \varpi)$ respectively as a random process and a random field. Then the Rice formula (Rice 1944) can be used to assess the probability that the random process crosses certain values, more specifically the Davies bound (Davies 1977, 1987, 2002) can be adapted to be applicable to the periodogram. The series of case considered by Baluev is best summarized in figure 1.22

More recently, Mortier & Collier Cameron (2017) recomputes the periodogram each time a point is added. If the amplitude increases approximately steadily, it is a sign of a planetary companion. If a peak alternatively appears and disappears, it might indicate a variability localized in time. Interestingly enough Schuster (1898) already considers computing the integrals of equation (1.31) on $[0, T]$, $[0, 2T]$, $[0, 3T]$... and checks if these increase more rapidly than $\sqrt{kT}$. Mortier & Collier Cameron (2017) does not use the Generalised Lomb-Scargle periodogram but the Bayesian Generalised Lomb-Scargle periodogram (Mortier et al. 2015), presented next section.

Bayesian analysis

As said in section 1.3.2, not using a prior probability has the disadvantage of disallowing to ask the right question: the probability of the hypothesis knowing the data. Bayesian methods were introduced in the field of exoplanets by Gregory (2005) and Ford (2005). They suggested to compute the posterior distribution of the orbital elements

$$p(\boldsymbol{\theta}|\boldsymbol{y}) = \frac{p(\boldsymbol{y}|\boldsymbol{\theta})p(\boldsymbol{\theta})}{\int_\Theta p(\boldsymbol{y}|\boldsymbol{\theta})p(\boldsymbol{\theta})d\boldsymbol{\theta}} \tag{1.34}$$

where $\boldsymbol{\theta} = (\theta)_{k=1..p}$ represents the orbital parameters, $\boldsymbol{y}$ is the vector of observations and $p(\boldsymbol{y}|\boldsymbol{\theta})$ is a Gaussian likelihood as in equation (1.24). One can then obtain the marginalized

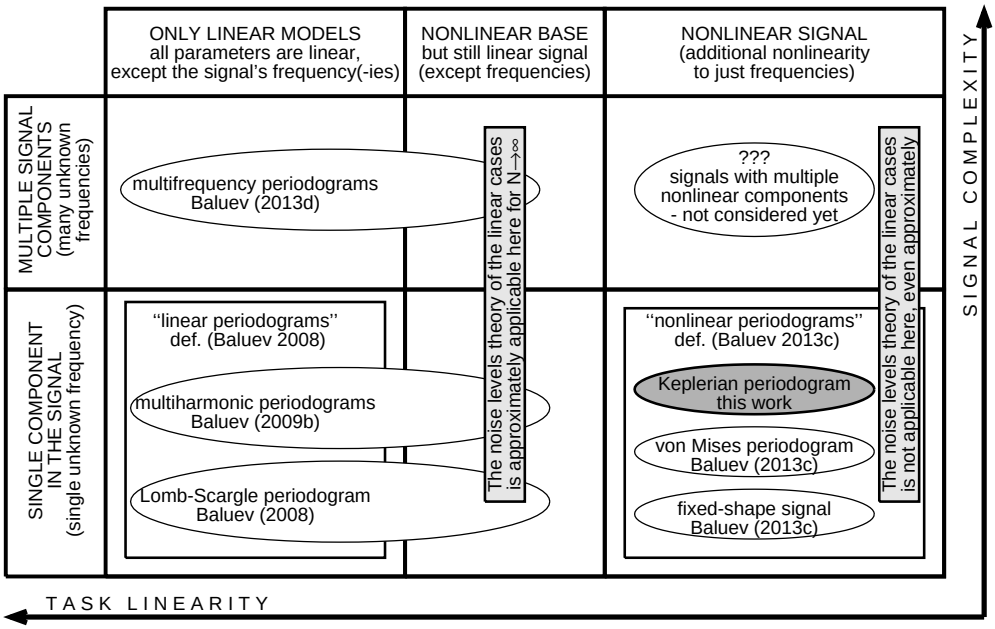

Figure 1.22: Classification of the works by Baluev as a function of two criterions, whether there is a single candidate signal (one planet) or several (vertical axis) and whether the model of the base model H is linear, non linear, or if both models H and $K(\theta)$ are fitted non-linearly. The fact that the ellipses on the left hand side reach the "non linear" column means that the theory of linear base models provides a good approximation when the model H is fitted non-linearly for each candidate signal θ, but when this one is fitted non linearly as well, the theory has to be modified (right column). The label "This work" refers to Baluev (2015b).

density for each parameter θ_k,

$$p(\theta_k|\boldsymbol{y}) = \int_{\Theta'(\theta_k)} p(\boldsymbol{\theta}'|\boldsymbol{y}) d\boldsymbol{\theta}' \tag{1.35}$$

where $\Theta'(\theta_k)$ is the set of all parameters θ_i, $i = 1..p$ $i \neq k$ and θ_k is fixed. From the marginalized posterior density one can obtain credible intervals $[\alpha_k, \beta_k]$ for all parameters θ_k. The interval is chosen so that the probability that θ_k is in $[\alpha_k, \beta_k]$ knowing $\boldsymbol{y}$ is γ, which does not depend on k,

$$\Pr\{\theta_k \in [\alpha_k, \beta_k]|\boldsymbol{y}\} = \int_{\alpha_k}^{\beta_k} p(\theta_k|\boldsymbol{y}) d\theta_k = \gamma. \tag{1.36}$$

We believe it is useful to consider this formula from a communication theory point of view. Let us suppose that an emitter sends us messages $\boldsymbol{\theta}$ with probability density $p(\boldsymbol{\theta})$. The messages here are not in English or binary code, but in "orbital parameters". The message goes through a channel (the observation process) and gives the receiver information $\boldsymbol{y}$, and we know the conditional density $p(\boldsymbol{y}|\boldsymbol{\theta})$. Formula (1.36) tells us that when $\boldsymbol{y}$ is observed, the true k^{th} component of the message θ_k is in $[\alpha_k, \beta_k]$ in $\gamma \times 100$ % of the cases. The tighter $[\alpha_k, \beta_k]$ and the higher γ, the more precise is the constraint. The confidence interval might

be chosen to be centred on the posterior mode (i.e. maximum value) and regions of equal probabilities on both sides, but other choices are possible (see Gregory 2005). [4]

One can also compute credible regions for several parameters $\theta^{(1)}...\theta^{(l)}$, $l \leqslant p$ chosen among the $(\theta_k)_{k=1..p}$. Denoting the marginalised distribution of those parameters by

$$p(\theta^{(1)}, ..., \theta^{(l)}|\boldsymbol{y}) = \int_{\theta \in \Theta'(\theta^{(1)},...,\theta^{(l)})} p(\boldsymbol{\theta}|\boldsymbol{y}) d\boldsymbol{\theta} \tag{1.37}$$

where $\Theta'(\theta^{(1)}, ..., \theta^{(l)}) \subset \Theta$ and the values of the parameters $(\theta^{(1)}, ..., \theta^{(l)})$ are fixed. The credible region R is such that

$$\Pr\{\theta^{(1)}, ..., \theta^{(l)} \in R|\boldsymbol{y}\} = \int_R p((\theta^{(1)}, ..., \theta^{(l)})|\boldsymbol{y}) d\theta^{(1)}..d\theta^{(l)} = \gamma. \tag{1.38}$$

To compute the posterior distribution (equation (1.34)), Ford (2005, 2006) use a Monte Carlo Markov Chain (MCMC) algorithm (Metropolis et al. 1953; Hastings 1970). This type of algorithms generates a sequence of pseudo-random numbers drawn from the posterior distribution (1.34). One ends up with a long sequence of parameter vectors $(\boldsymbol{\theta}_j)_{j=1..n_s}$, n_s being the total number of sample, which can reach several millions. The advantage of MCMC methods is that to generate a $\boldsymbol{\theta}_j$ sample, they do not necessitate to evaluate the whole posterior distribution. This distribution has a complicated integral at the denominator which is impractical to calculate. MCMC methods only necessitate to compute ratios of posterior distributions, in which the integral cancels out. Stopping criterion are crucial yet difficult points of the algorithms. To check that enough samples were generated to approximate the posterior distribution, several chains of $(\boldsymbol{\theta}_j)_{j=1..n_s}$ are run from different starting points. There are more formal way to check for convergence, such as computing the Gelman & Rubin (1992) test statistic.

The most widely used version of MCMC in the exoplanet community is EMCEE (Foreman-Mackey et al. 2013), based on affine-invariant sampling (Goodman & Weare 2010). We show here an example of its output on Proxima b dataset (Anglada-Escudé et al. 2016). The EMCEE version has been coded in matlab (Grinsted 2015) with a Keplerian model ($\boldsymbol{y}(\boldsymbol{\theta}) = A\dot{\boldsymbol{U}} + B\dot{\boldsymbol{V}}$ where $\dot{\boldsymbol{U}}, \dot{\boldsymbol{V}}$ is the velocity of the star in the orbital plane), a mean, trend and quadratic term. The results are plotted in figure 1.23. The plain green curves are the marginalised distribution of each parameters (equation (1.39)) and the grey areas are credible regions (equation (1.38)) for parameters taken two-by-two. To ensure that no region of the parameter space is left over, we consider a grid of values of the argument of periastron and eccentricity, minimize over all the other parameters and use the best fit values as starting points. As one can see the period is constrained between 11.184 and 11.188 days. Due to the fact that the local minima in periods are steep, the chain does not get out of the local minimum with high probability. To be meaningful, posteriors given by MCMC must be initialized at the correct periods. In this respect, this type of calculation is complementary to a periodogram approach. The periodogram can itself be computed in a Bayesian framework.

[4] Note that this reasoning holds when the messages emitted and received are in discrete spaces. If we think of the space of planetary parameters and possible observations as arbitrarily densely discretized, the interpretation is unchanged. We believe it is possible to extend this interpretation in the continuous case as a limit case (at least when the space of messages Θ is continuous) but did not investigate further the question.

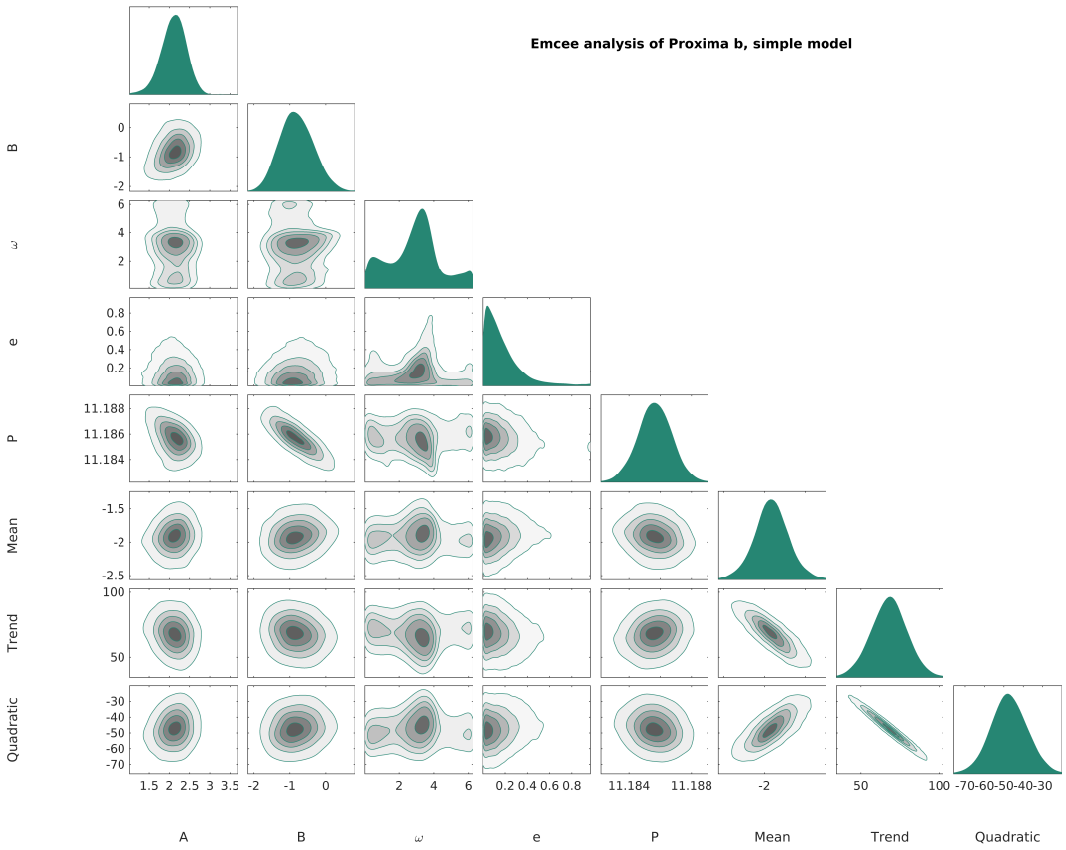

Figure 1.23: Example of MCMC analysis of Proxima b with EMCEE implemented by Grinsted (2015)

The Bayesian periodogram (Mortier et al. 2015) is the posterior density of a sinusoidal model $\boldsymbol{y}(\omega, A, B, C) = A\cos\omega\boldsymbol{t} + B\sin\omega\boldsymbol{t} + C$ marginalized over A, B, C

$$p(\omega|\boldsymbol{y}) = \int_{A,B,C} p(\omega|\boldsymbol{y}) dA dB dC, \tag{1.39}$$

for which an analytical expression is obtained. The "Keplerian periodogram" by Gregory (2007a) consists in computing the posterior density of the orbital elements, but so to explore all the parameter space by sampling the posterior via a technique called parallel tempering (Gregory 2005). The term periodogram is here to be taken in the sense of throughout exploration. It can be used to compute the evidence, as defined below.

Bayesian methods are also used to confirm planets. One computes the "evidence" also called the marginal likelihood of a model with a number of planets n_p, that we denote by $\mathcal{M}_{n_p}$, that corresponds to a parameter space Θ_{n_p}, which are all the possible combinations of orbital parameters of n_p planets plus the non-planetary effects modelled. The evidence of $\mathcal{M}_{n_p}$ is

$$p\{\boldsymbol{y}|\mathcal{M}_{n_p}\} = \int_{\boldsymbol{\theta}\in\Theta_{n_p}} p(\boldsymbol{y}, \boldsymbol{\theta}) d\boldsymbol{\theta} = \int_{\boldsymbol{\theta}\in\Theta_{n_p}} p(\boldsymbol{y}|\boldsymbol{\theta}) p(\boldsymbol{\theta}) d\boldsymbol{\theta} \tag{1.40}$$

where $p(\boldsymbol{y}, \boldsymbol{\theta})$ is the joint probability of $\boldsymbol{y}$ and $\boldsymbol{\theta}$. The denominator $p(\boldsymbol{y})$ is absent because evidences are used to compare models. A model with $n_p + 1$ planet is favoured

$$\frac{p\{\boldsymbol{y}|\mathcal{M}_{n_p+1}\}}{p\{\boldsymbol{y}|\mathcal{M}_{n_p}\}} \geqslant B, \tag{1.41}$$

The left-hand side term is sometimes called the "odds ratio" or "Bayes factor" (Kass & Raftery 1995). Evidence has been used in Gregory (2005); Tuomi & Kotiranta (2009); Tuomi (2011) and dozens of other discovery papers. In the exoplanet community, the threshold for a detection is in general $B = 150$. The ratio of posterior probability of the models

$$\frac{p\{\mathcal{M}_{n_p+1}|\boldsymbol{y}\}}{p\{\mathcal{M}_{n_p}|\boldsymbol{y}\}} = \frac{p\{\boldsymbol{y}|\mathcal{M}_{n_p+1}\}}{p\{\boldsymbol{y}|\mathcal{M}_{n_p}\}} \frac{p\{\mathcal{M}_{n_p+1}\}}{p\{\mathcal{M}_{n_p}\}} \tag{1.42}$$

so the odds ratio (1.41) is a posterior ratio when the prior probability of the models are assumed to be equal. Note that equation (1.41) is similar to the ratio (1.21) we encountered in section 1.3.2, page 27, except that to obtain $p\{\boldsymbol{y}|\mathcal{M}_{n_p+1}\}$, one has to use a prior on the orbital elements. However, here again we might want to ensure that B is large to have n_p+1 planet detection robust against a rare occurrence $n_p + 1$ planet model (that is a small value of $p\{\mathcal{M}_{n_p+1}\}/p\{\mathcal{M}_{n_p}\}$).

The main challenge of accurate evidence calculations is to explore wide parameter spaces to compute the integrals. Testing the reliability of the different techniques has been a specific discussion topic at the Extreme Precision Radial Velocity workshop in Penn State University, August 2017 (EPRV III). Six simulated signals were generated, and different teams will compute the evidence for models from 0 to three planets. The goal is to compare the different algorithms (Nested Sampling, Importance sampling, fusion MCMC...) and present the results in a forthcoming paper (Ford+ 2017).

A possible drawback of the evidence is that it might be high for a model with an additional planet because there are many candidates, but none of them concentrates the posterior probability. It must be ensured that not only the $n_p + 1$ model is likely, but that there is a strong mode of the posterior, ensuring the planet is well located in the parameter space. The question of whether having multiple candidates but none of them completely strongly supported could arise was explicitly discussed with Thomas Loredo and James Jenkins at the EPRV III. If Thomas Loredo did indeed see such situations in certain cases (not concerning exoplanets), James Jenkins said that he observed unclear situations for $B = 100$ but not for $B = 150$. Ensuring the uniqueness of a strong candidate is however a conceptual problem to keep in mind.

The computation of credible intervals as well as evidence ratios relies on a certain model, which is sometimes, following Jaynes & Bretthorst (2003), denoted by I in all the distributions mentioned above. For instance, the prior is denoted by $p(\boldsymbol{\theta}, I)$. In chapter 3, we show that one of the effect of mismodelling on average to lead to an overestimation of eccentricities.

Gaussian processes/correlated noise

In section 1.2.5, we saw that the RV variation induced by the star, that we denote here by $\boldsymbol{y}_\star(t)$, could be unpredictable, but present correlation. By this, we mean that the value $\boldsymbol{y}_\star(t)$ and $\boldsymbol{y}_\star(t+\tau)$ are not independent for τ in a certain time scale. Depending on the type of noise

(due to p-modes, granulation, activity...) the time scale may vary. A natural way to model such behaviour is to consider that $\boldsymbol{y}_\star(t)$, or more generally the stochastic part of the signal $\boldsymbol{\epsilon}$ is a multivariate Gaussian distribution with zero mean. It is then fully characterized by its covariance matrix defined in equation (1.25), that we denote by $\mathbf{V}$. The noise is in general assumed stationary, which means there exists a function c such that for any two instants t, t',

$$\mathbb{E}\{\epsilon(t)\epsilon(t')\} = c(|t - t'|).$$

Then the covariance matrix is

$$\mathbf{V}_{kl} = c(|t_l - t_k|)$$

We can think of the noise $\epsilon(t)$ as a continuous stochastic process (a function which associates a random variable to each real number t). If we impose that on any finite samples $\boldsymbol{t} = (t_i)_{i=1..N}$ it follows a Gaussian distribution, then $\epsilon(t)$ is said to be a Gaussian process. In the terms of Rasmussen & Williams (2005),

> *A Gaussian process is a collection of random variables, any finite number of which have a joint Gaussian distribution.*

Their use was first suggested in (Aigrain et al. 2011). Gaussian processes have attractive properties which makes them a useful tool for non-parametric inference. In the RV community, the theoretical property of Gaussian processes, seen as continuous processes, is the formula for the covariance function of a differentiated process. At least to the knowledge of the author, the other uses of Gaussian processes amount to manipulating parametrized multivariate Gaussian noise. Using non independent Gaussian noise has been done for instance in Baluev (2011), without reference to the term Gaussian process.

If one parametrizes the autocorrelation function by a vector $\boldsymbol{\eta}$, this one can be concatenated with orbital and non planetary effect parameters $\boldsymbol{\theta}$ to obtain an extended set of parameters. A popular form for the covariance matrix of the noise $\mathbf{V}$ is (Haywood et al. 2014, equation 4),

$$\mathbf{V}_{kl} = \eta_1^2 \exp\left(-\frac{(t_k - t_l)^2}{2\eta_2^2} - \frac{2\sin^2\left(\frac{\pi(t_k - t_l)}{\eta_3}\right)}{\eta_4^2}\right). \tag{1.43}$$

The idea behind this formula is to consider a pseudo-periodic correlation. The parameter η_1 is the amplitude of the noise, η_2 is a time-scale of correlation. We expect the features on the stellar surface (spots, faculae...) to re-appear every stellar rotation period η_3, and η_4 scales the amplitude of the periodic behaviour. If one wants to add a simpler red noise model to account for granulation (with an exponential kernel for instance), under the hypothesis that this noise is independent from the stellar activity the covariance matrices are simply summed.

The fact that $\boldsymbol{\eta}$ models the noise does not give it a status different from $\boldsymbol{\theta}$. The noise parameters $\boldsymbol{\eta}$ can be straightforwardly included in the analysis presented in the previous section. In figure 1.24, we show the same plot as in 1.23 but with more free parameters. The data comes from two instruments, and we include in the model the offsets of these two ("Mean HARPS" and "Mean Terra"). An additional variance σ^2

$$\mathbf{V}_{kl} = \delta_{kl}\sigma_k^2 + \sigma^2 \tag{1.44}$$

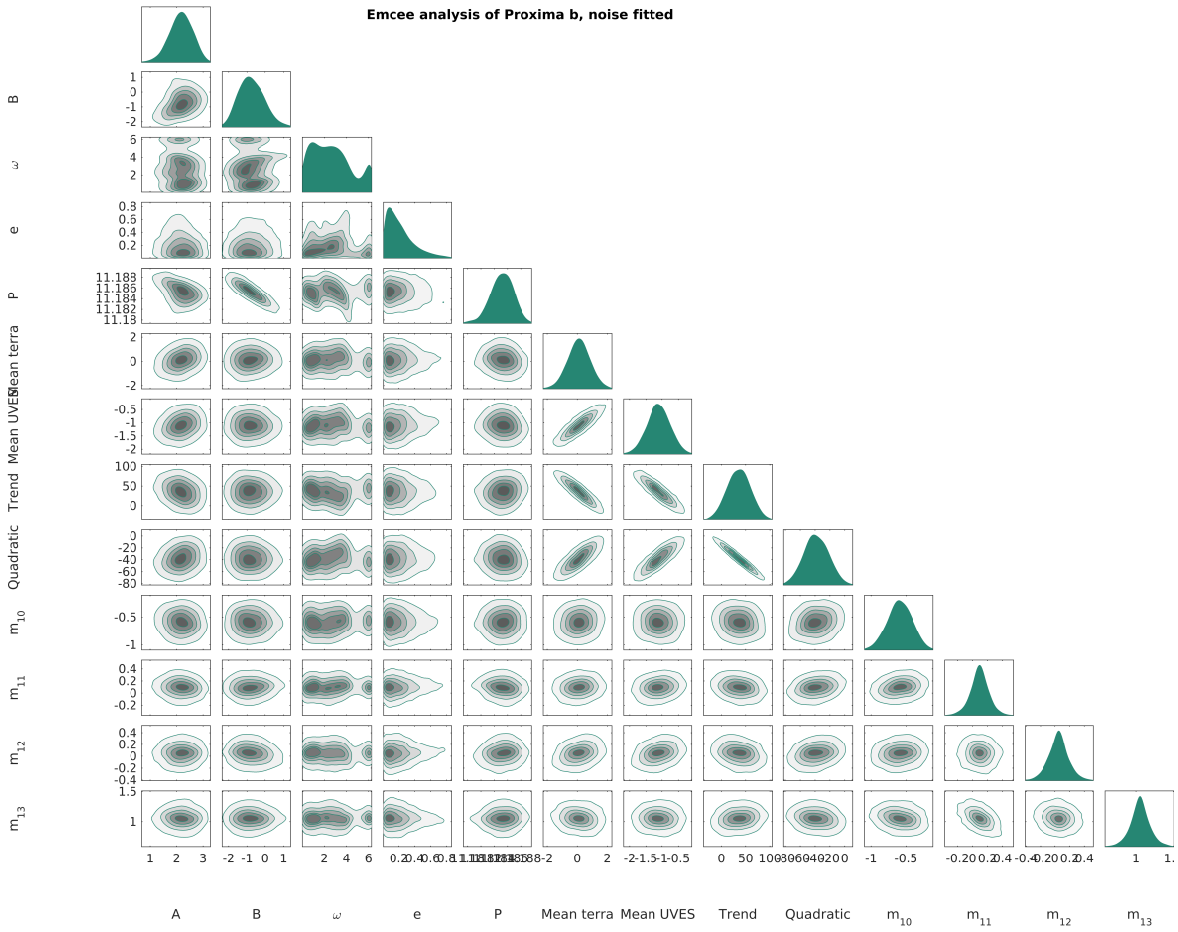

Figure 1.24: Example of MCMC analysis of Proxima b with EMCEE implemented by Grinsted (2015).

where σ_k is the nominal uncertainty on measurement k and δ_{kl} is the Kronecker symbol, equal to one if and only if $k = l$. We choose a narrow prior on σ around 1 m.s^{-1} σ as in the discovery paper (Anglada-Escudé et al. 2016). The posterior distribution of σ^2 is labelled "m_{14}" in figure 1.24.

This example does not include non diagonal terms in the matrix. Such terms would necessitate a matrix inversion, whose complexity scales approximately as N^3, to compute the Likelihood at each sample of the EMCEE code. Faria et al. (2016) uses a diffusive nested sampling algorithm (Brewer 2014). On the 177 public measurements of CoRoT-7, with the covariance function (1.43) plus a free extra variance term (as σ in equation (1.44)). They obtain 50,000 samples in four days [5]. Mitigating the effect of correlated noise can be done in the framework of residual periodograms, instead of random sampling. The procedure consists in directly maximizing the likelihood with the Keplerian model plus the free noise parameters at each trial frequency of the periodogram (Baluev 2011, 2013c). Estimating the noise from

[5] On an Intel® Core™ i5-4460 CPU running at 3.20 GHz and 4 GB of RAM. An interesting feature is that this algorithm treats the number of planet as a free parameter (it is trans-dimensional).

the smoothed residual periodogram is also possible (same references). Another solution is to keep only the diagonal terms of $\mathbf{V}$ and set $\mathbf{V}_{ij}$ equal to zero elsewhere except on a few minor diagonals closest to the main one. This allows to speed up the inversion and to fit a noise parameter (moving average, Tuomi et al. 2013). Finally, one can "train" the Gaussian process on another data set. For instance, Haywood et al. (2014) analyses CoRoT-7 in the following way. They have a RV and a photometric time series. They consider the photometric time series has a zero expectancy, noise parameters $\boldsymbol{\eta}$ and autocorrelation function given by equation (1.43). They fix η_3 as the estimated rotation period of the star and take $\boldsymbol{\eta}^\star$ that maximizes the marginal likelihood of the photometric time series. Local maxima are not expected to be a major issue (MacKay 1999). The covariance function of the radial velocity is then estimated with the FF' method (Aigrain et al. 2012). This one links the photometric variation to the RV variation through a 1 spot model. The covariance for the RV time series so obtained is held fix for the search of planetary companions. Let us finally mention that Rajpaul et al. (2015) goes a step further and models simultaneously the RV, BIS and $\log R'_{HK}$ with a single Gaussian process G and its first order derivative $\dot{G}$.

$$\Delta RV = V_c G(t) + V_r \dot{G}(t); \tag{1.45}$$

$$\log R'_{HK} = L_c G(t) \tag{1.46}$$

$$\mathrm{BIS} = B_c G(t) + B_r \dot{G}(t) \tag{1.47}$$

Using this modelling allows to constraint the parameters of the covariance function, which is one of the challenges of the use of correlated Gaussian noise. Indeed, Gaussian processes can easily absorb a great quantity of signal energy, to the point where planets can be eaten out as well. Gaussian processes are appealing since they seem to provide an account for both stochastic and physical behaviour.

1.4 Summary

In this chapter, our aim was to outline the elements that have to be taken into account when finally reporting a planet and its orbital elements. Besides the fact that a good understanding of the data and of the methodology is key for a reliable inference, it seemed particularly important to us to identify areas of improvements. Indeed, this work had initially no precise purpose but the analysis of data collected to detect and characterize exoplanets. The literature review was an important step to ensure not to reinvent the wheel and not to miss a crucial aspect of the inference process that could endanger our analysis.

Our starting point issued from two questions that emerged during the participation to the NEAT challenge (Anglada-Escude et al. 2014): how to circumvent aliasing and why are eccentricities overestimated? The NEAT challenge consisted in analysing 200 simulated data sets of 45 astrometric measurements to retrieve planets, the goal was to demonstrate the capacity of the community to exploit the data from a potential GAIA successor. Our first approach was an iterative process, where one would compute the periodogram, fit the best candidate and perform the same analysis on the residuals. We found that it often occurred that at some iteration, the tallest peak of the periodogram was spurious due to an unfortunate combination of several planetary signals and the noise. We had the idea of using ℓ_1 minimization methods, that allow to search for several signals simultaneously and

yet avoid brute force approaches. Tailoring and these tools to our purposes and studying their usefulness and reliability is the object of chapter 2, and of (Hara et al. 2017).

In the analysis of the NEAT challenge we encountered another interesting fact, which is that the best fit eccentricities were in general over-estimated. Furthermore, when initializing a non-linear fit of a Keplerian model at a spurious period, it often happened that the eccentricity converged to one. We also found that there is a detection bias against high eccentricity planets at a given semi-amplitude. The eccentricity estimation and detection bias has been studied in more or less depth in Cumming (2004); Shen & Turner (2008); Zakamska et al. (2011); Baluev (2015b). Our goal was to get a precise understanding of the origin of these biases, and more generally of the behaviour of least square fit estimates. The analysis of this bias is done chapter 3. We find that biases on projected mass and inclination (from astrometry) have the same origin. We argue that bias is not so much of a problem in itself, since inference on physical properties are done with interval estimates, be it for a given target star or on eccentricity distributions. We also consider the problem of designing an observation strategy to obtain guarantees on the resolution on the distribution of eccentricities.

This brings us to the third point we wish to highlight. We found that when the model is incorrect, the eccentricity tends to be even more over-estimated. However, in that case, the uncertainties are possibly underestimated. We then undertook to find ways to diminish the impact of a wrong model and to diagnose when models are incorrect. The discussion in section 1.3.2 is part of this goal, and aims at providing theoretical elements to consider when reasoning under wrong models. This discussion, as well as the tools suggested in section 3.4 are part of a work in progress (Hara et al. 2018, in prep.).

Chapter 2

Compressed Sensing/Sparse Recovery and Radial Velocity data

Je ne cherche pas, je trouve.

Pablo Picasso

The choice of a data analysis method includes a trade-off between accuracy and speed. In the analysis of radial velocity techniques, provided the errors are appropriately computed, improved robustness is obtained by several means: by searching directly for Keplerian signals instead of purely sinusoidal ones, combining RV measurements with other ones (e.g. transits), trying several noise models, possibly based on activity indicators derived from the spectra and finally, searching for several planets at once. All these improvements take additional computation time.

In any statistical analysis technique, the likelihood must be computed over the space of parameters. So we will have to find the maximum of a function, or use it in an integral to compute the posterior likelihood. In general, the time required to compute those quantities evolves exponentially with the number of parameters, a problem sometimes called the "curse of dimensionality". Where does the exponential increase comes from? If we could just evaluate the function but did not know anything about it, we would have to try every combination of parameters on a grid as dense as possible to find the maximum or compute an integral. If this was systematic, since every additional planet brings five parameters it would discourage us to attempt to fit many at once. However, knowing something about the function can drastically reduce research time. For instance, if the function is convex and the space of parameters is convex as well, finding the minimum of the function is drastically simplified: Newton descents will do.

What property of our problem could be used to speed up the search of signals? In full generality, this question remains open. In this chapter, we will exploit the assumption that the signal is sparse, which means that the number of planets is "much less" than the number of measurements. In section 2.1 this claim is precised and put in the broad context of sparse signals analysis. The goal of this section is to provide an overview of motivations and theoretical aspects of sparse recovery, which might be unfamiliar to the exoplanet community. In section 2.2 we present the core of this thesis, that is the application of this theory to radial velocity measurements as done in Hara et al. (2017), which is reproduced in appendix G.

Symbol	description
$\langle \boldsymbol{a}, \boldsymbol{b} \rangle_w$	Euclidian or Hermitian product respectively if a and b are real or complex valued. $\langle \boldsymbol{a}, \boldsymbol{b} \rangle_w = \sum_{k=1}^{n} w_k a_k^* b_k$ where the weights w_k are positive real numbers.
z^*	Conjugate of $z \in \mathbb{C}$
$\|\boldsymbol{a}\|_w$	$\|\boldsymbol{a}\|_w = \sqrt{\langle \boldsymbol{a}, \boldsymbol{a} \rangle_w}$
$\|\boldsymbol{x}\|_{\ell_0}$	ℓ_0-norm of vector $x \in \mathbb{C}^n$: $\|x\|_{\ell_0}$ is the number of non-zero components of x
$\|\boldsymbol{x}\|_{\ell_p}$, $p > 0$	ℓ_p-norm of vector $x \in \mathbb{C}^n$: $\|x\|_{\ell_1} = \left(\sum_{k=1}^{n} |x_k|^p \right)^{\frac{1}{p}}$
$\boldsymbol{x}^\star = \arg\min_{\boldsymbol{x} \in E} F(\boldsymbol{x})$	Value of $x \in E$ that minimizes the functional $F : E \to \mathbb{R}_+$
$\boldsymbol{x}_t$	True value of a variable x. In any case the subscript t holds for "true".
$\boldsymbol{y}$	Vector of observations.
$\mathbf{A}$	Dictionary: a $m \times n$ matrix, $m < n$ with non vanishing columns
$\boldsymbol{a}_k$	k-th column of $\mathbf{A}$
$\mathbf{V}$	correlation matrix of the noise
$\mathbf{W}$	$\mathbf{W} = \mathbf{V}^{-\frac{1}{2}}$ (exists as V is symmetric)
Ω	Discrete subset of $\mathbb{R}$
$\ker\mathbf{A}$	Null space of $\mathbf{A}$, $\ker\mathbf{A} = \{\boldsymbol{h} \in \mathbb{R}^n, \mathbf{A}\boldsymbol{h} = 0\}$
$f(\boldsymbol{t})$, $f : \mathbb{R} \to \mathbb{R}$	When $f : \mathbb{R} \to \mathbb{R}$ and $\boldsymbol{t}$ is a vector of M components, $f(\boldsymbol{t}) = (f(t_k))_{k=1..M}$ is the vector made of evaluations of f on each t_k. In particular $\mathrm{e}^{\mathrm{i}\omega\boldsymbol{t}} = (\mathrm{e}^{\mathrm{i}\omega t_k})_{k=1..m}$.
$p(x\|y)$	Probability density of random variable X knowing the value of random variable Y

Table 2.1: List of symbols

In section 2.3, possible improvements are discussed. Section 2.4 presents more in-depth the Radial Velocity Challenge, for which the most recent version of the tool is particularly efficient.

2.1 Sparsity, Compressed Sensing and ℓ_1 norms

2.1.1 Reformulation

The model for ideal radial velocity signals, given by Equation 1.20, is reproduced here

$$y(t, \boldsymbol{\theta}) = V + \sum_{i=1}^{n_p} K_i(\cos(\omega_i + \nu(t, e_i, P_i, \omega_i, M_{0,i})) + e_i \cos\omega_i).$$

When the orbits are circular, this becomes a simple sum of sines

$$y(t, \boldsymbol{\theta}) = V + \sum_{i=1}^{n_p} B_i \cos\frac{2\pi}{P_i}t + C_i \sin\frac{2\pi}{P_i}t.$$

The decomposition of the radial velocity signal as a sum of sinusoidal signals is in fact more general than the circular case. Indeed, since Keplerian signals are periodic, they can be expanded in Fourier series and even more generally, stable planetary systems can be approximated by quasi-periodic series (Laskar 1988; Laskar et al. 1992; Laskar 1993, 2003).

We will now use the notations of Hara et al. (2017) and denote by $\omega = 2\pi/P$ the mean motions and m the number of measurements. Let us denote by $y(t_i)$ the velocity evaluated at different instants $\boldsymbol{t} = (t_k)_{k=1..m}$, then

$$\begin{pmatrix} y(t_1) \\ \vdots \\ y(t_m) \end{pmatrix} = \begin{pmatrix} 1 & \cos(\omega_1 t_1) & \dots & \cos(\omega_{n_p} t_1) & \sin(\omega_1 t_1) & \dots & \sin(\omega_{n_p} t_1) \\ \vdots & \vdots & & & & & \vdots \\ 1 & \cos(\omega_1 t_m) & \dots & \cos(\omega_{n_p} t_m) & \sin(\omega_1 t_m) & \dots & \sin(\omega_{n_p} t_m) \end{pmatrix} \begin{pmatrix} V \\ B_1 \\ \vdots \\ B_{n_p} \\ C_1 \\ \vdots \\ C_{n_p} \end{pmatrix}. \tag{2.1}$$

When searching for the periodicities in the data, we obviously do not know where the $(\omega_i)_{i=1..n_p}$ are. An upper bound can be put though, since planets should not be observed below the Roche limit. The frequencies are to be searched in an interval $[0, \omega_{\max}]$ which we can first think of informally as a discrete grid $\Omega = (\omega_j)_{j=1..n}$ arbitrarily dense such that $\omega_1 = 0$, $\omega_n = \omega_{\max}$ and $\omega_{j+1} = \omega_j + \Delta\omega$. The situation we are in is rather that we have an additive noise $\boldsymbol{\epsilon} = (\epsilon_i)_{i=1..m}$ and an equation

$$\begin{pmatrix} y(t_1) \\ \vdots \\ y(t_m) \end{pmatrix} = \begin{pmatrix} \cos(\omega_1 t_1) & \dots & \cos(\omega_n t_1) & \sin(\omega_1 t_1) & \dots & \sin(\omega_n t_1) \\ \vdots & & & & & \vdots \\ \cos(\omega_1 t_m) & \dots & \cos(\omega_n t_m) & \sin(\omega_1 t_m) & \dots & \sin(\omega_n t_m) \end{pmatrix} \begin{pmatrix} B_1 \\ \vdots \\ B_n \\ C_1 \\ \vdots \\ C_n \end{pmatrix} + \begin{pmatrix} \epsilon_1 \\ \vdots \\ \epsilon_m \end{pmatrix}, \tag{2.2}$$

a solution of which we wish to determine. This is an inverse problem, but since the matrix is such that $m < n$, this equation is under-determined. Inversion might seem hopeless, but we have a key information: there are only a few planets. We want the solution not to have more than $\lesssim 2n_p + h$ non null coefficients, where n_p is the number of planets (times two for one sine and one cosine coefficient), and h is the number of additional harmonics. This one could originate from stellar rotation or eccentricities. Overall, the number of non zero coefficients we expect is $\lesssim 20$. In fact, equation (2.2) is a particular instance of a more generic problem

$$\begin{aligned} &\boldsymbol{y} = \mathbf{A}\boldsymbol{x} + \boldsymbol{\epsilon} \\ &\mathbf{A} \text{ is a } m \times n \text{ matrix and } n \gg m \\ &\boldsymbol{x} \text{ has a few non-zero components } (\boldsymbol{x} \text{ is sparse }) \end{aligned} \tag{2.3}$$

The study of sparsity in a broad sense became a field of its own in ≈ 1990 (Donoho & Stark 1989; Mallat & Zhang 1993; Tibshirani 1994; Chen et al. 1995). Let us note that searching for the simplest explanation in some sense is no new thing. One could argue that the choice of a scientific explanation is always a compromise between its simplicity and its precision.

Moreover, sparse solutions where used as soon as Prony (1795). In the modern literature, the matrix $\mathbf{A}$ is termed the "dictionary" and its columns are termed "atoms". We are basically searching for a linear combination of a few atoms of $\mathbf{A}$ that reproduces the data.

In this form, the problem is not precisely defined. A natural way to search for a sparse solutions is to look for vectors with a few non-zero entries whose distance to the data is bounded,

$$\underset{\boldsymbol{x}\in\mathbb{R}^n}{\arg\min} \quad \|\boldsymbol{x}\|_{\ell_0} \quad \text{s.t.} \quad \|\mathbf{A}\boldsymbol{x}-\boldsymbol{y}\|_{\ell_2} < \epsilon \tag{$P_{\ell_0,\epsilon}$}$$

where $\|\boldsymbol{x}\|_{\ell_0}$ is the number of non zero elements of a vector $\boldsymbol{x}$ and is termed the ℓ_0 norm, even though it is not a norm in a strict sense. Formula ($P_{\ell_0,\epsilon}$) reads: among the set of vectors $\boldsymbol{x}$ that satisfy $\|\mathbf{A}\boldsymbol{x}-\boldsymbol{y}\|_{\ell_2} < \epsilon$, we search those that have the minimal number of non-zero components. The notation $\arg\min f(x)$ is an abbreviation for "argument that minimizes the function $f(x)$" and s.t. stands for "subject to". Table 2.1 summarizes the notations used. Let us first examine the noiseless case, in which the problem becomes

$$\underset{\boldsymbol{x}\in\mathbb{R}^n}{\arg\min} \quad \|\boldsymbol{x}\|_{\ell_0} \quad \text{s.t.} \quad \mathbf{A}\boldsymbol{x}=\boldsymbol{y}. \tag{P_{ℓ_0}}$$

Let us denote by $\boldsymbol{x}_t \in \mathbb{R}^n$ the true coefficients of the signal, that is

$$\boldsymbol{y} = \mathbf{A}\boldsymbol{x}_t. \tag{2.4}$$

Do we have any chance to recover $\boldsymbol{x}_t$ by solving (P_{ℓ_0})? The success will depend on matrix $\mathbf{A}$, and more precisely on its spark, defined as the minimum number of linearly dependent columns. Denoting by $\boldsymbol{a}_j$ the j-th column of $\mathbf{A}$, J a subset of indices and $|J|$ the cardinal of J, $\text{spark}(\mathbf{A}) = \min\{k, (\boldsymbol{a}_j)_{j\in J}, |J| = k, \exists(\lambda_j)_{j\in J}, \sum_j \lambda_j \boldsymbol{a}_j = \mathbf{0}\}$. Then one obtains (this is a very slightly different formulation of theorem 1 in Kutyniok 2012)

Theorem 1. *The following propositions are equivalent*

(i) $\forall \boldsymbol{x}_t \in \mathbb{R}^n$ *such that* $\|\boldsymbol{x}_t\|_{\ell_0} \leqslant S$, $\boldsymbol{x}_t$ *is the unique solution of* (P_{ℓ_0}) *for* $\boldsymbol{y} = \mathbf{A}\boldsymbol{x}_t$.

(ii) $S < \frac{\text{spark}(\mathbf{A})}{2}$.

This theorem gives a condition which guarantees that the sparsest solution is the true one, but it does not tell us if (P_{ℓ_0}) is computationally tractable. Unfortunately, when no hypotheses are made on the dictionary $\mathbf{A}$, this is a NP-hard problem (Ge et al. 2011). As it is, this reformulation does not allow to bypass the thorough exploration of S-sparse signals.

2.1.2 Greedy algorithms

The result of Ge et al. (2011) tells us that no fast algorithm will solve (P_{ℓ_0}) for all dictionary $\mathbf{A}$. However the problem (2.4) is always solved in a specific context: the dictionary is known and its specific properties can be used to speed up algorithms. Astronomers in particular should not be too pessimistic, since there are cases where fast, iterative algorithms are able to retrieve efficiently harmonics when those are not too numerous. After all, we are doing nothing but fitting epicyclic motions to observations, just like Ptolemy. Although the interpretation was wrong, the frequency retrieved by his analysis were at least not always spurious. The

algorithms known to perform well — though not perfectly — are CLEAN (Roberts et al. 1987), frequency analysis (Laskar 1988, 1993, 2003) and CLEANest (Foster 1995).

In all these algorithms the dictionary is made of sine functions (complex or real) and they function on the same principle: the frequency corresponding to the maximum correlation is subtracted from the signal, then the search is performed on the residuals and so on. In the case of frequency analysis and CLEANest, at each iteration, all the atoms found so far are re-fitted, possibly in a non-linear way with re-adjustment of the frequencies, not only the coefficients. In CLEAN, only the last element of the dictionary found is adjusted and subtracted from the residuals.

Interestingly enough, these methods have a counterpart in the Statistical and Compressed Sensing literatures. The principle of CLEAN is exactly the one of Matching Pursuit (Mallat & Zhang 1993), CLEANest and frequency analysis are close to Orthogonal Matching Pursuit (Pati et al. 1993). These methods are known as "greedy algorithms" in the sparse recovery literature. Beforehand, adding features step by step was known in the statistical literature as forward stepwise regression (Efroymson 1960), and matching pursuit is a particular case of projection pursuit (Huber 1985). The method used to find exoplanets in RV data has the same structure: one computes the periodogram with a dictionary made of sinusoidal or Keplerian functions, then the signal that fits data the best is subtracted from the data, and a periodogram is computed on the residuals.

Unfortunately, all these methods share the same problem: at some iteration the vector of the dictionary that has the maximum correlation with data might be spurious[1]. We met such a case several times in section 1.3.3 : the tallest peak of the periodogram might be at a spurious frequency around a period of one day or at a period which could pass as one of an exoplanet. We are now going to try to understand why this happens.

The theory of compressed sensing gives conditions on the dictionary that guarantee success of greedy algorithms. These ones are expressed via the notion of mutual coherence, which is the maximum correlation between two columns of $\mathbf{A}$,

Definition 1 (Mutual coherence (Donoho & Elad 2002))**.** *The mutual coherence μ of a real or complex matrix $\mathbf{A}$ is defined as*

$$\mu(\mathbf{A}) = \max_{k,l\in\{1..n\},k\neq l} \frac{|\langle \boldsymbol{a}_k, \boldsymbol{a}_l\rangle|}{\|\boldsymbol{a}_k\|_{\ell_2}\|\boldsymbol{a}_l\|_{\ell_2}}. \tag{2.5}$$

This quantity is in fact familiar to astronomers. When the dictionary is made of complex sines, $\boldsymbol{a}_k = \mathrm{e}^{\mathrm{i}\,\omega_k \boldsymbol{t}} = (\mathrm{e}^{\mathrm{i}\,\omega_k t_i})_{i=1..m}$, so the right hand side of the equation comes down to the spectral window, which we denote by S_w.

$$\frac{|\langle \boldsymbol{a}_k, \boldsymbol{a}_l\rangle|}{\|\boldsymbol{a}_k\|_{\ell_2}\|\boldsymbol{a}_l\|_{\ell_2}} = \frac{1}{m}\left|\sum_{i=1}^{m} \mathrm{e}^{\mathrm{i}(\omega_l-\omega_k)t_i}\right| =: |S_w(\omega_l-\omega_k)|. \tag{2.6}$$

It is a remarkable property of dictionaries made of complex exponentials that the correlation is somewhat "stationary", that is the correlation between the columns k and l of the dictionary only depends on $|k-l|$.

[1] Although when all frequencies found at a given iteration are re-fitted, the problem is mitigated. Foster (1995) shows an example where the first peak selected is wrong but at the end of the iteratios the estimated amplitude at the corresponding frequency is much smaller than the one at the true frequencies.

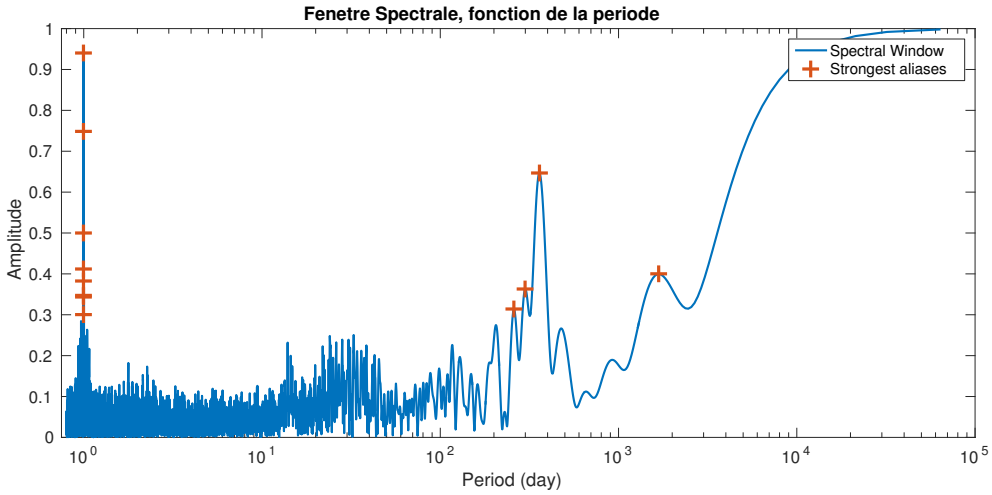

Figure 2.1: Spectral window of HD 10180

A sufficient condition of success of orthogonal matching pursuit can be expressed with the mutual coherence (Donoho & Huo 2001; Tropp 2004)

Theorem 2. *If* $\boldsymbol{x}_t \in \mathbb{R}^n \backslash \{0\}$ *verifies*

$$\|\boldsymbol{x}_t\|_{\ell_0} < \frac{1}{2}\left(1 + \frac{1}{\mu(\mathbf{A})}\right) \tag{2.7}$$

then it is the unique solution of (P_{ℓ_0}) *with* $\boldsymbol{y} = \mathbf{A}\boldsymbol{x}_t$ *and Orthogonal Matching Pursuit with error threshold 0 recovers it.*

It means that iterative searches will succeed if the columns of the dictionary are not too correlated two by two. Unfortunately, equating the right hand side of the inequality to two yields the condition $\mu(\mathbf{A}) < 1/3$. As soon as the mutual coherence is greater than this threshold, we are not sure that the Orthogonal Matching Pursuit will be able to retrieve correctly two components, even if there is no noise.

The condition of the theorem is never fulfilled for exoplanets searches in radial velocity. Indeed, if we take two close frequencies on a dense grid, ω_k and ω_{k+1}, equation (2.6) gives the correlation between the columns, $S_w(\omega_{k+1} - \omega_k)$, which tends to one as the grid gets finer. In simpler terms, neighbouring frequencies on the grid give almost identical models. Furthermore, due to the daily repetition of measurements and seasonality of observations, spectral windows often have spikes $\gtrsim 0.8$ and $\gtrsim 0.5$ respectively at one day and one year. As an example, we show the spectral window of HD 10180 measurements (Lovis et al. 2011b) in figure 2.1, that has ten peaks above 0.3 despite its relatively high number of measurements (190). Six of the eleven strongest aliases are packed around a period of 1 day, with a peak reaching 0.94. Even when disregarding close frequencies, the condition of the theorem is not fulfilled.

Since the condition is a sufficient one, its tightness can be questioned. Let us do a numerical experiment. We take the measurements times of HD 10180, $\boldsymbol{t} = (t_k)_{k=1..m}$ and generate

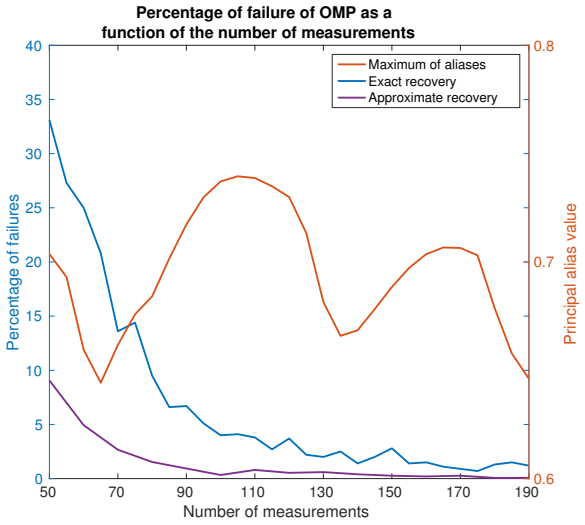

Figure 2.2: Percentage of failure to retrieve correct frequencies via Orthogonal Matching Pursuit as a function of the number of measurements.

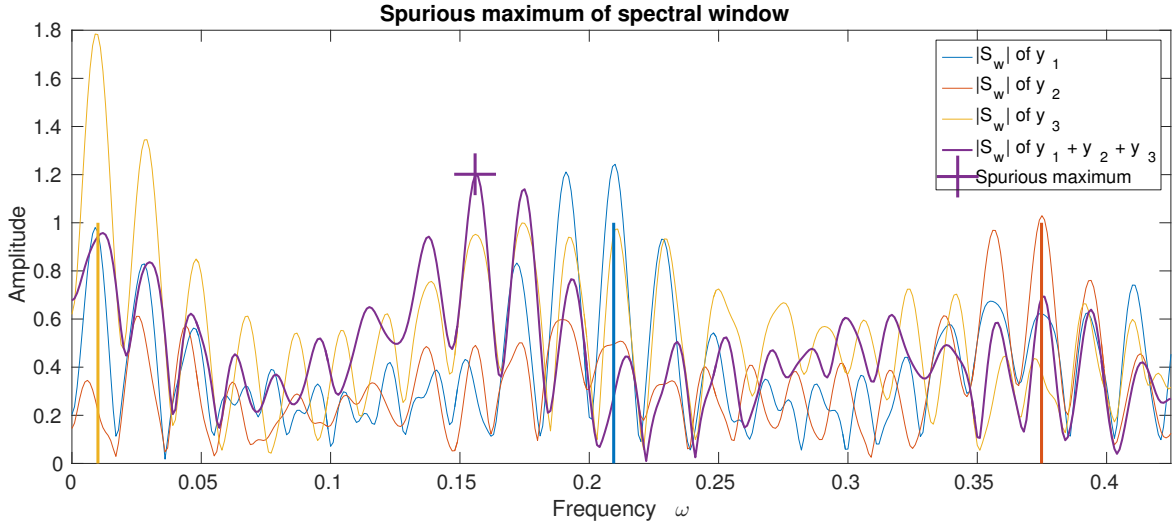

Figure 2.3: Case of spurious maximum of the spectral window. The coloured stem show the position of correct frequencies. The plain lines represent the absolute value of the correlation between $\boldsymbol{y}_i$, $i = 1, 2, 3$ or $\boldsymbol{y}$ and $\mathrm{e}^{\mathrm{i}\omega \boldsymbol{t}}$ as a function of ω, which is $|S_w(\omega_i + \omega) + S_w(\omega_i - \omega)|$ in the first three case (blue, red, yellow curves) and $\left|\sum_{i=1}^{3} S_w(\omega_i + \omega) + S_w(\omega_i - \omega)\right|$ for the purple curve (see equation (2.6)).

signals of type

$$\boldsymbol{y} = \sum_{k=1}^{n_p} c_k \, \mathrm{e}^{\mathrm{i}\,\omega_k \boldsymbol{t}} \tag{2.8}$$

with frequencies ω_k chosen randomly with uniform probability on the span 0 to 1.9π. First we take $n_p = 2$, $c_1 = c_2 = 1$ and generate 1000 realizations of signal (without noise). On each

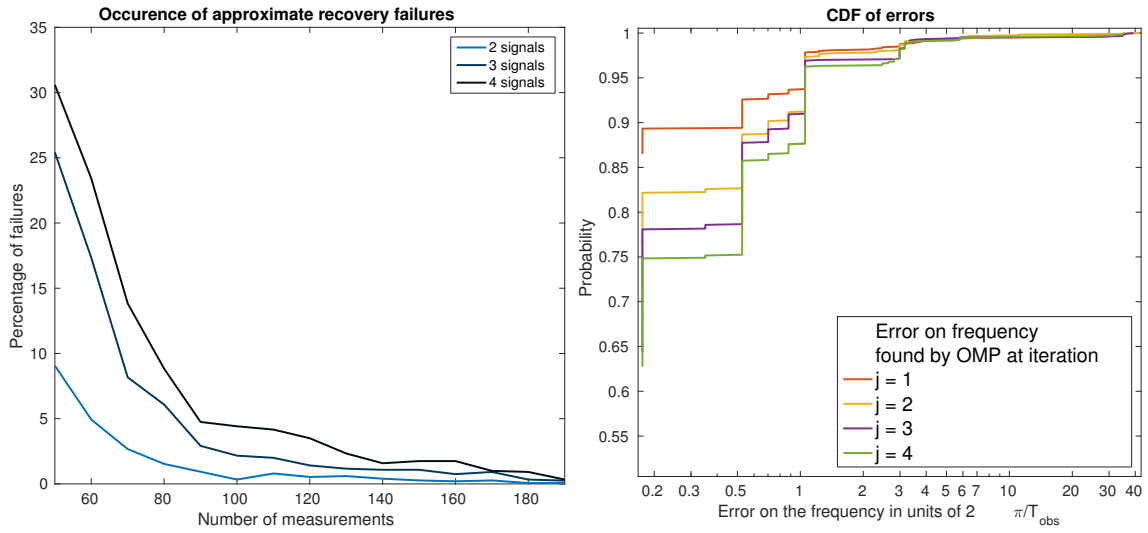

Figure 2.4: Cumulative distribution function (CDF) of errors on frequency found by orthogonal matching pursuit.

of those, an orthogonal matching pursuit algorithm is performed. We compute the number of failures to retrieve the correct frequencies divided by the number of simulations.

This method is performed for signals generated on the 50 first measurement dates of HD 10180, then the 55 first, 60 and so on, up to 190. The results are presented figure 2.2, where it is clear that though the value of the maximum alias of the periodogram stays relatively steady, the rate of failure decreases with the number of measurements. When failure is taken as a non exact retrieval of frequencies, the blue curve is obtained. In practice, this is too stringent a criterion, since the physical properties of planets will remain similar in a neighbourhood of the frequencies. If we only take as a criterion that one of the two true frequencies is not retrieved within $2\pi/T_{\mathrm{obs},m}$, where $T_{\mathrm{obs},m}$ is the observation timespan for m measurements, the purple curve is obtained. $2\pi/T_{\mathrm{obs},m}$ is a typical scale of local minima in frequency, hence its choice. Note that the smallest is the distance between subsequent frequencies on the grid Ω, the greater will be the rate of failure of exact recovery for reasons which will become clearer afterwards.

The method fails when at a certain iteration, the correlation of the residuals with a column of the dictionary is higher than a correct one. Since there is no noise, failures appear only because of correlations between columns. In the case of pure frequencies with no noise, the signal is given by equation (2.8), so the correlation with a column of the dictionary $\mathrm{e}^{\mathrm{i}\omega \boldsymbol{t}}$ is

$$|\langle \mathrm{e}^{\mathrm{i}\omega \boldsymbol{t}}, \boldsymbol{y}\rangle| = \left|\sum_{k=1}^{n_p} c_k \langle \mathrm{e}^{\mathrm{i}\omega \boldsymbol{t}}, \mathrm{e}^{\mathrm{i}\omega_k \boldsymbol{t}}\rangle\right| = \left|\sum_{k=1}^{n_p} c_k S_w(\omega_k - \omega)\right| \tag{2.9}$$

It means in particular that the correlation with a vector that truly is in the signal, $\mathrm{e}^{\mathrm{i}\omega_j \boldsymbol{t}}$ is

$$|\langle \mathrm{e}^{\mathrm{i}\omega_j \boldsymbol{t}}, \boldsymbol{y}\rangle| = \left|c_j + \sum_{k=1, k\neq j}^{n_p} c_k S_w(\omega_k - \omega_j)\right|. \tag{2.10}$$

It might turn out that a frequency ω corresponds to an alias of several true signals, that is $|S_w(\omega_k - \omega)|$ is high for several k. If the phases of $S_w(\omega_k - \omega)$ add coherently, one will observe a high correlation $|\langle \mathrm{e}^{\mathrm{i}\,\omega \boldsymbol{t}}, \boldsymbol{y}\rangle|$ that does not correspond to a true signal. Figure 2.3 shows an example of such a situation. We consider the thirty first measurements of HD10180 and generate signals with periods 30, 16.71 and 624 days, $\boldsymbol{y}_1 = \cos(2\pi/30\boldsymbol{t})$, $\boldsymbol{y}_2 = \cos(2\pi/16.71\boldsymbol{t})$, $\boldsymbol{y}_2 = -\cos(2\pi/624\boldsymbol{t})$ and consider $\boldsymbol{y} = \boldsymbol{y}_1 + \boldsymbol{y}_2 + \boldsymbol{y}_3$. Note that $\cos x = (\mathrm{e}^{ix} + \mathrm{e}^{-\mathrm{i}\,x})/2$, so that we are in the situation of equation (2.9). We represent $|\langle \mathrm{e}^{\mathrm{i}\,\omega \boldsymbol{t}}, \boldsymbol{y}_i\rangle|$ as a function of ω for $i = 1, 2, 3$ in respectively blue, red and yellow. Also we plot in purple $|\langle \mathrm{e}^{\mathrm{i}\,\omega \boldsymbol{t}}, \boldsymbol{y}\rangle|$, which is what any of the algorithm mentioned above would do at first iteration. Unfortunately, the maximum occurs at a spurious frequency, the algorithm would fail from the very first iteration. This example was designed to be pathological, but this situation does happen in a non negligible fraction of cases (see next paragraph and section 2.3.2).

Furthermore, equation (2.10) shows that the Fourier transform at ω_j is not exactly $|c_j|$. Other sources might had coherently and boost or cut off the signal at ω_j. It might result in a small shift in frequency of the maximum of correlation. Then even if the highest correlation happens not at a completely spurious peak but close to a true frequency ω_j, one does not subtract exactly $c_j\,\mathrm{e}^{\mathrm{i}\,\omega_j \boldsymbol{t}}$ but $\tilde{c_j}\,\mathrm{e}^{\mathrm{i}\,\tilde{\omega_j}\boldsymbol{t}}$ with some $\tilde{c_j}, \tilde{\omega_j} \neq c_j, \omega_j$. This distorts the residuals, the error propagates and amplifies in subsequent iterations. We show figure 2.4 (left) how the recovery failures evolve with the number of signals injected and the number of measurements. It is clear that the precision worsens with the number of signals injected. We do another experiment, which consists in injecting 4 random signals of frequency $\omega_1, \omega_2, \omega_3, \omega_4$. We calculate $e_j = \min_{i=1,2,3,4} |\omega^{(j)} - \omega_i|$ where $\omega^{(j)}$ is the j-th frequency found by the iterative process. We generate such four signals with uniform distribution in frequency 1200 times so that we have four series of 1200 errors. For each of them we plot the experimental cumulative distribution function, that is the fraction of signals that have an error below the level given on the x axis. We do not plot the distributions of e_j but of $e_j T_{\mathrm{obs}}/2\pi$ where T_{obs} is the observation timespan. Based on figure 2.3, the local minima are separated by ≈ 0.02 rad/s, that are $\approx 3 \times 2\pi/T_{\mathrm{obs}}$. The representation adopted therefore allows to count in terms of "how many local minima off" the estimate is. The conclusion is clear: as the iterations proceed, the error on the retrieved frequency spreads out.

This discussion points out that aliasing can create spurious peaks and error propagates with iteration of the algorithm. However, when looking at the numbers of figure 2.4 the situation does not seem too concerning, for the four iterations the error on the frequency is below $4 \times 2\pi/T_{\mathrm{obs}}$. For a period of 10 days this means we will measure between 9.83 and 10.16 days, which is not too bad. There seem to be only $\approx 2\%$ of the signals where the period retrieved is completely off the track.

More realistic setting

Let us now try to quantify how frequently we should expect to run in the situation of figure 2.3. To do so, we make simulations that have the following structure. Measurement times $(t_i)_{i=1..m}$ are taken from sets of observations. We then generate n_p sine signals (or circular planets) with amplitudes $(a_k)_{k=1..n_p}$, uniform random phases $(\phi_k)_{k=1..n_p}$, random frequencies $(\omega_k)_{k=1..n_p}$

and a noise $\boldsymbol{\epsilon}$ of covariance matrix $\mathbf{V}$, that is

$$y(t_i) = \sum_{k=1}^{n_p} a_k \cos(\omega_k t_i + \phi_k) + \epsilon_i. \tag{2.11}$$

We generate N such signals, and for each of them we perform an iterative search. The Generalized Lomb-Scargle periodogram (GLS) is computed, then the frequency corresponding to the maximum is fitted non linearly along with a constant vector and subtracted. The GLS is computed on the residuals, the frequency corresponding to the maximum is fitted with the first frequency found plus a constant, and this model is subtracted. This is basically algorithm orthogonal matching pursuit or frequency analysis with GLS instead of a Fourier transform and a non-linear fit of all frequencies found at each iteration (the non linear fit is included in the frequency analysis). The process is repeated until we find a maximum of the periodogram that is below a certain false alarm probability threshold α, computed by simulation.

At iteration j, the algorithm will select a frequency $(\widehat{\omega}_j)$. We then compute $d_j = \min_{k=1..n_p} |\widehat{\omega}_j - \omega_k|$, that is the smallest distance between $\widehat{\omega}_j$ and a true frequency. The algorithm stops at iteration j_f. If $j_f < n_p$, we missed a signal. By convention, we then take $d_j = -1$ for $j = j_f + 1..n_p$. If $j_f > n_p$, spurious signals have been selected, and we pose $d_j = -2$.

We do such simulations with three levels of noise. In each case, there are three sine functions with amplitudes 3.51, 2.66 and 2.2 m.s^{-1} evaluated on the 50 first measurements of GJ 876 (Correia et al. 2010). The periods are generated uniformly in $\log P$ with a criterion avoiding too close mean motions. The noise level taken are 0.1, 1 and 2 m.s^{-1}. The cumulative distribution function of the error on the mean motions in rad/day are shown respectively in figures 2.5, 2.6 and 2.7.

The blue, red and yellow curves show the distribution of errors at iteration 1, 2 and 3 respectively. Their value at the x axis label "missed" indicates that the corresponding iteration did not detect any signal (it was below the detection threshold). For instance in figure 2.7, there are $\approx 30\%$ of missed detections at the third iteration (yellow curve). The purple and green curves represent the distributions of fourth and fifth iteration. Their value at the x axis label "False" shows the proportion of cases where a signal was spuriously detected. In figure 2.7, this amounts to $\approx 45\%$ of false detections at fourth iteration.

2.1.3 ℓ_1 penalties

The numerical experiment done in previous section shows that without noise, with 50 measurements and without looking to periods close to one day, still in 10% of the cases a spurious peak was found. This will worsen with noise and as additional planets are added. As more measurements are taken, the components of the signal get more and more decorrelated which diminishes the probabilities of failure, but measurements are costly. Tackling this problem was in fact the initial motivation of this work.

We would like a method that searches for several frequencies at a time, but as we saw in section 1.3.3, the task is made complicated by the numerous local minima of the phase space (see for instance figure 1.19).

A class of algorithm, based on ℓ_1 norm penalties on the vector to be recovered will be useful. The basic idea is to replace the ℓ_0 norm by an ℓ_1 one in problem (P_{ℓ_0}). From the computational

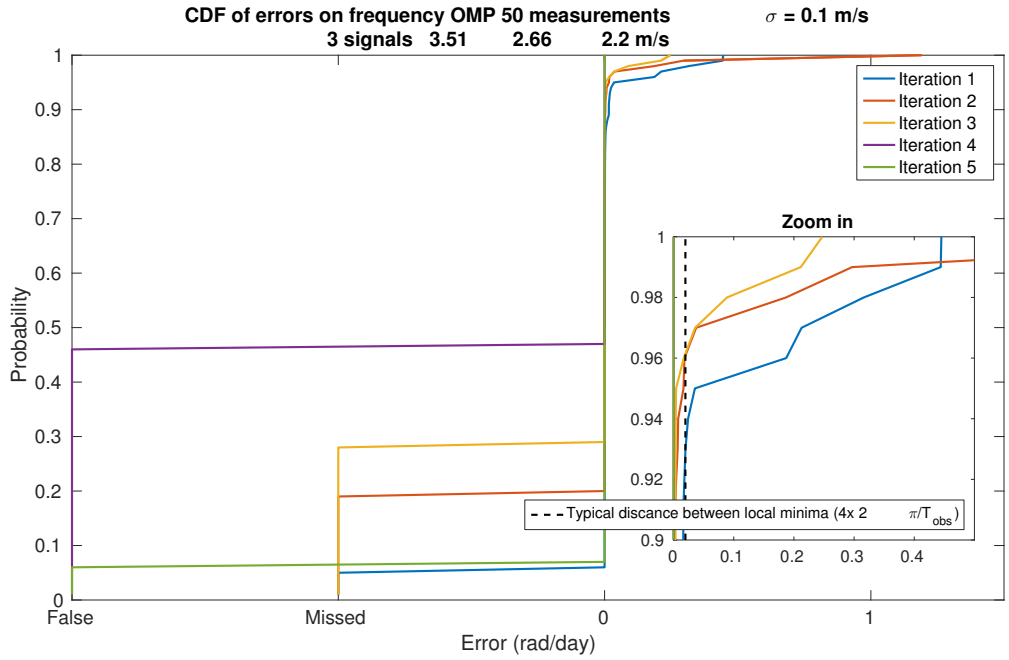

Figure 2.5: Cumulative distribution function of an iterative method, high signal-to-noise ratio.

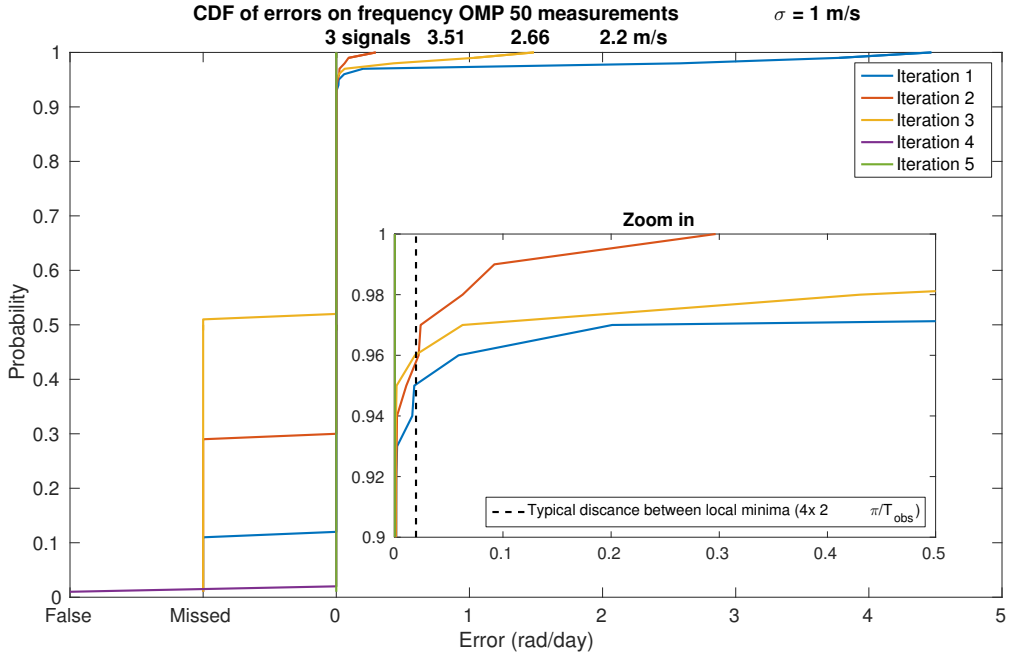

Figure 2.6: Cumulative distribution function of an iterative method, intermediate signal-to-noise ratio.

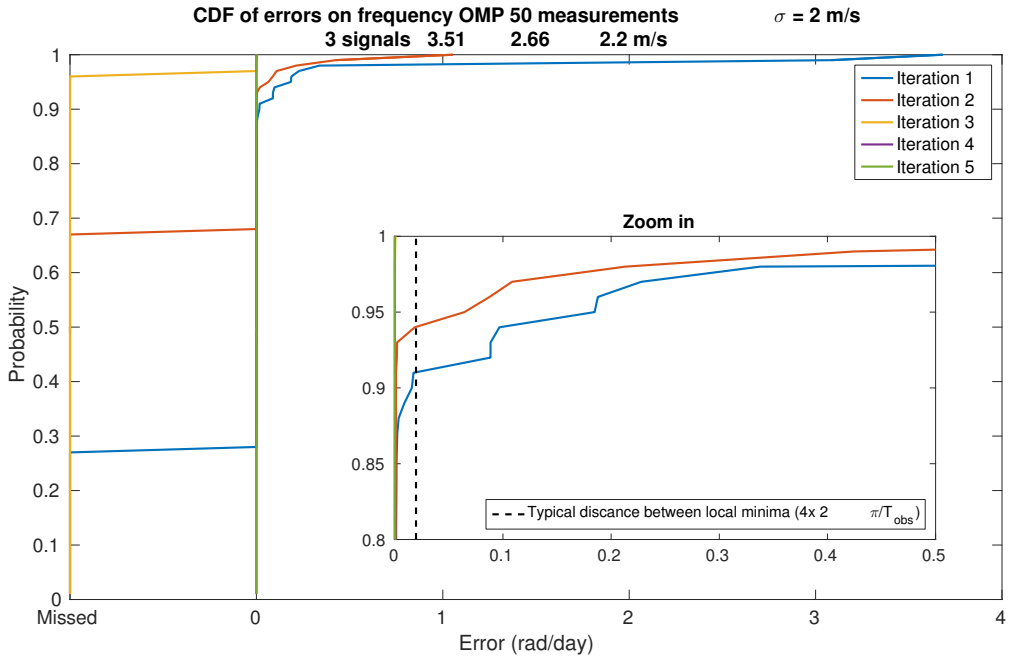

Figure 2.7: Cumulative distribution function of an iterative method, low signal-to-noise ratio.

point of view, the problem becomes a minimization of a convex function over a convex space and becomes much more attractive. This idea was first proposed in Tibshirani (1994); Chen et al. (1995) with different formulations equivalent in some sense, under respective names of Least Absolute Shrinkage and Selection Operator (LASSO) and Basis Pursuit. The definition of the latter is obtained by replacing the ℓ_0 norm by the ℓ_1 one in (P_{ℓ_0}).

$$\underset{\boldsymbol{x}\in\mathbb{R}^n}{\arg\min} \quad \|\boldsymbol{x}\|_{\ell_1} \quad \text{s. t.} \quad \mathbf{A}\boldsymbol{x} = \boldsymbol{y} \tag{P_{ℓ_1}}$$

Changing the norm might seem artificial, but the geometry of the ℓ_1 norm indeed enhances sparsity. To see that, let us consider a simple instance of the problem,

$$\arg\min_{x\in\mathbb{R}^2} \|\boldsymbol{x}\|_{\ell_p} \quad s.t. \quad 1 = \begin{pmatrix} a & b \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}.$$

The right hand side of the equation has an infinity of solutions that belong to the affine space defined by $1 - ax_1 - bx_2 = 0$. Among them, we select the one with the smallest ℓ_p norm. Figure 2.8 shows the solution of this problem for $p = 1$ and $p = 2$ in two cases, $(a, b) = (0.3, 1)$ (left) and $(a, b) = (1, 0.3)$ (right). The ℓ_1 norm selects one coefficient and sets the other to zero (black markers) while the ℓ_2 norm selects two (red markers).

Let us think of the space generated by $(0, 1)$, that is the ordinate axis, as the space where the coefficients of the signals are supposed to be non zero. As long as the angle between the green line and the abscissa is below 45°, the ℓ_1 minimization will select a coefficient on the ordinates. If this angle is more than 45°, the situation is inverted, the coefficient on the abscissa is selected (see figure 2.8, right). Generalizing this idea, it seems that as long as the null space of $\mathbf{A}$ is "sufficiently far" from the space where the signal is supposed to be sparse, minimizing the ℓ_1 norm will yield the correct coefficients.

This intuition can be formalised through the notion of null space property (Cohen et al. 2009), whose definition is given below.

Definition 2. *Let* $\mathbf{A}$ *be a* $m \times n$ *matrix.* $\mathbf{A}$ *has the null space property relative to a set* T *if for all vector* $\boldsymbol{h} \in \ker \mathbf{A} \backslash \{0\}$

$$\|\boldsymbol{h}_T\|_{\ell_1} \leqslant \frac{1}{2}\|\boldsymbol{h}\|_{\ell_1} \tag{2.12}$$

or equivalently

$$\|\boldsymbol{h}_T\|_{\ell_1} \leqslant \|\boldsymbol{h}_{\bar{T}}\|_{\ell_1} \tag{2.13}$$

where $\boldsymbol{h}_T$ *is the vector equal to* $\boldsymbol{h}$ *on indices of* T *and 0 elsewhere.*

The following theorem gives a sufficient and necessary condition for uniqueness of the solution, which we state as in Foucart & Rauhut (2013),

Theorem 3 (Cohen et al. (2009))**.** *Given* $\mathbf{A}$ *a* $m \times n$ *matrix, every vector* $\boldsymbol{x} \in \mathbb{R}^n$ *supported on a set* T *is the unique solution of* (P_{ℓ_1}) *with* $\boldsymbol{y} = \mathbf{A}\boldsymbol{x}$ *if and only if* $\mathbf{A}$ *satisfies the null space property relative to* T*.*

Note that the same proposition holds if one replaces ℓ_1 by ℓ_q, $0 \leqslant q \leqslant 1$ (Foucart et al. 2010). This theorem implies that if the true coefficients $\mathbf{x}_t$ are in T and $\mathbf{A}$ satisfies the null space property relative to the set of indices T, then (P_{ℓ_0}) recovers the true coefficients. Note that if $\mathbf{A}$ satisfies the null space property relative to all subsets of indices of cardinality k, it means that any k-sparse vector is exactly recovered by (P_{ℓ_1}). The proof is brief and gives some geometric intuition, therefore we are going to reproduce it here.[2]

Let us denote by $\text{supp}(\boldsymbol{x})$ the support of a vector $\boldsymbol{x}$, that is the set of indices of its non-zero components. The property "any element $\boldsymbol{x}$ whose support is in T is the unique solution of (P_{ℓ_1})" can be stated as "$\forall \boldsymbol{x}$ such that $\text{supp}(\boldsymbol{x}) \in T$, $B_{\boldsymbol{x}} \cap K_{\boldsymbol{x}} = \{\boldsymbol{x}\}$" where $B_{\boldsymbol{x}}$ is the ℓ_1 ball of radius $\|\boldsymbol{x}\|_{\ell_1}$ and $K_{\boldsymbol{x}} = \{\boldsymbol{h}, \mathbf{A}\boldsymbol{h} = \mathbf{A}\boldsymbol{x}\}$.

Proof. We prove each implication by contradiction. Let us first suppose that $\mathbf{A}$ does not verify the null space property relative to T. Then there exists a vector $\boldsymbol{h} \in \ker\mathbf{A}$ such that $\|\boldsymbol{h}_{\bar{T}}\|_{\ell_1} \leqslant \|\boldsymbol{h}_T\|_{\ell_1}$ (see figure 2.9, left). Since $\mathbf{A}\boldsymbol{h} = 0$, we have $\mathbf{A}\boldsymbol{h}_T = \mathbf{A}(-\boldsymbol{h}_{\bar{T}})$. In other words we have found another element than $\boldsymbol{h}_T$ in $B_{\boldsymbol{h}_T} \cap K_{\boldsymbol{h}_T}$, we do not have exact recovery.

Conversely, let us suppose that for some $\boldsymbol{x}_0$ with $\text{supp}(\boldsymbol{x}_0) \in T$, $B_{\boldsymbol{x}_0} \cap K_{\boldsymbol{x}_0}$ contains an element $\boldsymbol{x}_0' \neq \boldsymbol{x}_0$. Then we have $\mathbf{A}\boldsymbol{x}_0 = \mathbf{A}\boldsymbol{x}_0'$ therefore $\boldsymbol{h} := \boldsymbol{x}_0 - \boldsymbol{x}_0' \in \ker\mathbf{A}$ and $\boldsymbol{h} \neq 0$. (see figure 2.9, right). We can write $\boldsymbol{h} = \boldsymbol{h}_T + \boldsymbol{h}_{\bar{T}}$, and

$$\begin{aligned}
\|\boldsymbol{h}_{\bar{T}}\|_{\ell_1} &= \|\boldsymbol{x}_0'\|_{\ell_1} - \|\boldsymbol{x}_0 - \boldsymbol{h}_T\|_{\ell_1} && (2.14)\\
&\leqslant \|\boldsymbol{x}_0'\|_{\ell_1} - |\|\boldsymbol{x}_0\|_{\ell_1} - \|\boldsymbol{h}_T\|_{\ell_1}| && (2.15)\\
&\leqslant \|\boldsymbol{x}_0'\|_{\ell_1} - \|\boldsymbol{x}_0\|_{\ell_1} + \|\boldsymbol{h}_T\|_{\ell_1} && (2.16)
\end{aligned}$$

By definition of B_{x_0}, since $\boldsymbol{x}_0' \in B_{x_0}$, $\|\boldsymbol{x}_0'\|_{\ell_1} \leqslant \|\boldsymbol{x}_0\|_{\ell_1}$ and

$$\|\boldsymbol{h}_{\bar{T}}\|_{\ell_1} \leqslant \|\boldsymbol{h}_T\|_{\ell_1} \tag{2.17}$$

$\mathbf{A}$ does not satisfy the null space property relative to T. □

[2] Largely inspired by Dustin G. Mixon's a particularly clear proof.

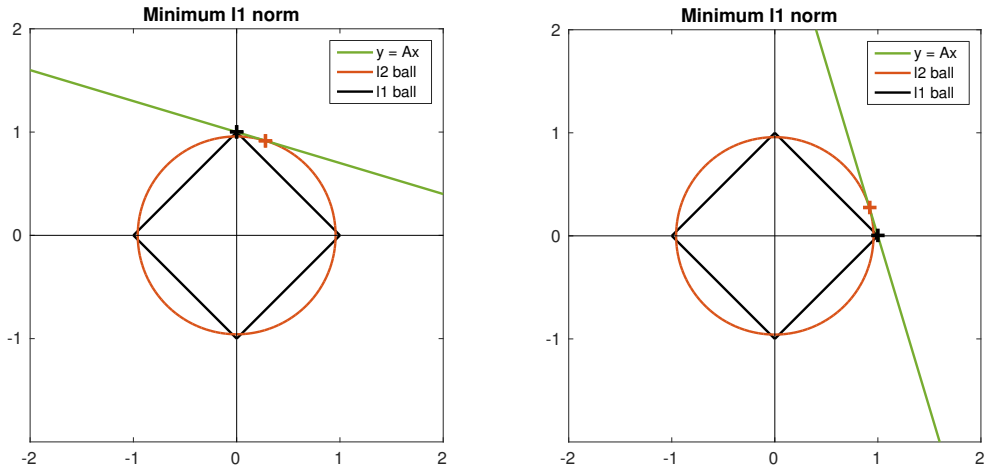

Figure 2.8: Example of ℓ_1 minimization in the case $1 = ax_1 + bx_2$ for $(a, b) = (0.3, 1)$ (left) and $(a, b) = (1, 0.3)$ (right). The red and black markers correspond respectively to the solutions of equation (P_{ℓ_q}) for $q = 2$ and $q = 1$.

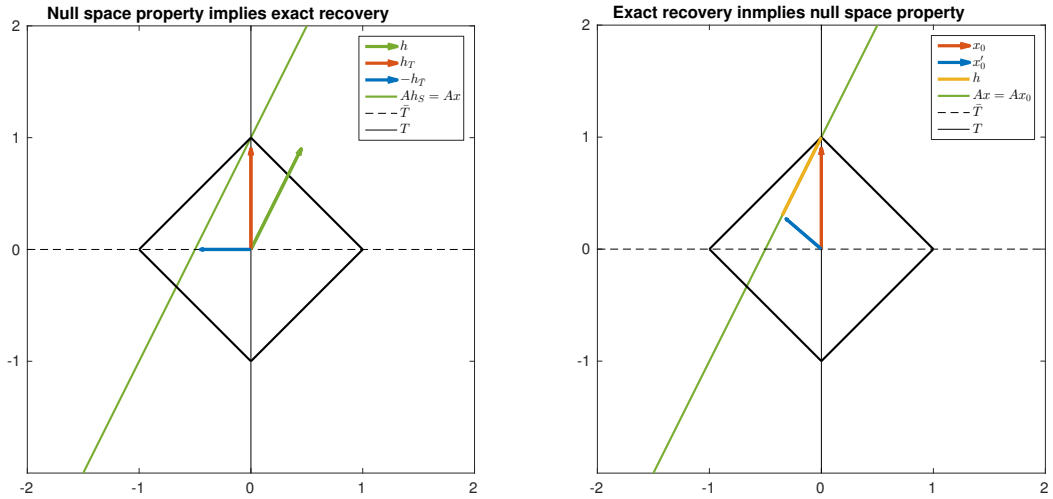

Figure 2.9: Illustrations for the proof of theorem 3

The diagram used for the proof seems also to suggest that the ℓ_1 minimization will either find the correct elements of the dictionary or fail completely. It has been shown that when the columns of $\mathbf{A}$ are randomly distributed, ℓ_1 methods tend to fail below a certain number of measurements and succeed otherwise with a sharp transition (Amelunxen et al. 2013).

The null space property gives some geometric intuition on the possibilities of success of ℓ_1 minimization. By modifying slightly the null space property definition, one can also obtain guarantees on the stability of the recovery when noise corrupts the data set. Since these are not going to be used, we refer the reader to chapter 4 of Foucart & Rauhut (2013). Theoretical results on compressed sensing can be obtained in terms of how well conditioned are the submatrices made of columns of $\mathbf{A}$. We present this approach and how it can be used in practice next section.

Conditioning of submatrices/Restricted isometry constants

Let us go back to problem ($P_{\ell_0,\epsilon}$), page 58. We have measurements $\boldsymbol{y} = \mathbf{A}\boldsymbol{x}_t + \boldsymbol{e}$ where $\boldsymbol{e}$ is some noise and $\boldsymbol{x}_t$ are the true coefficients. If there were no time constraints, we could select any set of (non redundant) s indices $T \in \{1..n\}^s$, the corresponding columns of $\mathbf{A}$, whose concatenation will be denoted by $\mathbf{A}_T$.

There exists $\boldsymbol{x}$ such that $\|\mathbf{A}_T\boldsymbol{x} - \boldsymbol{y}\|_{\ell_2} \leqslant \epsilon$ if and only if $\min_{\boldsymbol{x}\in\mathbb{R}^s} \|\mathbf{A}_T\boldsymbol{x} - \boldsymbol{y}\|_{\ell_2} \leqslant \epsilon$. Such a quantity can be computed explicitly since it is a linear least square problem. Provided $\mathbf{A}_T$ is full rank,

$$\min_{\boldsymbol{x}\in\mathbb{R}^s} \|\mathbf{A}_T\boldsymbol{x} - \boldsymbol{y}\|_{\ell_2} = \boldsymbol{y}^*(\mathbf{I} - \mathbf{A}_T(\mathbf{A}_T^*\mathbf{A}_T)^{-1}\mathbf{A}_T^*)\boldsymbol{y} \tag{2.18}$$

$$\arg\min_{\boldsymbol{x}\in\mathbb{R}^s} \|\mathbf{A}_T\boldsymbol{x} - \boldsymbol{y}\|_{\ell_2} =: \widehat{\boldsymbol{x}_T} = (\mathbf{A}_T^*\mathbf{A}_T)^{-1}\mathbf{A}_T^*\boldsymbol{y} \tag{2.19}$$

where * designates the matrix transposition and not ^{T} like in the rest of the text, to avoid confusion with T, the subset of indices. In particular when the columns corresponding to the true indices are selected,

$$\widehat{\boldsymbol{x}_T} = \boldsymbol{x}_t + (\mathbf{A}_T^*\mathbf{A}_T)^{-1}\mathbf{A}_T^*\boldsymbol{e}. \tag{2.20}$$

To have any chance of recovering the coefficients of $\boldsymbol{x}_t$, the matrix $\mathbf{A}_T^*\mathbf{A}_T$ should not be "too close" to a non-invertible matrix. This intuitive idea can be made precise through the notion of condition number. This one is defined as

Definition 3 (Condition number)**.** *The condition number of an invertible matrix* $\mathbf{M}$ *associated to the norm* ℓ_p *is defined as*

$$C_p(\mathbf{M}) = \sup_{\boldsymbol{u}\neq 0, \boldsymbol{v}\neq 0} \frac{\|\mathbf{M}^{-1}\boldsymbol{u}\|_{\ell_p}}{\|\boldsymbol{u}\|_{\ell_p}} \frac{\|\boldsymbol{v}\|_{\ell_p}}{\|\mathbf{M}^{-1}\boldsymbol{v}\|_{\ell_p}} = \|\mathbf{M}^{-1}\|_{p\to p}\|\mathbf{M}\|_{p\to p} \tag{2.21}$$

where

$$\|\mathbf{M}\|_{p\to p} = \sup_{\boldsymbol{u}\neq 0} \frac{\|\mathbf{M}\boldsymbol{u}\|_{\ell_p}}{\|\boldsymbol{u}\|_{\ell_p}} \tag{2.22}$$

is the operator norm.

When $p = 2$, it is well known that

$$\sigma_{\min} \leqslant \|\mathbf{A}\|_{2\to 2} \leqslant \sigma_{\max} \tag{2.23}$$

where $\sigma_{\min}$ and $\sigma_{\max}$ are respectively the minimum and maximum singular values of $\mathbf{A}$. These ones are the square root of the eigenvalues of $\mathbf{A}^*\mathbf{A}$. As a consequence, $C_2(\mathbf{A}) = \sigma_{\max}/\sigma_{\min}$. The smaller the condition number is, the better will be the precision on $\boldsymbol{x}$.

To illustrate the importance of the condition number, let us consider the problem of the stability of least square estimates. We consider a signal $\boldsymbol{y} = \mathbf{A}_T\boldsymbol{x}_t + \boldsymbol{e}$ where $\boldsymbol{e}$ is a Gaussian vector independent and identically distributed of variance σ_e^2. The least square estimate of $\boldsymbol{x}_t$ is given by equation (2.20). The covariance matrix $\mathbf{V}$ of $\widehat{\boldsymbol{x}}$ is equal to $\sigma_e^2(\mathbf{A}_T^*\mathbf{A}_T)^{-1}$. We take

as a measurement of the dispersion of a vector the square root of the sum of the variances of its components, so that

$$\sigma_{\boldsymbol{y}} := \sqrt{\sum_{i=1}^{m} \mathrm{Var}\{e_i\}} = \sqrt{m}\sigma_e \tag{2.24}$$

$$\sigma_{\widehat{\boldsymbol{x}}} := \sqrt{\sum_{i=1}^{s} \mathrm{Var}\{\widehat{x}_i\}} = \sqrt{\mathrm{tr}(\mathbf{V})} = \sigma_e \sqrt{\sum_{i=1}^{s} \frac{1}{\sigma_i^2}} = \frac{\sigma_{\boldsymbol{y}}}{\sqrt{m}} \sqrt{\sum_{i=1}^{s} \frac{1}{\sigma_i^2}} \tag{2.25}$$

Let us denote by $\mathbf{A}_T^+ := (\mathbf{A}_T^* \mathbf{A}_T)^{-1}$. Since $(\mathbf{A}_T^+)^* \mathbf{A}_T^+ = (\mathbf{A}_T^* \mathbf{A}_T)^{-1}$ we can apply (2.23) to $\mathbf{A}_T^+ \boldsymbol{y} = \widehat{\boldsymbol{x}}$, and get

$$\frac{1}{\sigma_{\max}} \|\boldsymbol{y}\|_{\ell_2} \leqslant \|\widehat{\boldsymbol{x}}\|_{\ell_2} \leqslant \frac{1}{\sigma_{\min}} \|\boldsymbol{y}\|_{\ell_2}.$$

Dividing by $\sigma_{\widehat{\boldsymbol{x}}}$ we obtain

$$\frac{\sqrt{m}}{\sqrt{\sum_{i=1}^{m} \frac{\sigma_{\max^2}}{\sigma_i^2}}} \frac{\|\boldsymbol{y}\|_{\ell_2}}{\sigma_{\boldsymbol{y}}} \leqslant \frac{\|\widehat{\boldsymbol{x}}\|_{\ell_2}}{\sigma_{\widehat{\boldsymbol{x}}}} \leqslant \frac{\sqrt{m}}{\sqrt{\sum_{i=1}^{m} \frac{\sigma_{\min^2}}{\sigma_i^2}}} \frac{\|\boldsymbol{y}\|_{\ell_2}}{\sigma_{\boldsymbol{y}}},$$

therefore,

$$\sqrt{\frac{m}{s}} \frac{1}{C_2(\mathbf{A}_T)} \frac{\|\boldsymbol{y}\|_{\ell_2}}{\sigma_{\boldsymbol{y}}} \leqslant \frac{\|\widehat{\boldsymbol{x}}\|_{\ell_2}}{\sigma_{\widehat{\boldsymbol{x}}}} \leqslant \sqrt{\frac{m}{s}} C_2(\mathbf{A}_T) \frac{\|\boldsymbol{y}\|_{\ell_2}}{\sigma_{\boldsymbol{y}}}. \tag{2.26}$$

Defining $\mathrm{SNR}(\boldsymbol{z}) = \sqrt{p} \frac{\|\boldsymbol{z}\|_{\ell_2}}{\sigma_{\boldsymbol{z}}}$ the signal-to-noise ratio of a vector $\boldsymbol{z}$ with p components, equation (2.26) can be re-written

$$\frac{1}{C_2(\mathbf{A}_T)} \mathrm{SNR}(\boldsymbol{y}) \leqslant \mathrm{SNR}(\widehat{\boldsymbol{x}}) \leqslant C_2(\mathbf{A}_T) \mathrm{SNR}(\boldsymbol{y}) \tag{2.27}$$

The inverse of the condition number then determines how close to zero can be the signal to noise ratio for a given estimate $\widehat{x}$. When the noise has a component in the direction of the eigenvector corresponding to the singular value $\sigma_{\min}$, it is amplified in the estimation process and might dominate the signal. The lower bound provides worst case scenario. On the other hand a high condition number also provides best case scenario if it is in the direction of $\sigma_{\max}$, it will be damped in the inversion.

A remarkable fact is that the conditioning of submatrices not only ensures stability with respect to noise, but also uniqueness of the solution of (P_{ℓ_1}) (see page 66), as we shall see theorem 4. Before presenting a few results, let us define an analogous of the null space property. We follow Foucart & Lai (2009) for the definition and notation, and name the notion "Condition number bound" to avoid confusion with other metrics.

Definition 4 (Condition number bound)**.** *Let s be a positive integer, $\mathbf{A}$ a $m \times n$ real or complex matrix, and $\alpha_s, \beta_s \geqslant 0$ be respectively the superior and inferior bounds of α, β such that*

$$\forall \boldsymbol{z} \in \mathbb{R}^n \text{ or } \mathbb{C}^n, \|\boldsymbol{z}\|_{\ell_0} \leqslant s, \qquad \alpha \|\boldsymbol{z}\|_{\ell_2} \leqslant \|\mathbf{A}\boldsymbol{z}\|_{\ell_2} \leqslant \beta \|\boldsymbol{z}\|_{\ell_2}, \tag{2.28}$$

the s-condition number bound of the matrix $\mathbf{A}$ is defined as $\gamma_s = \frac{\beta_s^2}{\alpha_s^2}$.

Foucart & Lai (2009) considers the more general problem of ℓ_q minimization

$$\underset{\boldsymbol{x}\in\mathbb{R}^n}{\arg\min} \quad \|\boldsymbol{x}\|_{\ell_q} \quad \text{s. t.} \quad \mathbf{A}\boldsymbol{x} = \boldsymbol{y} \tag{P_{ℓ_q}}$$

and show the following theorem,

Theorem 4 (Foucart & Lai (2009))**.** *Given $0 < q \leqslant 1$, if for some $t \geqslant s$,*

$$\gamma_{2t} - 1 < 4(\sqrt{2} - 1)\left(\frac{t}{s}\right)^{\frac{1}{q}-\frac{1}{2}} \tag{2.29}$$

then all s-sparse vector $\boldsymbol{x}_t$ is exactly recovered by solving (P_{ℓ_q}) with $\boldsymbol{y} = \mathbf{A}\boldsymbol{x}_t$.

Note that the smaller q is, the less stringent the requirement is. It means that ℓ_q minimization should be a little bit more efficient. The problem is that for $q < 1$ the objective function is not convex anymore. Another problem, common to all such results, is that when the dictionary $\mathbf{A}$ is not drawn from random variables, the conditions are never met. Furthermore, for $s > 2$, computing s-minimal condition numbers is very lengthy. What follows is an attempt to remedy this problem by putting constraints on the true signal rather than on all s sparse signals. In general situations, we do not know where the true signal is, therefore the result might seem useless. However, there are situations where some planets are almost certainly in the signal and one might want to check whether it is possible to detect an additional one, or if it is absorbed in the fit of the other planets. The slight modification of theorem 4 motivates the use of a somewhat generalized spectral window in that situation.

In the new hypotheses, we require information only on the mutual coherence defined in equation (2.5) and on the condition number of the matrix made of the vectors truly in the signal concatenated with another vector of the dictionary. We first define a notion analogous to the null space property relative to S.

Definition 5 (Condition number bound relative to S)**.** *Let S be a subset of s indices of $\{1..n\}$. Let be $\gamma_{S,1} = \beta_{S,1}^2/\ \alpha_{S,1}^2$ where $\beta_{S,1}$ and $\alpha_{S,1}$ are the tightest constants such that for any vector $\boldsymbol{a}_k$, k^{th} column of $\mathbf{A}$ with $k \in \bar{S}$, the concatenation of $\mathbf{A}_S$ and $\boldsymbol{a}_k$, $[\mathbf{A}_S, \boldsymbol{a}_k]$ verifies*

$$\forall \boldsymbol{z} \in \mathbb{R}^n \text{ or } \mathbb{C}^n, \|\boldsymbol{z}\|_{\ell_0} \leqslant s+1, \qquad \alpha_{S,1}\|\boldsymbol{z}\|_{\ell_2} \leqslant \|[\mathbf{A}_S, \boldsymbol{a}_k]\boldsymbol{z}\|_{\ell_2} \leqslant \beta_{S,1}\|\boldsymbol{z}\|_{\ell_2}. \tag{2.30}$$

Then we can state

Theorem 5 (Foucart & Lai (2009), modified)**.** *Provided that*

$$\gamma_{S,1} < s^{1/q-1/2}4(\sqrt{2}-1) + 1 \tag{2.31}$$

and

$$\beta_2^2 - \alpha_2^2 \leqslant \beta_{S,1}^2 - \alpha_{S,1}^2 \tag{2.32}$$

then all vector $\boldsymbol{x}_t$ with support in S is exactly recovered by solving (P_{ℓ_q}) with $\boldsymbol{y} = \mathbf{A}\boldsymbol{x}_t$.

The proof is given in appendix B. In fact, provided the columns of $\mathbf{A}$ are normalized in ℓ_2 norm, $\beta_2^2 - \alpha_2^2$ is twice the maximum of the spectral window, as shown below. Instead of computing the condition numbers of all matrices of size $2t$, $t \geqslant s$, we need the correlation of columns two by two and the condition of the set of interest plus one vector. This is incomparably faster to compute as soon as $s \geqslant 2$.

The notion of condition number is scarcely used directly on matrix $\mathbf{A}_T$. The theory has been made mostly with the notion of restricted isometry constant (Candes & Tao 2005). Since it is ubiquitous in Compressed Sensing literature, we give the definition below.

Definition 6 (Restricted isometry constant)**.** *For a $m \times n$ matrix $\mathbf{A}$, the s-restricted isometry constant is the smallest $\delta_s < 1$ such that for any vector $\boldsymbol{c}$ with at most s non zero components,*

$$(1 - \delta_s)\|\boldsymbol{c}\|_{\ell_2}^2 \leqslant \|\mathbf{A}\boldsymbol{c}\|_{\ell_2}^2 \leqslant (1 + \delta_s)\|\boldsymbol{c}\|_{\ell_2}^2 \tag{2.33}$$

for all s-sparse vector. Or equivalently,

$$\delta_s = \max_{T \subset \{1..n\}, |T| = s} \|\mathbf{A}_T^* \mathbf{A}_T - \mathbf{I}\|_{2 \to 2} \tag{2.34}$$

where $\mathbf{I}$ is the identity matrix of size s. When $\mathbf{A}$ satisfies an inequality such as (2.33) with $\delta_s < 1$, it is said to have verify the s-restricted isometry property.

It is easy to check that when the columns of $\mathbf{A}$ are normalized by their ℓ_2 norms and $s = 2$, the restricted isometry constant is equal to the mutual coherence, that is $\delta_2 = \mu$. We have also $(\beta_2^2 - \alpha_2^2)/(\beta_2^2 + \alpha_2^2) = ((1 + \delta_2) - (1 - \delta_2))/2 = \delta_2 = \mu$. We have the following, though trivial, important property (e.g. Foucart & Rauhut 2013)

$$\mu = \delta_2 \leqslant \delta_3 \leqslant ... \leqslant \delta_m \tag{2.35}$$

The problem with the restricted isometry constant is that its value changes with the units of $\mathbf{A}$. We have shown that when $s = 2$, the way to multiply the columns of $\mathbf{A}$ by constants that give the smallest δ_2 is simply to normalize $\mathbf{A}$ columns in norm ℓ_2. The proof was not included in this work, since it is not essential for our purposes. The case of $s \geqslant 3$ is more complicated and no result was found, but it seems reasonable to conjecture that normalizing columns of $\mathbf{A}$ gives the smallest δ_s for all s. Therefore we will systematically normalize the columns of the dictionary in what follows.

To conclude, let us assume that the columns of $\mathbf{A}$ are normalized, defining $\delta_{S,1} = (\gamma_{S,1} - 1)/(\gamma_{S,1} + 1)$ we have

Theorem 6 (Foucart & Lai 2009, adapted)**.** *Provided that*

$$\delta_{S,1} < \frac{1}{1 + \frac{\sqrt{s}}{2(\sqrt{2}-1)}} \approx \frac{1}{1 + 0.207\sqrt{s}} \tag{2.36}$$

and

$$|S_w(\omega)| \leqslant \delta_{S,1} \quad \forall \omega \in \Omega \backslash \{0\} \tag{2.37}$$

then all vector with support in $S \subset \Omega$, $|S| = s$, is exactly recovered by solving (P_{ℓ_1}).

So if we want to ensure that four vectors with support in S will be exactly recovered, we need to have a maximum of the spectral window and $\delta_{S,1}$ lower than 0.29. This condition will never be satisfied because of a fact mentioned earlier: close frequency have correlation almost one. However, the study of practical cases suggested that the real problem is the spurious adding of distant aliases (as in figure 2.3). The theorem above is an invitation to consider the restricted isometry constant as a spectral window that can take in argument more than two frequencies. Like most of the compressed sensing theoretical results, it can very scarcely been shown to hold in real cases.

2.1.4 Noisy data

Theorem 5 was stated in this way to motivate the use of the conditioning of submatrices to investigate possible correlations with the signals (as done in Hara et al. (2017), appendix D which is reproduced in the Appendix G.10), just like the spectral window, whose information is valuable regardless of the noise. This one has an important effect on sparse recovery, which we now discuss. The privileged framework to consider noise is the problem termed Basis Pursuit De-noising (BPDN, Chen et al. 1998), that is

$$\boldsymbol{x}^\star = \underset{\boldsymbol{x}\in\mathbb{C}^n}{\arg\min} \quad \|\boldsymbol{x}\|_{\ell_1} \quad \text{s. t.} \quad \|\mathbf{A}\boldsymbol{x} - \boldsymbol{y}\|_{\ell_2} \leqslant \epsilon \tag{BP$_\epsilon$}$$

where $\boldsymbol{x}_t$ is not necessarily exactly sparse, but is well approximated by a sparse vector (it is said to be compressible). Numerous results on the stability if the recovery have been obtained. They consist in

- Bounding the error on the reconstructed signal $\|\mathbf{A}\boldsymbol{x}^\star - \mathbf{A}\boldsymbol{x}_t\|_{\ell_2}$. These results are in particular useful in the case of data de-noising.
- Bounding the error on the $\|\boldsymbol{x}^\star - \boldsymbol{x}_t\|_{\ell_2}$ or $\|\boldsymbol{x}^\star - \boldsymbol{x}_t\|_{\ell_1}$, which is useful for inverse problems.
- Establishing conditions to have $\text{supp}(\boldsymbol{x}^\star) \subset \text{supp}(\boldsymbol{x}_t)$ where supp denotes the non zero elements of a vector.

One of the most famous theorem of the Compressed Sensing theory is stated in terms of restricted isometry property

Theorem 7. *(Candès et al. 2006b) Suppose that $\boldsymbol{x}_t$ is an arbitrary vector in $\mathbb{R}^n$ and let $\boldsymbol{x}_{t,s}$ be the truncated vector corresponding to the s largest values of $\boldsymbol{x}_t$ (in absolute value). We have measurements $\boldsymbol{y} = \mathbf{A}\boldsymbol{x}_t + \boldsymbol{e}$ with $\|\boldsymbol{e}\|_{\ell_2} \leqslant \epsilon$. Under the hypothesis that $\delta_{3s} + \delta_{4s} < 2$, the solution $\boldsymbol{x}^\star$ of the (BP$_\epsilon$) minimisation obeys*

$$\|\boldsymbol{x}^\star - \boldsymbol{x}_t\|_{\ell_2} \leqslant C_{1,s}\epsilon + C_{2,s}\frac{\|\boldsymbol{x}_t - \boldsymbol{x}_{t,s}\|_{\ell_2}}{\sqrt{s}} \tag{2.38}$$

with $C_{1,s} \approx 12.04$ and $C_{2,s} \approx 8.77$ for $\delta_{4s} = 1/5$.

In other words, the error on the reconstructed vectors can be bounded by the noise level and the error made by only selecting s coefficients of $\boldsymbol{x}$. Other results have been obtained in terms of mutual coherence by Donoho et al. (2006), and to the knowledge of the author, the sharpest bounds so far can be found in Foucart (2010). Let us note that the solution of (BP$_\epsilon$) will always satisfy $\|\mathbf{A}\boldsymbol{x}^\star - \boldsymbol{y}\|_{\ell_2} = \epsilon$ unless $\|\boldsymbol{y}\|_{\ell_2} \leqslant \epsilon$, i.e., $\boldsymbol{x}^\star = 0$. In all non trivial cases, the boundary will be reached.

2.1.5 A note on vocabulary

So far the terms sparse recovery and compressed sensing have been used indifferently, which seems to have become gradually the case in the literature. At its origin, compressed sensing or compressive sampling was used in a stricter sense, which is the recovery of sparse signals when matrix $\mathbf{A}$ has random columns Candes & Tao (2005); Donoho (2006). This concerns situations where the experimenter has the possibility to obtain random measurements of a vector $\boldsymbol{z}_t \in \mathbb{R}^n$ or $\mathbb{C}^n$ plus a noise $\boldsymbol{e}$. The k-th observation is then $y_k = \langle \boldsymbol{\Phi}_k, \boldsymbol{z}_t \rangle + e_k$. If it is known that $\boldsymbol{z}$ has a sparse representation in some basis Ψ, $n \times n$ matrix, then denoting by Φ the matrix whose rows are the $\boldsymbol{\Phi}_k$ and we have $k = 1..m$ measurements, then

$$\boldsymbol{y} = \Phi\Psi\boldsymbol{x}_t + \boldsymbol{e}. \tag{2.39}$$

The fact that $\boldsymbol{\Phi}_k$s are drawn randomly causes the matrix $\Phi\Psi$ to satisfy null space properties or restricted isometry properties with an overwhelming probability, provided enough measurements are taken. One can show that under suitable conditions the sparse recovery will succeed with high probability if the number of measurements satisfies a certain inequality. For instance if the matrix is drawn from Gaussian variables, recovery will succeed with high probability if the number of measurement m satisfies $m > 2Cs(1 + \log(n/s))$ for some constant C (e.g. Foucart & Rauhut 2013, chap. 9, the example given is theorem 9.29). Applications of such methods are found in electronics (e.g. Mishali et al. 2008; Tropp et al. 2010), radar (e.g. Potter et al. 2010), MRI (e.g. Haldar et al. 2011), astronomy (Bobin et al. 2008)... As an example, it allows to sample time continuous signals with a rate much lower than Nyquist frequency and still allow precise reconstruction of the signal.

In the terms of Donoho (2006), the original idea came from the fact that to store data, much of its original information is thrown away. This happens when music is stored as MP3 or AAC, or images are represented by wavelets. The total number of wavelet coefficients used is much lover than the precise WAV time series, or pixel by pixel representation. The compressed sensing theory was an attempt to avoid throwing away most of the data, and bridge the gap between the complexity of the sensing and the storage. Let us finally note that the sensing can be viewed as a communication channel. The random sensing can be viewed as a way to encode the symbols transmitted as far as possible from each others to avoid error, as wished in communication theory (Shannon 1948). Connections between this field and ℓ_1 minimization is done by Donoho (2005) for instance.

2.1.6 Continuous dictionaries

In this section we assumed that matrix $\mathbf{A}$ is of the form (2.2) where the frequencies are on a grid $\Omega = (\omega_j)_{j=1..n}$, arbitrarily dense such that $\omega_1 = 0$, $\omega_n = \omega_{\max}$ and $\omega_{j+1} = \omega_j + \Delta\omega$. However, the true frequencies of the planets are real numbers, and will therefore necessarily fall between the grid elements.

Furthermore, we also stated that correlation between subsequent columns of $\mathbf{A}$ is deadly to our ability to verify the hypotheses of theoretical results. Why not simply spacing out the grid to reduce correlation? Anticipating on the next section, we observed that in practice, loosing precision on the frequencies is worse than correlation of the dictionary.

There are theories for continuous dictionary in the sparse recovery literature but so far, the theories we found could not directly be applied to the radial velocity measurements. The

mathematical tool they use are significantly more advanced than the discrete case, and we quickly gave up the purpose of extending them formally to fit our problem. However, we took inspiration from a theoretical result of Tang et al. (2013a) to perform a smoothing of the output of (BP_ϵ) described section 2.2.3. Our reasoning is presented below.

Chandrasekaran et al. (2012) suggests to use an "atomic norm" that extends (P_{ℓ_1}) to infinite dictionaries. Practical methods to solve the new minimization problem are designed in Candès & Fernandez-Granda (2013) and Tang et al. (2013b). The atomic norm $\|\boldsymbol{y}\|_{\mathcal{A}}$, of $\boldsymbol{y} \in \mathbb{R}^m$ or $\mathbb{C}^m$ defined for a dictionary $\mathcal{A}$ is the smallest ℓ_1 norm of a combination of vectors of the dictionary reproducing $\boldsymbol{y}$:

$$\|\boldsymbol{y}\|_{\mathcal{A}} = \inf \left\{ \sum_j |x_j|, \boldsymbol{y} = \sum_j x_j \boldsymbol{a}_j(t) \right\} \tag{2.40}$$

If the observations were not noisy, computing the atomic norm of $\boldsymbol{y}$ would be sufficient. As this is obviously not the case, the following problem is considered.

$$\boldsymbol{u}^\star = \underset{\boldsymbol{u}\in\mathbb{C}^m}{\arg\min} \quad \|\boldsymbol{u} - \boldsymbol{y}(t)\|_{\ell_2}^2 + \lambda \|\boldsymbol{u}\|_{\mathcal{A}} \tag{2.41}$$

where λ is a positive real number fixed according to the noise. The coefficient λ can be interpreted as a Lagrange multiplier, and this problem can be seen as maximizing a posterior likelihood with a prior on u. The quantities we are interested in are the dictionary elements $\boldsymbol{a}_j^\star$ and the coefficients $\boldsymbol{x}^\star$ selected by the minimization, where $\boldsymbol{u}^\star = \sum_{j=1}^{S^\star} x_j^\star \boldsymbol{a}_j^\star(t)$.

The problem of the algorithms of Candès & Fernandez-Granda (2013) and Tang et al. (2013b) is that they rely on evenly sampled signals, with possibly missing measurements. Since radial velocity signals are unevenly spaced, the time samples have to be seen as approximately located on a grid with many missing samples. This would require to manipulate large matrices and we did not attempt to further investigate these algorithms.

However, Tang et al. (2013a) shows that when the function associating a column of the dictionary $\boldsymbol{a}(\omega)$ to ω is continuous (which is the case for us), as the maximum difference between frequencies on the grid tends to zero, the solution of (BP_ϵ) tends to the one of (2.41). Especially, corollary 1 of Tang et al. (2013a) states that the summed amplitudes of coefficients of $\boldsymbol{x}^\star$ within a certain distance η_0 from the actual peak in the signal tend to the appropriate value. In the proof, they choose a tolerance ϵ (appearing in (BP_ϵ)) such that the balls of width η_0 centred around the true peaks have a null intersection. Thus, it seems reasonable to select η as the largest interval within which the probability to distinguish frequencies is low. Values such as $\approx 0.5\pi/T_{\mathrm{obs}}$ to π/T_{obs} are robust in practice.

There is other algorithms and theories for spectral line estimation that accounts for the continuous nature of the dictionary. Chen & Chi (2014) builds on the so-called structured matrix completion but again, the algorithm requires the signal samples to be evenly spaced. This is also the case for the theory developed for shift-invariant signals in Eldar (2009).

2.2 Application to radial velocity data

2.2.1 Problem formulation

Among the minimization problems presented last section, the most natural to try out is (BP_ϵ) (page 73). Indeed, unlike greedy algorithms, several coefficients are searched at once. Un-

like (P_{ℓ_q}) with $q < 1$, (page 71), (BP_ϵ) is a convex problem and is therefore expected to be relatively fast to solve. Finally, the formulation accounts for possible noises, which are unavoidable in real life. It can be shown that (BP_ϵ) is equivalent to two other formulations, each of which leads to several algorithms as long as $\epsilon \neq 0$. First the regularized Least-Square,

$$\underset{\boldsymbol{x}\in\mathbb{C}^n}{\arg\min} \quad \|\mathbf{A}\boldsymbol{x}-\boldsymbol{y}\|_{\ell_2}{}^2 + \lambda\|\boldsymbol{x}\|_{\ell_1} \tag{QP_λ}$$

proposed by Chen et al. (1998). Interestingly enough, this can be seen as searching the posterior maximum of the spectrum with a Laplace prior and Gaussian ii.d. noise. Indeed, (QP_λ) is equivalent to

$$\underset{\boldsymbol{x}\in\mathbb{C}^n}{\arg\max} \quad \frac{1}{\sqrt{2\pi}^m}\,\mathrm{e}^{-\frac{1}{2}\|\mathbf{A}\boldsymbol{x}-\boldsymbol{y}\|_{\ell_2}^2}\,\mathrm{e}^{-\frac{\lambda}{2}\|\boldsymbol{x}\|_{\ell_1}} \tag{2.42}$$

(QP_λ) is also equivalent to solving a least square problem restricted to a ℓ_1 ball of given radius

$$\underset{\boldsymbol{x}\in\mathbb{C}^n}{\arg\min} \quad \|\mathbf{A}\boldsymbol{x}-\boldsymbol{y}\|_{\ell_2}{}^2 \quad \text{s. t.} \quad \|\boldsymbol{x}\|_{\ell_1} \leqslant \tau \tag{LS_τ}$$

This formulation was introduced by Tibshirani (1994) under the name LASSO (Least Absolute Shrinkage and Selection Operator), ubiquitous in statistical learning.

These methods yield the same result for an appropriate choice of ϵ, λ and τ. The equivalence simply follows from the fact that (BP_ϵ) and (LS_τ) have a Lagrangian formulation of the form (QP_λ). Unfortunately, given one of them there is no mean to detemine the values of the other scheme parameters that lead to the same solution. Since Basis Pursuit offers directly a control on how far from the data we allow the solution to be, it was selected as a starting point. In this section, we provide some background and some modifications of the algorithm made in Hara et al. (2017).

2.2.2 A search for an efficient algorithm

There is a considerable literature on the ways to solve (BP_ϵ). Our method of selection consisted basically in trying public codes in several situations, with simulated and real systems. The decisive part was a blind test on five systems sent by Alexandre Correia (Aveiro Universidade), which were HD 69830 (Lovis et al. 2006), HD 40307, HD 10180 (Lovis et al. 2011b), 55 Cnc, GJ 876 (Correia et al. 2010). The analysis was done with several algorithms and sent to A. Correia, who could confirm if the periods retrieved were correct.

Testing the method on real data sets first might be puzzling since it is unsure whether the reported signals truly are in the data or not. However, these systems have a relatively high number of high precision measurements and show high significance detection with several techniques. In particular 55 Cnc and GJ 876 were studied in-depth in Baluev (2013a) and Baluev (2015a). The method was tailored on these application to ensure it is efficient in real cases. We test its reliability on simulated cases section 2.3.2.

The existing codes we have tested for analysing radial velocity data sets are: ℓ_1-magic (Candès et al. 2006a), SparseLab (Donoho 2006), NESTA (Becker et al. 2011), CVX (Grant & Boyd 2008), Spectral Compressive Sampling (Duarte & Baraniuk 2013) and SPGL1 (van den Berg & Friedlander 2008). Since SparseLab is a collection of algorithms, we tested them even though they are not intended for solving (BP_ϵ). We tested also CoSaMP (Needell & Tropp 2009) and the algorithm MUSIC (Schmidt 1986). SPGL1 gave the best results in general for

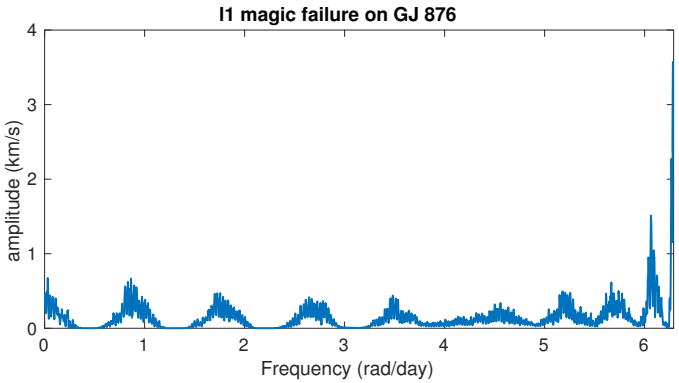

Figure 2.10: Failure of ℓ_1-magic on the 214 points data set of GJ 876. Obtained with 2000 frequencies.

exoplanetary data and consequently is the one we selected (the code can be downloaded from this link [3]). We then made several modifications to improve the performance of the algorithm for the radial velocity data analysis. We will not dwell on the details of the selection process, which would require to present each algorithm and the adjustments that were made, but highlight key reasons why the algorithms were not selected.

The first algorithm considered was ℓ_1-magic (Candès et al. 2006a). This one performed well for small data sets but is unstable when considering more than ≈ 100 data points. We re-coded the algorithm in the TRIP language (Gastineau & Laskar 2011). We found that the results were slightly different from the matlab public version. A close inspection showed that the two versions diverged because of a matrix inversion that occurs in a sub-routine, necessary for a gradient descent. When the matrix $\mathbf{A}$ is large, the matrix to be inverted, denoted by $\mathbf{B}$, is ill-conditioned. This causes a numerical instability that amplifies with iterations. We show an example of failure in figure 2.10. Matrix $\mathbf{A}$ has columns $\cos\omega_k$ and $\sin\omega_k$ for $(\omega_k) = k\Delta\omega$, $\Delta\omega = 2\pi/2000$, $k = 0..1999$. We plot the amplitude of $a(\omega_k)^2 + b(\omega_k)^2$ as a function of ω_k where $a(\omega_k)$ and $b(\omega_k)$ are respectively the coefficients of $\cos\omega_k$ and $\sin\omega_k$ of the solution of (BP_ϵ). The oscillations that are spaced with ≈ 0.8 rad/day are due to the numerical instability. Other tuning were tried, but these were found unable to retrieve correctly the frequencies in 55 Cnc and GJ 876.

The algorithms of SparseLab, MUSIC and Spectral Compressive Sampling yielded less clean figures and were missing some of the periods even on simulated data. After trying several configurations we simply dismissed those without further investigations. Let us note that the performance of Spectral Compressive Sampling is surprising, since it relies on SPGL1.

NESTA and CVX were performing closely to SPGL1. The problem with the former is that it relies on $\mathbf{A}$ to be expressed as an orthonormal basis $\mathbf{A}'$, which is a square matrix. To obtain a good resolution on time and frequency, the matrix $\mathbf{A}'$ to set in input becomes large when there are numerous measurements. In the case of 55 Cnc, that has 663 measurements available, the frequency resolution was either too low, or the algorithm too long to converge. Finally, CVX and SPGL1 performed closely, but the latter was faster.

[3] https://www.math.ucdavis.edu/~mpf/spgl1/supplement.html

2.2.3 Tuning

Regardless of the algorithm chosen for solving (BP_ϵ), we have to make several choices. These concern the form of the dictionary $\mathbf{A}$, the tolerance ϵ, whether any processing is performed on the raw data before setting it in input of the solver, or if post-processing on the output of the solver must be made. We first searched for a matrix $\mathbf{A}$ of the form (2.2). We then have to define the frequency grid $\Omega = (\omega_k)_{k=1..n}$.

Since there are many parameters to tune and they have interconnected effects, it is uneasy to optimize jointly their choices. We basically made several adjustments until the tool was found stable and efficient in all configurations in order to demonstrate its usefulness. A thorough exploration of the parameters is yet to be done.

Dictionary elements

We have to determine whether $(\cos\omega_k\boldsymbol{t}, \sin\omega_k\boldsymbol{t})$ or $(\mathrm{e}^{-\mathrm{i}\omega_k\boldsymbol{t}}, \mathrm{e}^{\mathrm{i}\omega_k\boldsymbol{t}})$ should be taken as the columns of $\mathbf{A}$. Interestingly enough, $(\mathrm{e}^{-\mathrm{i}\omega_k\boldsymbol{t}}, \mathrm{e}^{\mathrm{i}\omega_k\boldsymbol{t}})$ is more robust. So far we could not explain this feature, but made the choice to go along with it.

Figure 2.11 shows an example of the differences between the two options on HD 10180 data (Lovis et al. 2011b). To highlight the difference between the dictionary, we solve basis pursuit, which is (BP_ϵ) with $\epsilon = 0$. When $\mathbf{A}$ is made of sines and cosines, we plot $\sqrt{a(\omega_k)^2 + b(\omega_k)^2}$ where $a(\omega_k)$ and $b(\omega_k)$ are respectively the coefficients of $\cos\omega_k\boldsymbol{t}$ and $\sin\omega_k\boldsymbol{t}$, which gives figure 2.11.a. When the columns of $\mathbf{A}$ are complex, we plot as a function ω_k

$$A(\omega_k) = \sqrt{[\Re(c(\omega_k) + d(\omega_k))]^2 + [\Im(c(\omega_k) - d(\omega_k))]^2} \tag{2.43}$$

where $c(\omega_k)$ and $d(\omega_k)$ are respectively the coefficients of $\mathrm{e}^{-\mathrm{i}\omega_k\boldsymbol{t}}$ and $\mathrm{e}^{\mathrm{i}\omega_k\boldsymbol{t}}$. The result is plotted figure 2.11.b. Finally, figure 2.11.c, we plot the amplitudes of the best fit initialized at the periods found significant in the discovery paper Lovis et al. (2011b).

The aspects of figures 2.11.a and b is similar, with a clear appearance of the five planets with the highest amplitude (at 5.759, 16.35, 49.74, 122.7 and 2222 days). However, the amplitudes are incorrect (up to ≈ 0.7 m.s^{-1} instead of ≈ 4.5 m.s^{-1}). We also observe a forest of peaks in the one-day region, which will be discussed later. Let us note that figure 2.11.b is slightly cleaner at low amplitudes. In particular between one and ≈ 5 days, the "forest" of peaks has a lower amplitude in the complex case. Secondly, it seems like even though in both cases amplitudes are under-estimated, the pattern of relative amplitudes is closer in the complex case.

Grid spacing and tolerance, first approach

The theory of ℓ_1 minimization offers guarantees of recovery if the columns are not too correlated. As a consequence, we tried to construct the grid so that the correlation is kept low. We also tried random grids and evenly spaced ones in two configurations. First, with a great number of frequency, such that the gap between neighbouring frequencies is very small. Secondly, we explored the strategy of computing the solution of (BP_ϵ) with several coarse grids to limit the correlation between the columns, and average out the solutions. It turned

a)

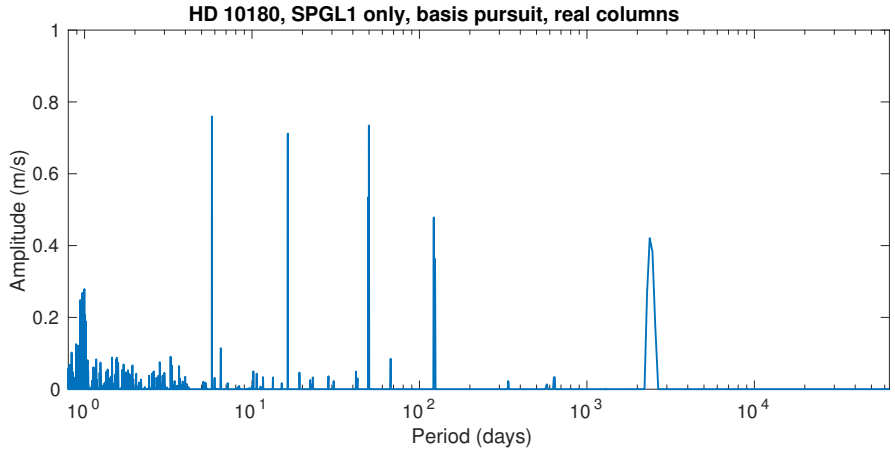

b)

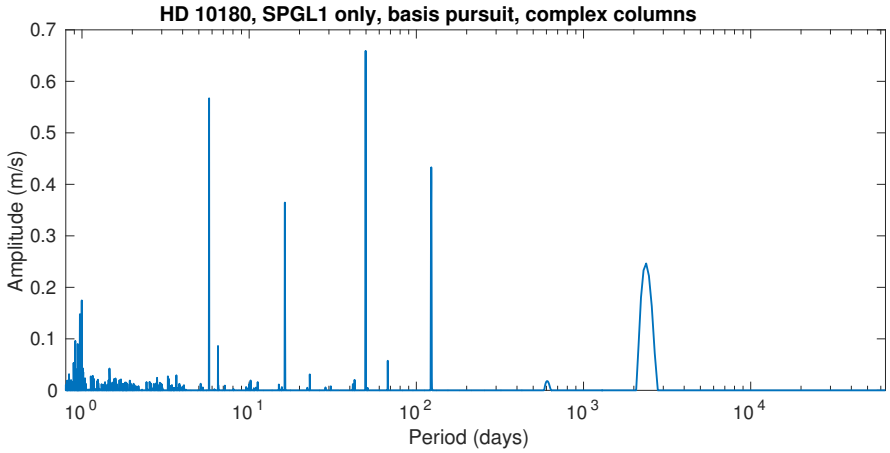

c)

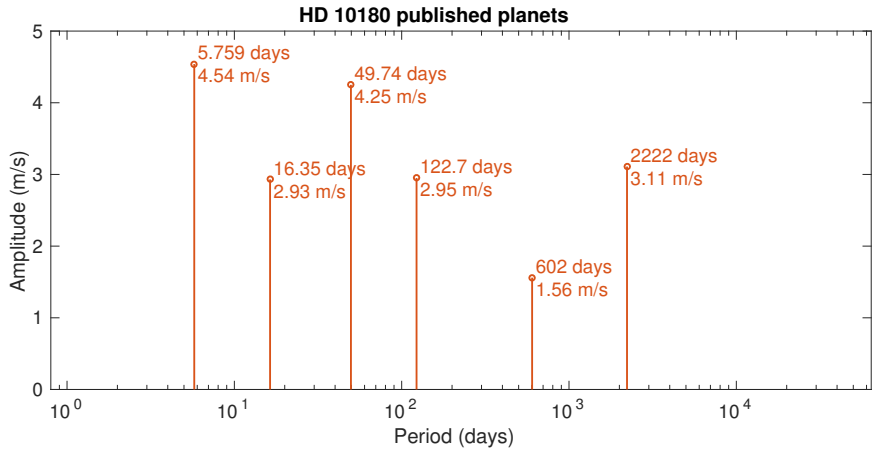

Figure 2.11: Basis pursuit solutions of SPGL1 for HD 10180. a) the columns of $\mathbf{A}$ are $\cos(\omega t), \sin(\omega t)$, $\omega \in \Omega$, b) the columns of $\mathbf{A}$ are $e^{i\omega t}, e^{i\omega t}$, $\omega \in \Omega$ c) Published planets and amplitudes of the best fit.

out that a dense evenly spaced grid performs best, which is also the simplest solution. We are left with the choice of the grid span $[\omega_{\min}, \omega_{\max}]$ and the grid spacing $\Delta\omega$.

The problem of the grid spacing has already been tackled in the Compressed Sensing literature. For instance, (Stoica & Babu 2012) suggests to choose the spacing such that the matrix $\mathbf{M}(\omega) = \mathrm{e}^{\mathrm{i}\,\omega \boldsymbol{t}}\,\mathrm{e}^{\mathrm{i}(\omega+\Delta\omega)\boldsymbol{t}^T}$ has a practical rank equal to one for all $\omega \in \Omega$. The authors do not choose the usual definition of rank since $\mathbf{M}(\omega)$ is in general full rank in the strict sense. However, "most" of its eigenvalues might be close to zero. The practical rank as defined by Stoica & Babu (2012) is the number of eigenvalues of $\mathbf{M}(\omega)$ that are greater than $\lambda_{\max} \times 10^{-3}$, $\lambda_{\max}$ being the maximum eigenvalue. On the other hand, Sahnoun et al. (2012) suggests to select initially only one vector and add columns to $\mathbf{A}$ iteratively.

It was unclear to us whether these criteria help recovering the true frequencies, therefore we adopted another rationale. The idea is to guarantee that the true coefficients $\boldsymbol{x}_t$ are an acceptable solutions, more precisely that they belong to the set $\{\boldsymbol{x}, \|\boldsymbol{y} - \mathbf{A}\boldsymbol{x}\|_{\ell_2} < \epsilon\}$ where $\boldsymbol{y} = \mathbf{A}\boldsymbol{x}_t + \boldsymbol{e}$ with a certain probability. To achieve this, the choices of ϵ and $\Delta\omega$ are interconnected.

Denoting by $\boldsymbol{y}_t$ the signal we would obtain if there was no noise, the random variable $\boldsymbol{y} - \boldsymbol{y}_t$ follows the distribution of the noise. If this distribution is Gaussian of covariance matrix $\mathbf{V}$ and $\mathbf{W} = \mathbf{V}^{-\frac{1}{2}}$ then $\mathbf{W}(\boldsymbol{y} - \boldsymbol{y}_t)$ is Gaussian with covariance matrix identity. As a consequence, $\|\mathbf{W}(\boldsymbol{y} - \boldsymbol{y}_t)\|_{\ell_2}^2$ is distributed according to a χ^2 law with m degrees of freedom, where m is the number of measurements. The tolerance ϵ is selected so that the probability of the true signal to be outside the constraint is α, that is $\Pr\{\|\mathbf{W}(\boldsymbol{y} - \boldsymbol{y}_t)\|_{\ell_2}^2 \geqslant \epsilon^2\} < \alpha$, that is

$$F_{\chi^2_m}(\epsilon^2_{\text{noise}}) = 1 - \alpha \tag{2.44}$$

where $F_{\chi^2_m}$ is the cumulative distribution function (CDF) of a χ^2 distribution with m degrees of freedom.

We now want to ensure that there exists a vector $\boldsymbol{x}$ close to the true coefficients $\boldsymbol{x}_t$ and ϵ that verifies $\|\mathbf{W}(\mathbf{A}\boldsymbol{x} - \boldsymbol{y})\|_{\ell_2}^2 < \epsilon$ that has the same number of non-zero components. If this condition is fulfilled, we ensure that there is a vector with the correct sparsity that is within the bounds of the constraints. We select a positive real number ϵ_{grid} and take a grid spacing to ensure there exists $\boldsymbol{x}$ with same sparsity as $\boldsymbol{x}_t$ such that

$$\|\mathbf{W}(\mathbf{A}\boldsymbol{x} - \boldsymbol{y}_t)\|_{\ell_2} < \epsilon_{\text{grid}}. \tag{2.45}$$

Thus, there exists $\boldsymbol{x}$ with same sparsity as $\boldsymbol{x}_t$

$$\|\mathbf{W}(\mathbf{A}\boldsymbol{x} - \boldsymbol{y})\|_{\ell_2} < \|\mathbf{W}(\mathbf{A}\boldsymbol{x} - \boldsymbol{y}_t)\|_{\ell_2} + \|\mathbf{W}(\boldsymbol{y}_t - \boldsymbol{y})\|_{\ell_2} < \epsilon_{\text{grid}} + \epsilon_{\text{noise}} := \epsilon \tag{2.46}$$

with probability at least $1 - \alpha$.

We show in Hara et al. (2017), Appendix A, that for a grid spacing $\Delta\omega$, we find an appropriate $\boldsymbol{x}$ that verifies

$$\|\mathbf{W}(\mathbf{A}\boldsymbol{x} - \boldsymbol{y}_t)\|_{\ell_2}^2 \leqslant 4m\|\mathbf{W}\|_{2\to 2}^2 \sin^2 \frac{\Delta\omega T}{4} \sum_{j=1}^{p} |c_j|^2 \tag{2.47}$$

where $\|\mathbf{W}\|_{2\to 2}$ is defined as in equation (2.34), and p, c_j are such that $\boldsymbol{y}_t = \sum_{j=1}^{p} c_j \, \mathrm{e}^{\mathrm{i}\,\omega_j \boldsymbol{t}}$, T is the observation timespan. We then select a grid spacing

$$\Delta\omega \leqslant \frac{4}{T} \arcsin \frac{\epsilon_{\text{grid}}}{2\sqrt{\sum_{j=1}^{p} |c_j|^2}\sqrt{\sum_{k=1}^{m} \frac{1}{\sigma_k^2}}}. \tag{2.48}$$

Obviously we do not know the c_j, so we used the approximation $\sqrt{\sum_{j=1}^{p} |c_j|^2} \approx \|\boldsymbol{y}\|_{\ell_2}/\sqrt{m}$ which comes from

$$\|\boldsymbol{y}\|_{\ell_2}^2 = \sum_{k=1}^{m} \left| \sum_{j=1}^{p} c_j \, \mathrm{e}^{\mathrm{i}\,\omega_j t_k} \right|^2 = m \sum_{j=1}^{p} |c_j|^2 + 2\Re \left(\sum_{k=1}^{m} \sum_{j=1}^{p-1} \sum_{i=j+1}^{p} c_j c_i \, \mathrm{e}^{\mathrm{i}(\omega_j - \omega_i) t_k} \right) \tag{2.49}$$

and gives

$$\epsilon_{\text{grid}} = 2\|\mathbf{W}\|_{2\to 2} \|\boldsymbol{y}\|_{\ell_2} \sin \frac{\Delta\omega T}{4} \tag{2.50}$$

We usually took $\epsilon_{\text{grid}} \leqslant \epsilon_{\text{noise}}/3$ and checked visually that the final figure was not too dependent of the parameters by varying them and re-computing the plots.

Normalizing the input

We noticed that the number of iteration is not independent of the units of $\boldsymbol{y}$ and $\mathbf{A}$, which is an undesirable property. SPGL1 stops whenever the estimate of the output $\boldsymbol{x}^\star$ at iteration k, $\boldsymbol{x}_k$, verifies $\frac{|\|\mathbf{A}\boldsymbol{x}_k - \boldsymbol{y}\|_{\ell_2} - \epsilon|}{\max(1, \|\mathbf{A}\boldsymbol{x}_k - \boldsymbol{y}\|_{\ell_2})} < \text{tol}$. By default, tol$=10^{-4}$.

To make the precision of the SPGL1 solver independent from the value of $\mathbf{W}\boldsymbol{y}$, the weighted observations $\mathbf{W}\boldsymbol{y}$ are normed by $\|\mathbf{W}\boldsymbol{y}\|_{\ell_2}$, the columns of the matrix $\mathbf{WA}$ are also normed. Denoting by $\boldsymbol{y}' = \frac{1}{\epsilon}\mathbf{W}\boldsymbol{y}/\|\mathbf{W}\boldsymbol{y}\|_{\ell_2}$ and $\mathbf{A}' = \frac{1}{\epsilon}(\mathbf{W}\boldsymbol{a}_k/\|\mathbf{W}\boldsymbol{a}_k\|_{\ell_2})_{k=1..n}$, where $\boldsymbol{a}_k$ is the k-th column of $\mathbf{A}$. We set in input of the solver:

$$\arg\min_{\boldsymbol{x}\in\mathbb{C}^n} \|\boldsymbol{x}\|_{\ell_1} \quad \text{s.t.} \quad \left\|\mathbf{A}'\boldsymbol{x} - \boldsymbol{y}'\right\|_{\ell_2} \leqslant 1. \tag{2.51}$$

Averaging the output

We now use SPGL1 in the basis pursuit de-noising mode ((BP$_\epsilon$), $\epsilon > 0$) on HD10180 data. The output is used to compute the amplitude $A(\omega_k)$ per period via formula (2.43). Figure 2.12.a shows the result of the computation along with the periods and amplitudes of the published planets. Figure 2.12.b shows the same plot but with enhanced scale on the y-axis. Though the structure of the periods is globally appropriate, the amplitude of the peaks is under-estimated. To circumvent that problem, we average the coefficients of neighbouring frequencies with the following procedure.

Denoting by $\boldsymbol{x}^\star$ the solution to (BP$_\epsilon$), and by $\boldsymbol{x}^\star(\omega)$ the coefficient corresponding to frequency ω (c and d in the notations of formula (BP$_\epsilon$)), we compute for a certain η

$$\hat{\boldsymbol{y}}_\omega(t) = \|\mathbf{W}\boldsymbol{y}\|_{\ell_2} \sum_{\substack{\omega' \in \Omega \\ \omega - \eta \leqslant |\omega'| \leqslant \omega + \eta}} \frac{\boldsymbol{x}^\star(\omega')\boldsymbol{a}_{\omega'}(t)}{\|\mathbf{W}\boldsymbol{a}_{\omega'}(t)\|_{\ell_2}} \tag{2.52}$$

a)

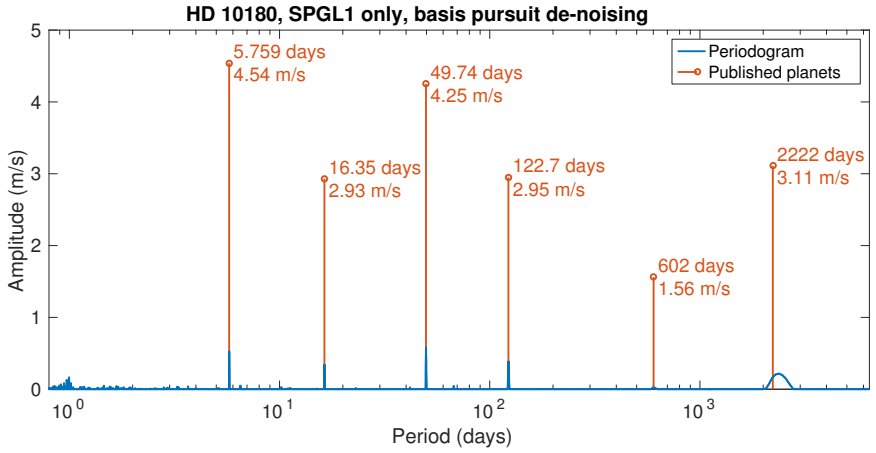

b)

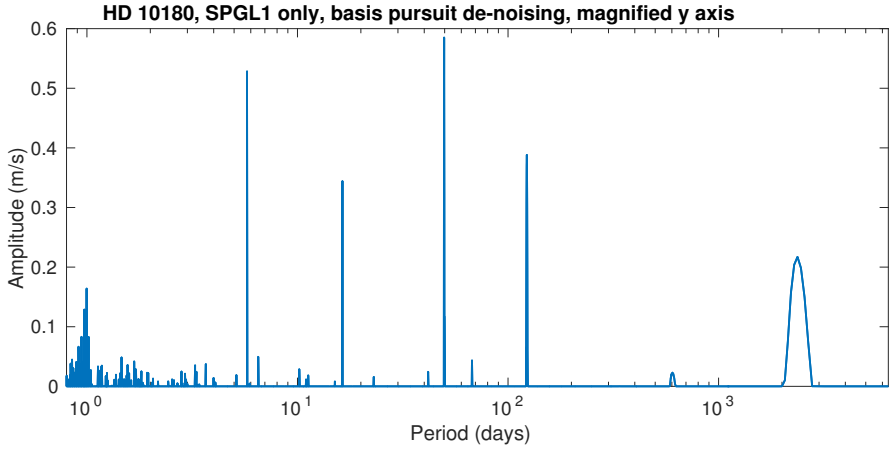

c)

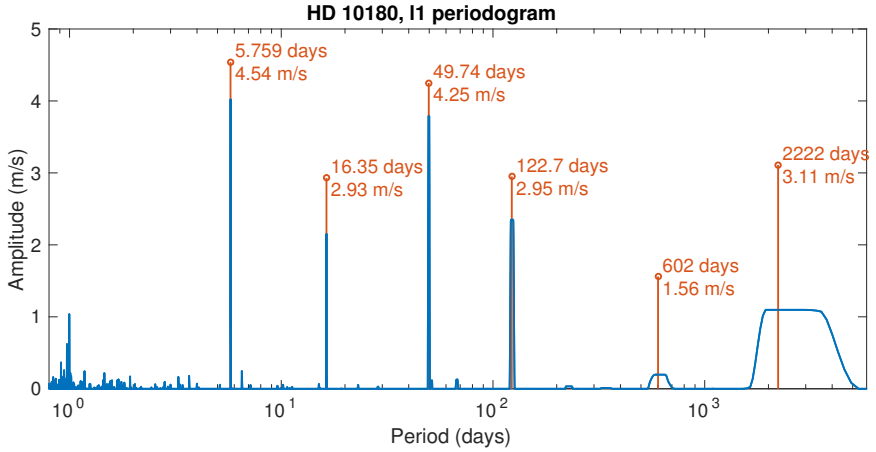

Figure 2.12: Basis pursuit denoising with SPGL1 ($\mathbf{A}$ has complex columns) a) Without normalization of the input of the solver nor averaging. b) Same figure as a) with zoom on the y axis c) "ℓ_1-periodogram" with weighting by matrix $\mathbf{W}$, normalization of the input and averaging of the solution.

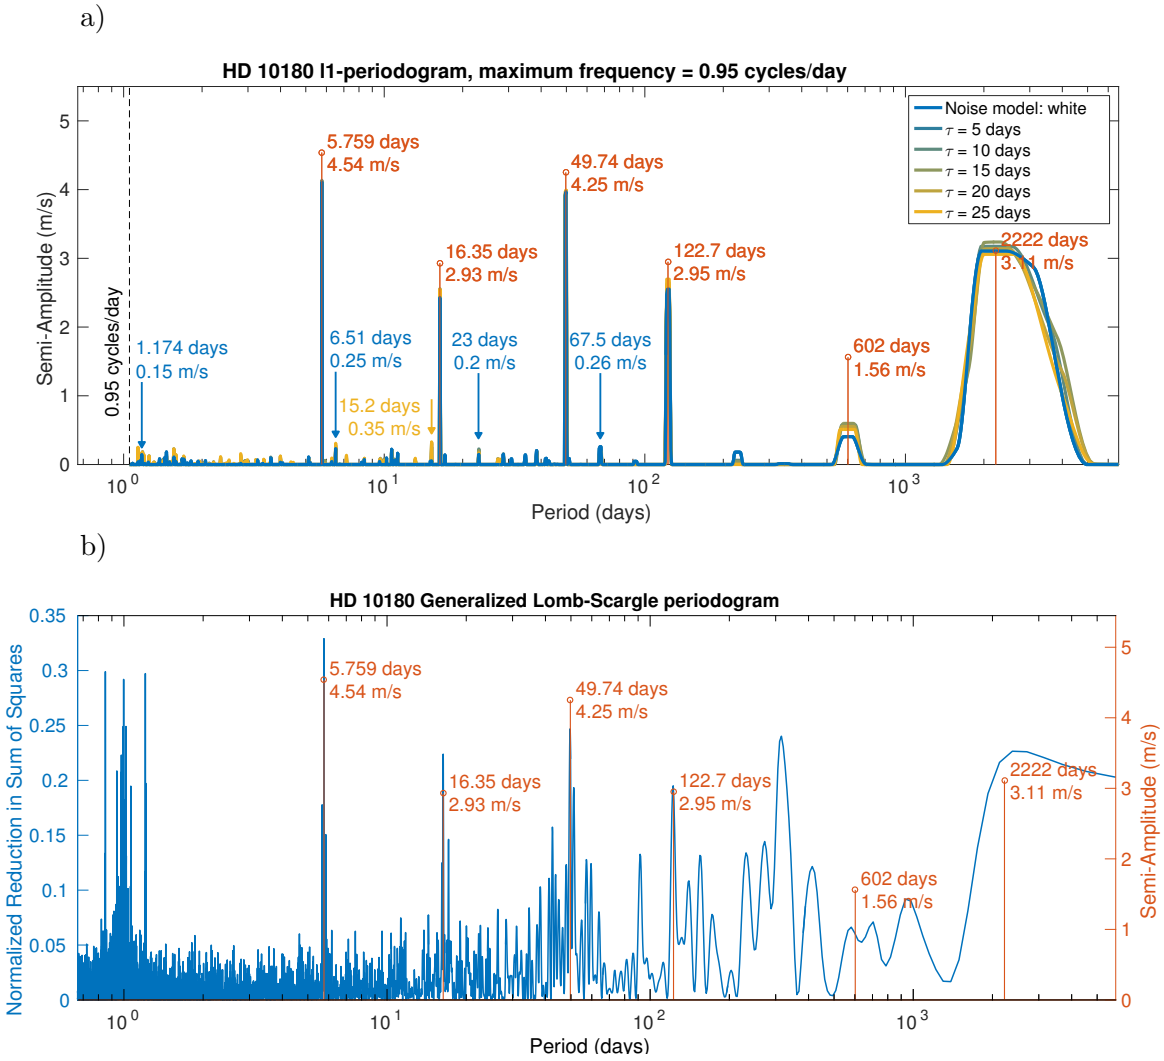

Figure 2.13: a) ℓ_1-periodograms of HD 10180 data set with mean subtracted, b) Generalized Lomb-Scargle periodogram of the same data set. The red stems have the periods and amplitude of published planets. For all the noise model considered for matrix W, $\sigma_W = 0$, $\sigma_R = 1$ m.s^{-1}.

Where $\boldsymbol{a}_{\omega'}(t)$ is the column of A corresponding to frequency ω'. The terms $\|\mathbf{W}\boldsymbol{y}\|_{\ell_2}$ and $1/\|\mathbf{W}\boldsymbol{a}_{\omega'}(t)\|_{\ell_2}$ appear because the columns of $\mathbf{WA}$ and the weighted observations $\mathbf{W}\boldsymbol{y}$ were normalized. The vector $\hat{\boldsymbol{y}}_\omega(t)$, $t = t_1..t_m$ is approximately a sine function, the new estimation of the signal power is:

$$\boldsymbol{x}^\sharp(\omega) = \max_{t_1..t_m} |\hat{\boldsymbol{y}}_\omega(t_k)|. \tag{2.53}$$

Given the regular spacing of periodogram peaks, we can take a constant η equal to a fraction of a typical peak width. We use $\eta = 0.7\pi/T_{\text{obs}}$ where T_{obs} is the observation timespan, but any value between $0.5\pi/T_{\text{obs}}$ and π/T_{obs} is acceptable.

Besides providing better estimates of the main peaks amplitude, the process has also the advantage of bringing up the smaller signals at 602 days. This is a general pattern: averaging the output helps sorting out spurious peaks and true signals that have low amplitude.

Grid span

Now that the grid step $\Delta\omega$ is set, we must select the grid span $[\omega_{\min}, \omega_{\max}]$. To include low frequencies, the grid must contain 0. Since the dictionary would be redundant if we had $\omega_{\min} < 0$ and $\omega_{\max} > 0$, we choose to consider grids of the form $[0, \omega_{\max}]$.

The main issue comes from the nearly daily repetition of the measurements, which leads to a strong correlation of a vector $e^{i\omega \boldsymbol{t}}$ and $e^{i(\omega+\omega_0)\boldsymbol{t}}$ where $\omega_0 \approx 2\pi$. If $\omega_{\max} > \omega_0$, we observe aliases in the one day region. On the other hand, if $\omega_{\max} < \omega_0$ very short period planets such as 55 Cnc e (0.73 days) remain unseen.

Since solving (BP_ϵ) takes only up to a few minutes, the workaround adopted is simply to solve it twice, once with $\omega_{\max} = 3\pi$ to check for short period companions and once with $\omega_{\max} = 1.9\pi$, to avoid to introduce very correlated columns that yield spurious peaks.

To show the effect of the grid span, we compute the ℓ_1 periodogram as in figure 2.13.a but on $[0, 1.9\pi]$. It obviously suppresses the aliases in the one day region and enhances the signal at 602 days. The interpretation of such a feature is that the power of the signal at this period was split between the ≈ 600 days coefficients and their aliases.

Discussion on the algorithm tuning

The reasonings above were first attempts to find an automatic procedure to define the parameters of the algorithm, which performed satisfactorily in general. We believe however that it could be improved. Our procedure, as well as the one of Stoica & Babu (2012), does not take into account all the parameters that have an influence on the performance of the algorithm.

One of the aspects we disregarded was the speed of the algorithm. The smaller the tolerance (ϵ) and the denser the grid ($\Delta\omega$), the lengthier the computation is. Imposing ϵ_{grid} small in comparison of ϵ_{noise} leads to a small $\Delta\omega$, and since the total grid span is fixed, to grids with several hundreds of thousands of frequencies and twice as much columns for $\mathbf{A}$. Even in that configuration, the algorithm takes only up to a few minutes to run, but this can be improved to tens of seconds.

Secondly, we chose the tolerance to bound by a certain α the probability of the true signal to be outside the constraint, or $\Pr\{\|\mathbf{W}(\boldsymbol{y} - \boldsymbol{y}_t)\|_{\ell_2}^2 \geqslant \epsilon^2\} \leqslant \alpha$. However as discussed in section 2.1.4, when the solution $\boldsymbol{x}^\star$ of (BP_ϵ) is non zero it always verifies $\|\mathbf{W}(\mathbf{A}\boldsymbol{x}^\star - \boldsymbol{y})\| = \epsilon$. As a consequence, the algorithm will always select $\boldsymbol{x}^\star$ that has the maximal allowed distance to data, therefore minimal likelihood, which is proportional $\propto e^{-\frac{1}{2}\|\mathbf{W}(\mathbf{A}\boldsymbol{x}^\star - \boldsymbol{y}_t)\|_{\ell_2}^2}$.

If the tolerance is reduced, to ensure the existence of a solution that has the correct ℓ_0 norm, we should take a smaller grid step. However, simulations show that this is unnecessary in general. We found the observation of Stoica & Babu (2012) valid in general, for an observation timespan T_{obs}, taking $\Delta\omega \approx 2\pi/(10T_{\text{obs}})$ is efficient in general. For HD 10180, this reduces the number of frequency from 95,000 to 20,000.

Since we expect the noise not to be represented sparsely on our dictionary, choosing a small tolerance should not affect our inferences in general. The representation of the noise will be spread out onto many atoms, giving forests of peak as in figure 2.11.a and b.

If we are to try several matrices $\mathbf{W}$ to search for correlated noise, then we should not choose $\epsilon = 0$. We showed in Hara et al. (2017), Appendix B (section G.8, page 209) that when the noise is correlated of covariance matrix $\mathbf{V}$ and $\mathbf{W} \approx \mathbf{V}^{-1/2}$, low amplitude signals otherwise buried in the noise can be brought up. It then makes sense to define ϵ not on a confidence intervals, but so that the true $\boldsymbol{x}_t$ is most likely to be in the region $\{\boldsymbol{x}, \|\mathbf{A}\boldsymbol{x} - \boldsymbol{y}\|_{\ell_2} = \epsilon\}$. Thus, $\epsilon = \sqrt{m-2}$ seems to be a reasonable choice since it is the mode of the χ^2 law, and it therefore defines the region were it is most likely to find the true coefficients.

2.3 Developments

2.3.1 Example

Reducing the grid size allows us to include vectors in the dictionary that represent eccentric orbits. On top of atoms of the form $\mathrm{e}^{-\mathrm{i}\omega \boldsymbol{t}}$ and $\mathrm{e}^{-\mathrm{i}\omega \boldsymbol{t}}$ we add $\mathrm{e}^{-\mathrm{i}\nu_e(\boldsymbol{t})}$ and $\mathrm{e}^{\mathrm{i}\nu_e(\boldsymbol{t})}$ where $\nu_e(\boldsymbol{t})$ is a true anomaly computed as a fifth order expansion in eccentricity e. Such atoms are included for a grid of a few eccentricity values (typically $e = 0, 0.2, 0.4, 0.6$). For a full account of eccentricity, the atoms should also depend on the argument of periastron. However, in practice, the inclusion of such atoms was found to improve the representation of some systems. Figure 2.14 we plot the coefficients of all eccentric orbits in the case of HD 10180 with a frequency span of 1.9π. The amplitude of the 600 days signal is brought up and other low-amplitude candidate signals appear.

These signals are not all as steady when the parameters of the algorithm are changed. For instance, when the frequency grid is changed from 1.9π to 2π, we obtain figure 2.14, the 92 days signal vanishes. To interpret this, we use the notion of restricted isometry constant defined in equation (2.33). We take a frequency ω and the best fit frequencies of the six planets found so far, $(\omega_1, ...\omega_6)$ and form the matrix $\mathbf{M}(\omega) = [\mathrm{e}^{-\mathrm{i}\omega_1 \boldsymbol{t}}, \mathrm{e}^{\mathrm{i}\omega_1 \boldsymbol{t}}, ..., \mathrm{e}^{-\mathrm{i}\omega_6 \boldsymbol{t}}, \mathrm{e}^{\mathrm{i}\omega_6 \boldsymbol{t}}, \mathrm{e}^{-\mathrm{i}\omega \boldsymbol{t}}, \mathrm{e}^{\mathrm{i}\omega \boldsymbol{t}}]$. We compute the eigenvalues of $\mathbf{G}(\omega) = (\mathbf{WM}(\omega))^*(\mathbf{WM}(\omega))$. Since $\mathbf{G}(\omega)$ is symmetric positive definite, it has non negative eigenvalues $(\lambda_i)_{i=1..14}$ of minimal and maximal values $\lambda_{\min}$ and $\lambda_{\max}$. We then compute $\delta(\omega) = (1-\lambda_{\max}/\lambda_{\min})/(1+\lambda_{\max}/\lambda_{\min})$. We plot δ as a function of period in figure 2.14.d.

This function goes to one each time $\omega = \omega_i$ since the matrix becomes degenerate and $\lambda_{\max}/\lambda_{\min}$ tends to infinity. It also displays a peak at 92 days, which indicates a strong correlation between signals at this position and the existing. In figure 2.14.c we show another example of algorithm tuning ($\epsilon = 2\epsilon_{\text{noise}}$) where a peak at 1.206 days appears, this one corresponds to a value of δ of 0.95. These examples, along with other ones we do not present for the sake of brevity, show that plotting δ could be a useful indicator of where aliases could be in cases with multiplanetary signals. We however acknowledge that the use of δ , though motivated by theoretical result and analogy with the spectral window, is heuristic and should be therefore used with care. In particular, it is insensitive to the amplitudes of the signals. To take these into account, a useful indicator of where aliases could be is more simply the (generalized) Lomb-Scargle periodogram itself. Figure 2.13.b indeed displays peaks at 92 and 1.206 days.

All these examples show that the estimation of low amplitude signals depend on the choice of the grid and should be taken with care. From the theoretical point of view, Reeves & Gastpar (2009) compute the optimal error rate of support recovery with information theory tools. They show that it is asymptotically equal to a trivial thresholding (assimilable to a

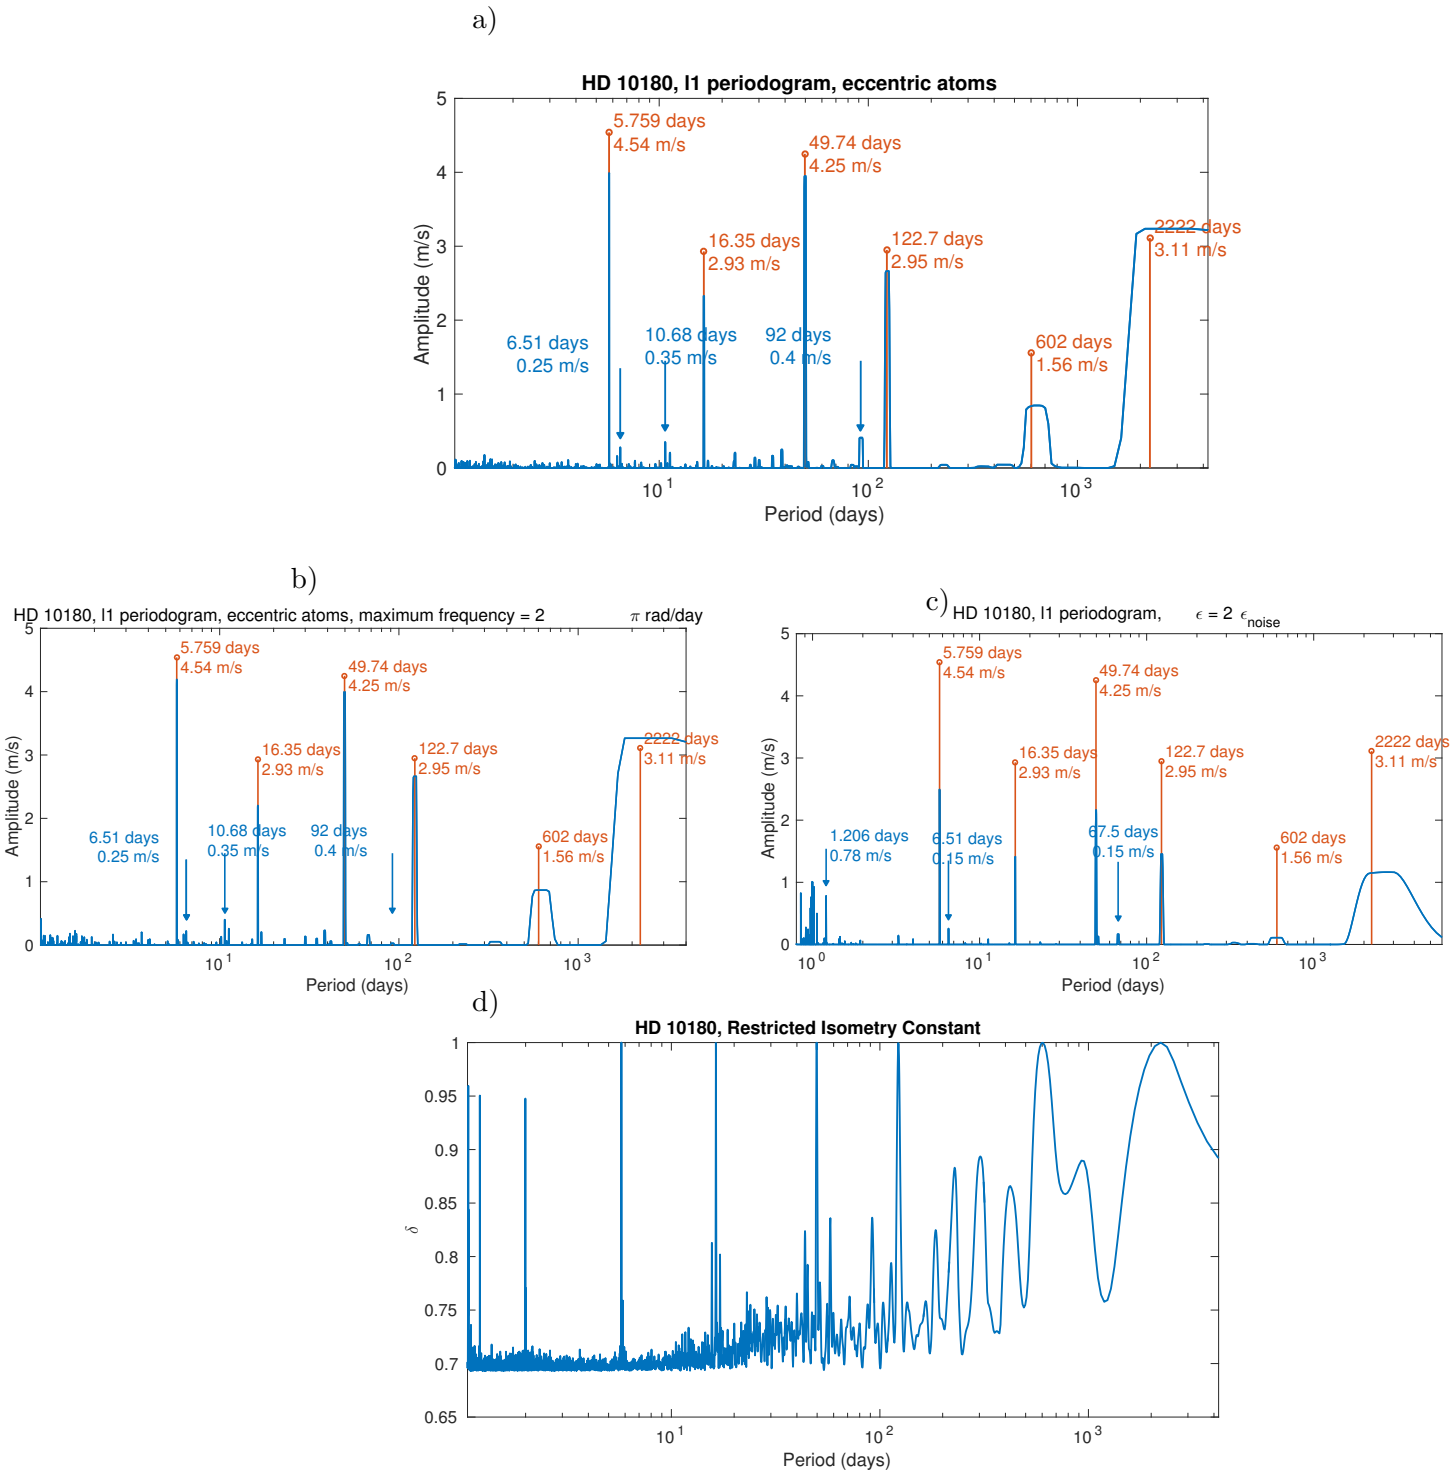

Figure 2.14: ℓ_1-periodogram of HD 10180 with eccentric dictionary elements.

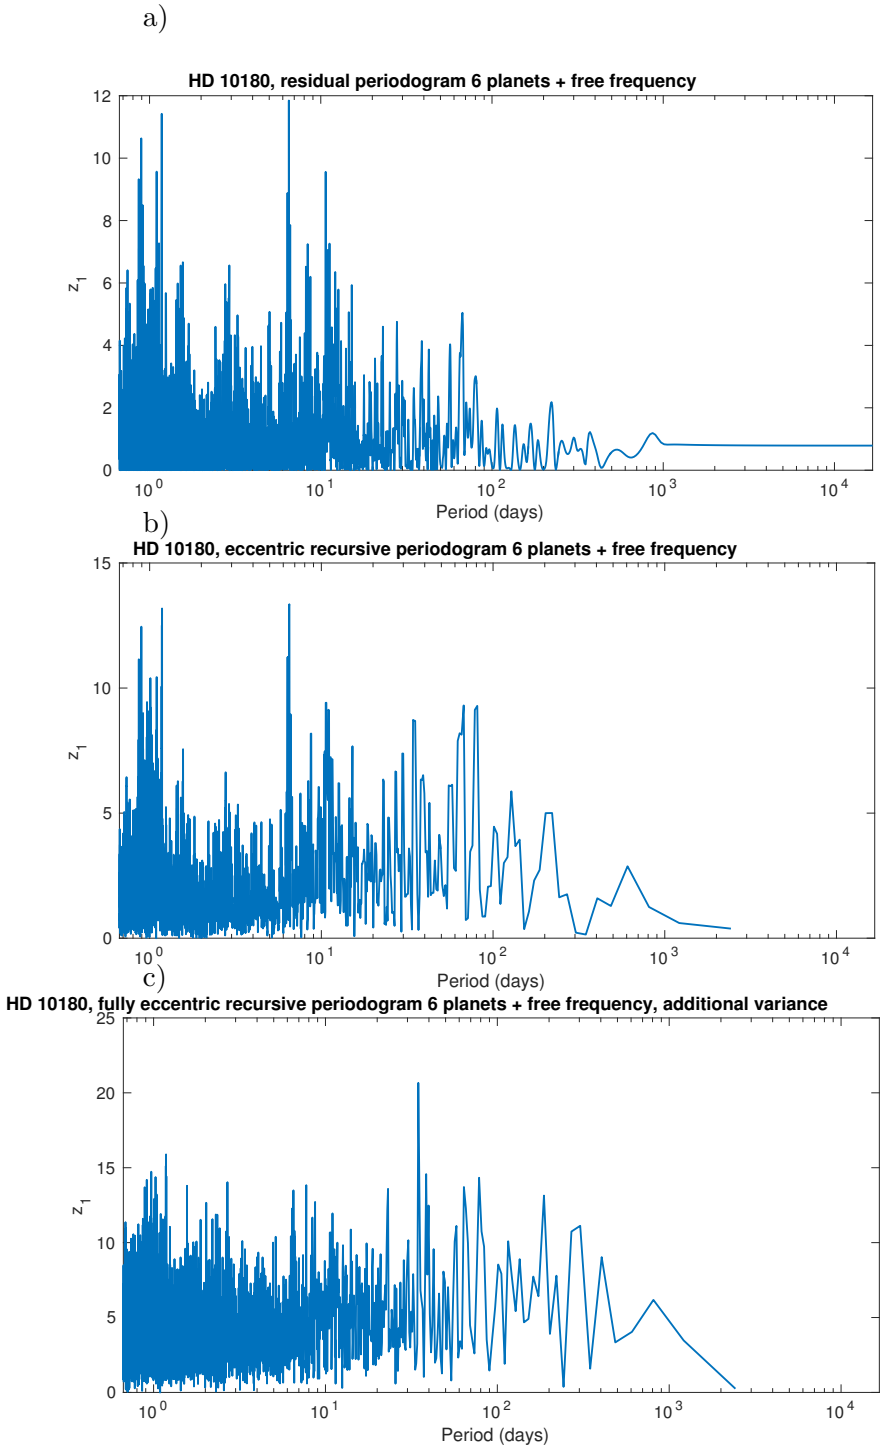

Figure 2.15: ℓ_1-periodogram of HD 10180 with eccentric dictionary elements.

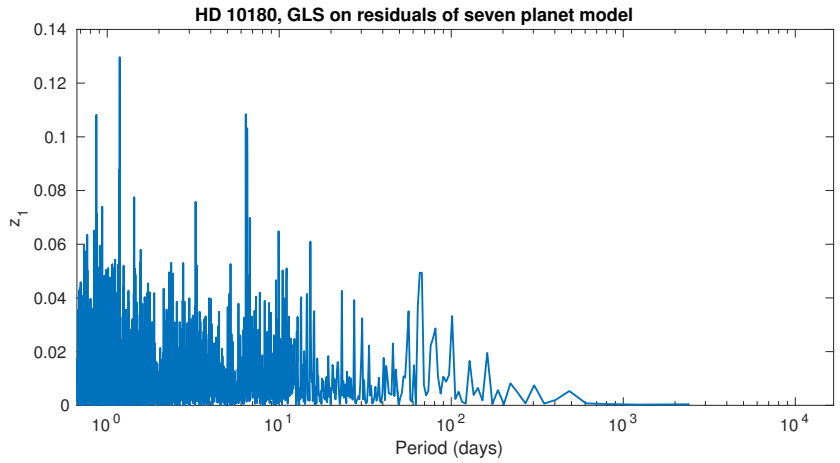

Figure 2.16: GLS of the residuals after seven signals are removed.

less sophisticated version of orthogonal matching pursuit) respectively in the limit of 0 and infinite signal to noise ratio (amplitude of the signal divided by noise)[4]. Therefore, we should expect that the ℓ_1 minimization will not help us much to dig in low amplitude signals. The ℓ_1 minimization will greatly increase our confidence in the high amplitude peaks, to dig in the noise we can move to periodogram methods.

For the search of exoplanets, astronomers have developed tools that push one step further the idea of orthogonal matching pursuit with the following principle. At each trial frequency of the periodogram, the previously found signals are re-adjusted, not only when the maximal correlation between residuals and the dictionary is found. Let denote by $\omega_1, ..., \omega_k$ the frequencies that have been stored up to iteration k. At the trial frequency ω we can compute

$$z_1(\omega) = \alpha \frac{\chi^2_{k+1} - \chi^2_k}{\chi^2_k} \tag{2.54}$$

where χ^2_k and χ^2_{k+1} are the χ^2 of residuals respectively when the models at period $\omega_1, ..., \omega_k$ and $\omega_1, ..., \omega_k, \omega$ are fitted, α is a constant that will be made explicit later on. Depending on what is intended by "fitted", different answers are obtained.

[4] More precisely, Reeves & Gastpar (2009) consider signals of the form $\boldsymbol{y} = \mathbf{A}\boldsymbol{x}_t + \boldsymbol{\epsilon}$ where $\boldsymbol{\epsilon}$ is a Gaussian independent noise. They suppose that the number non zero components of $\boldsymbol{x}_t$ is known, that the support is uniformly distributed over all support of size s and that the amplitudes of the coefficients is distributed with an independent Gaussian law. The signal to noise ratio S is defined as the ratio of the variances of the signal amplitude and of the noise $\boldsymbol{\epsilon}$. They consider a fraction of errors α on the indices of the support and the optimal estimator that minimized the probability to have an error on the support $> \alpha$, $\delta_{\text{opt}}(\boldsymbol{y}; \alpha, S)$ that minimizes the probability that the error is above α. They consider the trivial thresholding support estimator $\delta_{\text{TH}}(\boldsymbol{y})$, which consists in taking the s largest elements of $\mathbf{A}^T\boldsymbol{y}$. They show the following theorem.

Theorem ((Reeves & Gastpar 2009)). *Let* $\mathbf{A}$ *and* $\boldsymbol{y}$ *be fixed. If* $\delta_{\text{TH}}(\boldsymbol{y})$ *is unique, then there exists* $S_{\mathbf{A},\boldsymbol{y}} > 0$ *such that for all* $S < S_{\mathbf{A},\boldsymbol{y}}$*, we have* $\delta_{\text{opt}}(\boldsymbol{y}; \alpha, S) = \delta_{\text{TH}}$

In other words, as long as the matric $\mathbf{A}$ has non degenerate rows, one can always find a signal to noise ratio above $S_{\mathbf{A},\boldsymbol{y}}$ above which the thresholding estimator and the optima one estimate exactly the same subset. The first step of an iterative method is to find the maximum of a periodogram is assimilable to finding the maximum correlation with a column of $\mathbf{A}$ (this is what is done in frequency analysis).

The first form of the model we consider is

$$\boldsymbol{y}^{(k+1)}(\boldsymbol{t},\omega) = A\cos(\omega\boldsymbol{t}) + B\sin(\omega\boldsymbol{t}) + \sum_{i=1}^{k} A_i\cos(\omega_i\boldsymbol{t}) + B_i\sin(\omega_i\boldsymbol{t}) + \mathbf{M}\boldsymbol{\theta} \tag{2.55}$$

where $\mathbf{M}$ is a $m \times p$ matrix and $\boldsymbol{\theta}$ is a $p \times 1$ vector that account for non-planetary effects. This might include a constant, a linear or quadratic trend, etc. When only the constants $A, B, (A_i, B_i)_{i=1..k}$ are fitted, the false alarm probability has been computed by Baluev (2008) under the hypothesis that $\omega_1 ... \omega_k$ and M truly are in the data[5]. For $\alpha = m - 2k - p$ one obtains the definition (2) of Baluev (2008).

We assess the statistical significance of the maximum of the periodogram in the following manner. Let us denote by $Z_{\max}$ and $z_{\max}$ the maximum of the periodogram respectively as a random variable under the hypothesis that the signal has k planets and the observed value of the maximum. We compute the p-value $\Pr\{Z_{\max} \geqslant z_{\max}|\text{k planets}\}$ with formula (5) of Baluev (2008).

We represent z_1 as a function of frequency in figure 2.15.a on the HD10180 data sets were the frequencies $\omega_1 ... \omega_6$ are initialized at the best fit frequencies of the six published planets. We find a maximum of the periodogram at 6.51 days (11.85) and a close peak at 1.178 days. These were already found in the discovery paper Lovis et al. (2011b). Using formula (5) of Baluev (2008) we find a false alarm probability of ≈ 0.029, which is approximately twice the value of 0.014 found in Lovis et al. (2011b). The discrepancy might be explained by the fact that the formula we use is an upper bound of the false alarm probability. It might also originate from the fact that 0.014 is obtained by simulations and depends on the choice of the noise level. With the values set in Lovis et al. (2011b), the seven planet model yields a reduced χ^2 of 1.5>1, which indicates with high probability an under-estimated noise level. As noted by Lovis et al. (2006), 6.51 and 1.178 are aliases of each other. To evaluate the likelihood of both options, the stability of the system is explored in both cases. Only the system with a planet at 1.178 days is found to be stable. Note that a signal at 10.68 days also appear, which also the case on the ℓ_1 periodogram.

A slightly more sophisticated version assumes a model

$$\boldsymbol{y}^{(k+1)}(\boldsymbol{t},\omega) = \boldsymbol{y}_{\text{Kep}}(\boldsymbol{t},\omega) + \sum_{i=1}^{k} \boldsymbol{y}_{\text{Kep}}(\boldsymbol{t},\omega_i) + \mathbf{M}\boldsymbol{\theta} \tag{2.56}$$

where $\boldsymbol{y}_{\text{Kep}}(\boldsymbol{t},\omega_i)$ indicates a Keplerian model initialized at a circular orbit at frequency ω_i. When $\boldsymbol{y}_{\text{Kep}}(\boldsymbol{t},\omega)$ is restricted to a circular orbit, the periodogram so obtained is the recursive periodogram suggested by Anglada-Escudé & Tuomi (2012). This one is plot figure 2.15.b. Again, the signals at 6.51 and 1.178 days do appear. Interestingly enough, when zooming on the latter, it is found that it is split in two and a peak at 1.182 appears.

This finding motivates us to do another analysis: we compute the recursive periodogram but allow a Keplerian fit for the trial frequency. We plot the result in figure 2.15.c. A period at 34.69 days stands out with an eccentricity of 0.7. When subtracting that signal – which could be due to stellar rotation – and computing the periodogram of residuals (see figure 2.16), the 1.1816 periodicity clearly stands out. The peak at 6.51 also splits and the

[5]In fact it has been computed in a more general setting: the comparison of a linear model and this linear model plus a frequency.

maximum now occurs at 6.39 days. Unfortunately, the amplitude of the peaks does not allow to conclude on the nature of the signal.

The purpose of this analysis is first to introduce the use of the restricted isometry constant as an indicator of the correlation between potential signal, secondly, to show an example of an ℓ_1-periodogram with Keplerian atoms and finally, to demonstrate that interesting features can be unveiled by a residual/recursive approach as a complement of the ℓ_1 minimization method. A more in depth study with an appropriate Keplerian periodogram is left for future work.

2.3.2 Statistical tests

To compare the performances of the ℓ_1 minimization with those of an iterative scheme, we test both methods on simulated data sets. We consider the measurement times of GJ 876 and select the 50 or 100 first dates. We consider two noise levels: 10 cm/s and 2 m.s^{-1} (high and low signal-to-noise ratio). Then we generate 2, or 4 planets in a circular orbit (i.e. sine waves), all with semi-amplitude 3 m.s^{-1}, uniformly distributed phases and periods distributed uniformly in $\log P$. This gives $2 \times 2 \times 2 = 8$ simulation configurations, on each of which 100 systems are generated.

The methods are used in the following setting. In the iterative case, periods are searched with the Generalized Lomb-Scargle periodogram (Zechmeister & Kürster 2009a). When a frequency is found, all the frequencies found so far are re-fitted non-linearly with a sum of sinusoidal model plus a constant. We compute equation (2.54) with χ^2 of models with k and $k+1$ planets. The false alarm probability is computed with the formula (5) of Baluev (2008). When the maximum of the periodogram of the residuals corresponds to a non-significant period, the algorithm stops. Since the periods are fitted non linearly, the FAP calculation of Baluev (2013b) would have been more appropriate, but we stuck to the choice made for Hara et al. (2017).

In the ℓ_1 minimization case, the k frequencies corresponding to the peaks ordered in decreasing amplitude are fitted non-linearly as well as the $k+1$ ones. The k and $k+1$ models are compared with the same formula as the iterative method, until a non significant peak is found. As a consequence, the only difference in the FAP assessment between the ℓ_1 minization and the iterative method lies in the selected peaks.

The performance is evaluated as in section 2.1.2. At iteration j, the algorithm selects a frequency ($\widehat{\omega}_j$). We then compute $d_j = \min_{k=1..n_p} |\widehat{\omega}_j - \omega_k|$, that is the smallest distance between $\widehat{\omega}_j$ and a true frequency. The algorithm stops at iteration j_f. If $j_f < n_p$, we missed a signal. By convention, we then take $d_j = -1$ for $j = j_f + 1..n_p$ (label "Missed" on the x axis). If $j_f > n_p$, spurious signals have been selected, and we pose $d_j = -2$ (label "False" on the x axis). We plot the cumulative distribution function of errors of the frequencies retrieved at each iteration. It is considered that the frequency retrieved are incorrect when a frequency is more than $4 \times 2\pi/T_{\text{obs}}$ away from any of the true frequency, which is the measured width of the periodogram peak. When two frequencies are further apart it means they are not in the same local minimum. We indicate the $4 \times 2\pi/T_{\text{obs}}$ threshold with a dotted black line.

We do the experiment twice. First, the FAP threshold is set to one: we accept n_p frequencies retrieved by the algorithms where n_p is the number of planets injected (2 or 4). With this tuning, we see what are the output frequencies with the algorithms, regardless of the statistical significance assessment. The cumulative distribution functions are shown in figures 2.17

and 2.18. The procedure is then performed as described above with a FAP threshold equal to 10^{-3}, which is a common threshold for exoplanet detection. On each figure, we show a zoom of the CDF close to 0. The iterative method is labelled OMP for "'Orthogonal Matching Pursuit'. Note that the non-linear minimization step is not included in the original algorithm, but since the spirit is very close we stick to this label.

Let us first consider figure 2.17: n_p selected frequencies retrieved when n_p are injected. The signals have a 3 m.s^{-1} amplitude and the noise level is 0.1 m.s^{-1}. The plots on the left side are the result of the ℓ_1 minimization and on the right side of the iterative method. The plots are ordered from top to bottom from easiest to most challenging setting (row 1: 100 measurements, 2 signals, row 2: 100 measurements, 4 signals, row 3: 50 measurements, 2 signals, row 4: 50 measurements, 4 signals).

On row one (figures 2.17.a and b), both algorithms perform well, retrieving the correct frequencies except one error for the iterative method on first iteration. When doubling the number of signals (figures 2.17.c and d), the differences of performance start appearing. The ℓ_1 minimization method still has no mistakes, while $\approx 8\%$ of frequencies obtained at first iterations and $\approx 1-3\%$ of subsequent iterations are incorrect. When halving the number of measurements (figures 2.17.e and f), the ℓ_1 minimization still performs well, with only one frequency out of a hundred slightly off the error threshold. Finally, in figures 2.17.g we see that in more complicated case, 50 measurements and 4 signals, the ℓ_1 method starts showing some mistakes: $\approx 4\%$ on the highest peak found $\approx 19\%$ on third and fourth iteration. This still outperforms the iterative method, which has an error rate from $\approx 25\%$ to $\approx 45\%$ percents (first and fourth iteration respectively, see 2.17.h). Note that the error rate is higher as the iteration proceed, which is inverted when the number of measurent is 50 (figure 2.17.d). This suggests that when there are only 50 measurements, subtracting a spurious signal distorts the residuals to a greater extent and modifies the spectral content of the signal, while when there are more measurements when a spurious signal is fitted, the energy of a few true signals is absorbed, but the remaining true peaks that are not correlated with the spurious signal are not affected by the fitting process and are visible in the next iterations.

In the low signal to noise regime (figure 2.18), the differences in performance become less clear. In both figures 2.18.a and b and figures 2.18.e and f, the error rate is less on the first peak for the ℓ_1 minimization and higher on the second peak. The difference in performance starts appearing when the number of signals is doubled (figures 2.18.c and d and figures 2.18.g and h). In the 100 measurements case, the ℓ_1 method shows better performances on the three first peaks selected and on all peaks for 50 measurements. The fact that the performance of the two methods gets closer as the noise increases is compatible to the result of (Reeves & Gastpar 2009), which was mentioned in the previous section page 88 (footnote).

Independenly of the method, there is a clear degradation of the performance between the 2 and 4 signal cases for 50 measurements. This invites to be suspicious about periodogram peaks at high noise level. Even if the signal is statistically significant in the sense that the white noise hypothesis can be rejected, the principal peak and subsequent ones can be spurious. This is particularly visible in figure 2.17.h), where the noise standard deviation is thirty times less than the amplitude of the signal.

We now move to the case where the FAP threshold is 10^{-3}. Figures 2.19 and 2.20 represent the same quantities as figures 2.17 and 2.18. Let us first note that there are only a very few occurrences of "false" frequencies detected (when at least n_p+1 frequencies are retrieved).

Furthermore, the errors are much less common than in the case where all the signals are accepted, but this originate from the fact that much more signals are rejected. Even in the high signal to noise ratio regime, for 50 measurements and 4 signals 2.19.g and h, almost all signals are rejected but $\approx 20\%$ and $\approx 30\%$ of the first peaks (respectively in the ℓ_1 minimization and iteratve methods). We interpret this feature by the fact that the definition of false alarm probability that we use to accept a planet at iteration $k+1$ (equation (2.54)) is scaled by the χ^2 of residuals with k planets. The more planets there are, the less relative improvement there is in the fit when only one planet is fitted.

Note that figure 2.19.d is not very different from figure 2.17.d. When there are 100 measurements, the signals are still significant with the metric (2.54), but the errors due to aliasing are still present. In those cases, the ℓ_1 minimization outperforms iterative methods, as expected.

As a conclusion of these numerical experiments, we highlight the following points: ℓ_1 minimization exhibits better performances than iterative search in general. As the signal to noise ratio decreases, the difference becomes less apparent, yet the ℓ_1 minimization still shows slightly better results. When using a common FAP threshold for detection of exoplanets, selection of spurious peaks can have statistical significance. The false alarm probability definition seems also to lead to missed detections when there are many signals in the data. It should be noted that false alarm probabilities based on combination of signals as in Baluev (2013) and Baluev (2013a) may be more appropriate and will be considered in future work.

2.4 Return on the Radial Velocity Fitting Challenge

2.4.1 Presentation

The radial velocity fitting challenge was organized by Xavier Dumusque to test existing RV data analysis methods on simulated data sets dominated by stellar activity. The simulated systems were multiplanetary with planets of semi-amplitude $\lesssim 5$ m.s^{-1} and the activity was simulated with SOAP 2.0 software (Dumusque et al. 2014), presented in section 1.2.5, along with time series of the bisector span and $\log R'_{HK}$. Red noise with a correlation of a few days was also added to the signal. The full method of data generation is presented in Dumusque (2016).

Fifteen data sets were generated in various situations and submitted to eight teams who tried to retrieve the correct signals and were ranked as function of the number of true positives and false positives. In the following sections we present our first approach, which was not very efficient, though we obtained one of the lowest rate of false positives. The new approach is presented section 2.4.3.

2.4.2 First approach

Our first approach was to include in the dictionary columns supposed to represent the activity of the star and some aiming at represent the planets. To obtain a representation of the planets, we chose $\mathrm{e}^{-\mathrm{i}\omega \boldsymbol{t}}$ and $\mathrm{e}^{\mathrm{i}\omega \boldsymbol{t}}$ for ω lying in a certain grid (typically made of $n_\omega = 10^5$ frequencies). To represent the activity, we used frequency filters on the FWHM, bisector span and $\log R'_{HK}$. The frequency filters were obtained by projecting onto orthogonal polynomials of the time, higher degrees corresponding to higher frequencies. We considered the m vectors $(\boldsymbol{t}^k)_{k=0..m-1}$ and separated them in c classes: $C_1 = (\boldsymbol{t}^k)_{k=0..j_1}$, $C_2 = (\boldsymbol{t}^k)_{k=j_1+1..j_2}$ and so on. We then

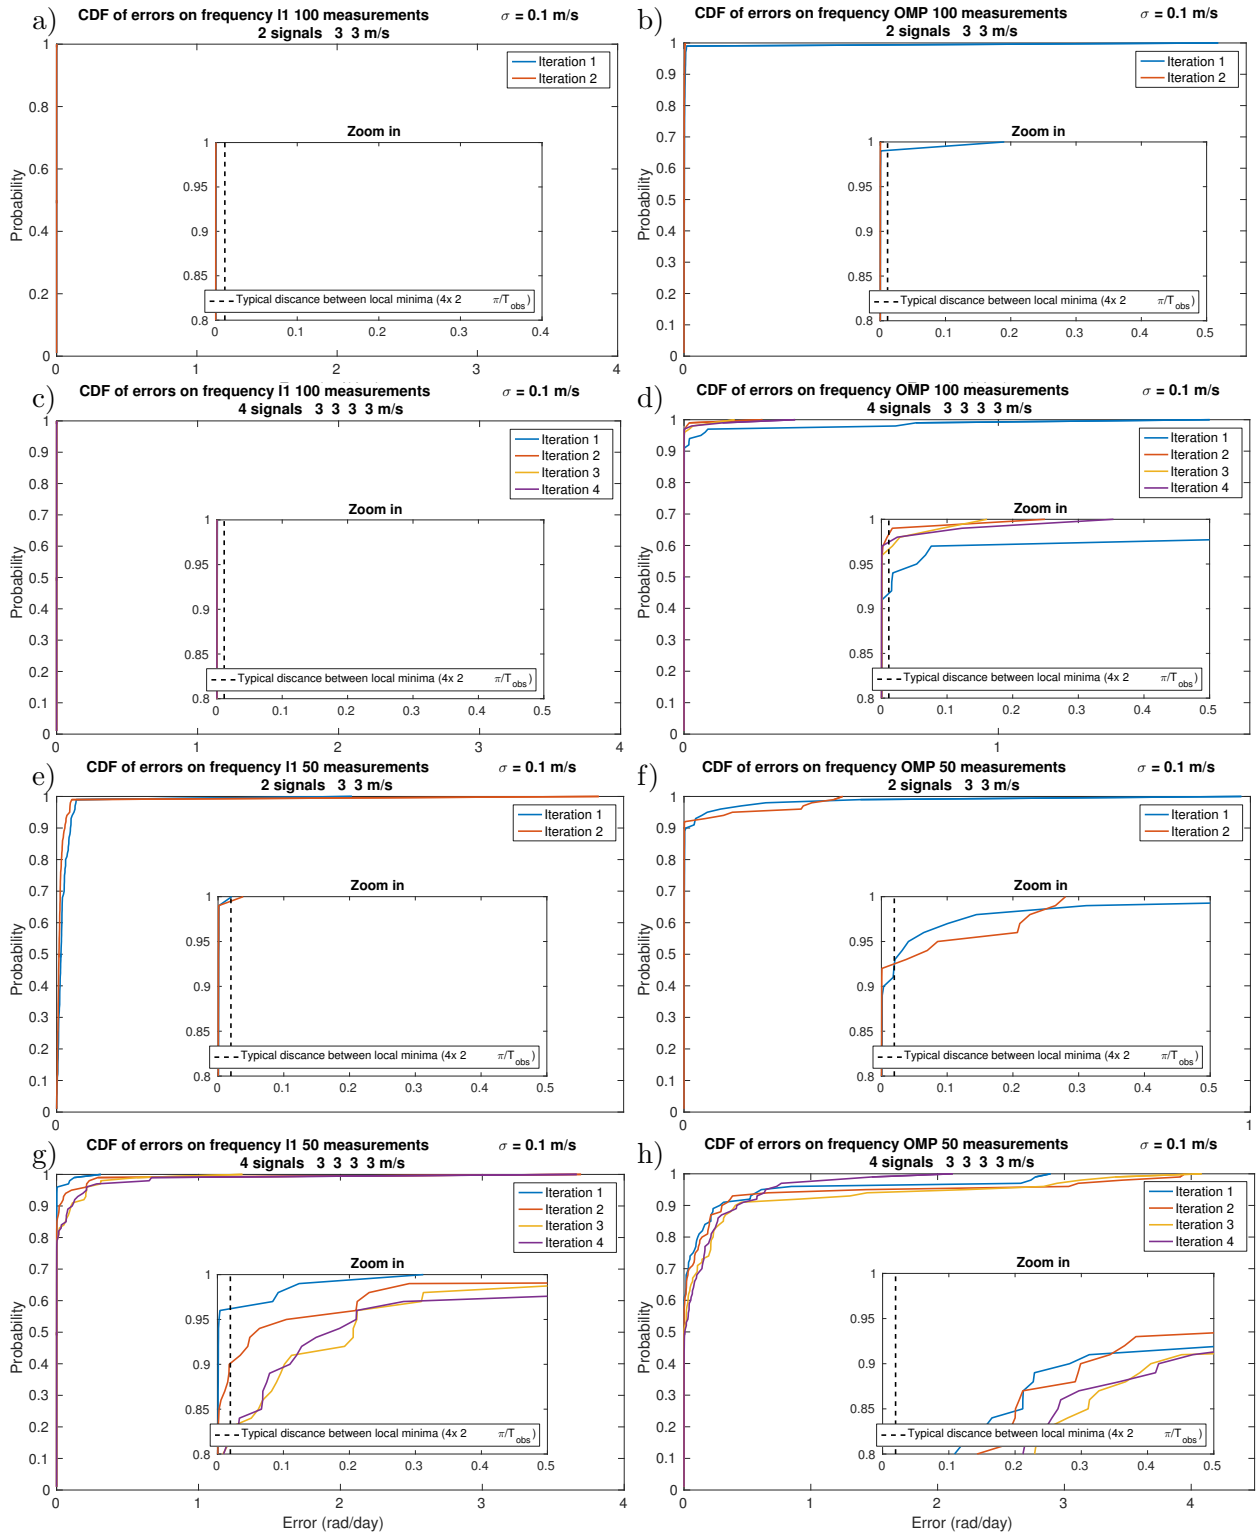

Figure 2.17: Cumulative distribution functions of errors, FAP threshold = 1.

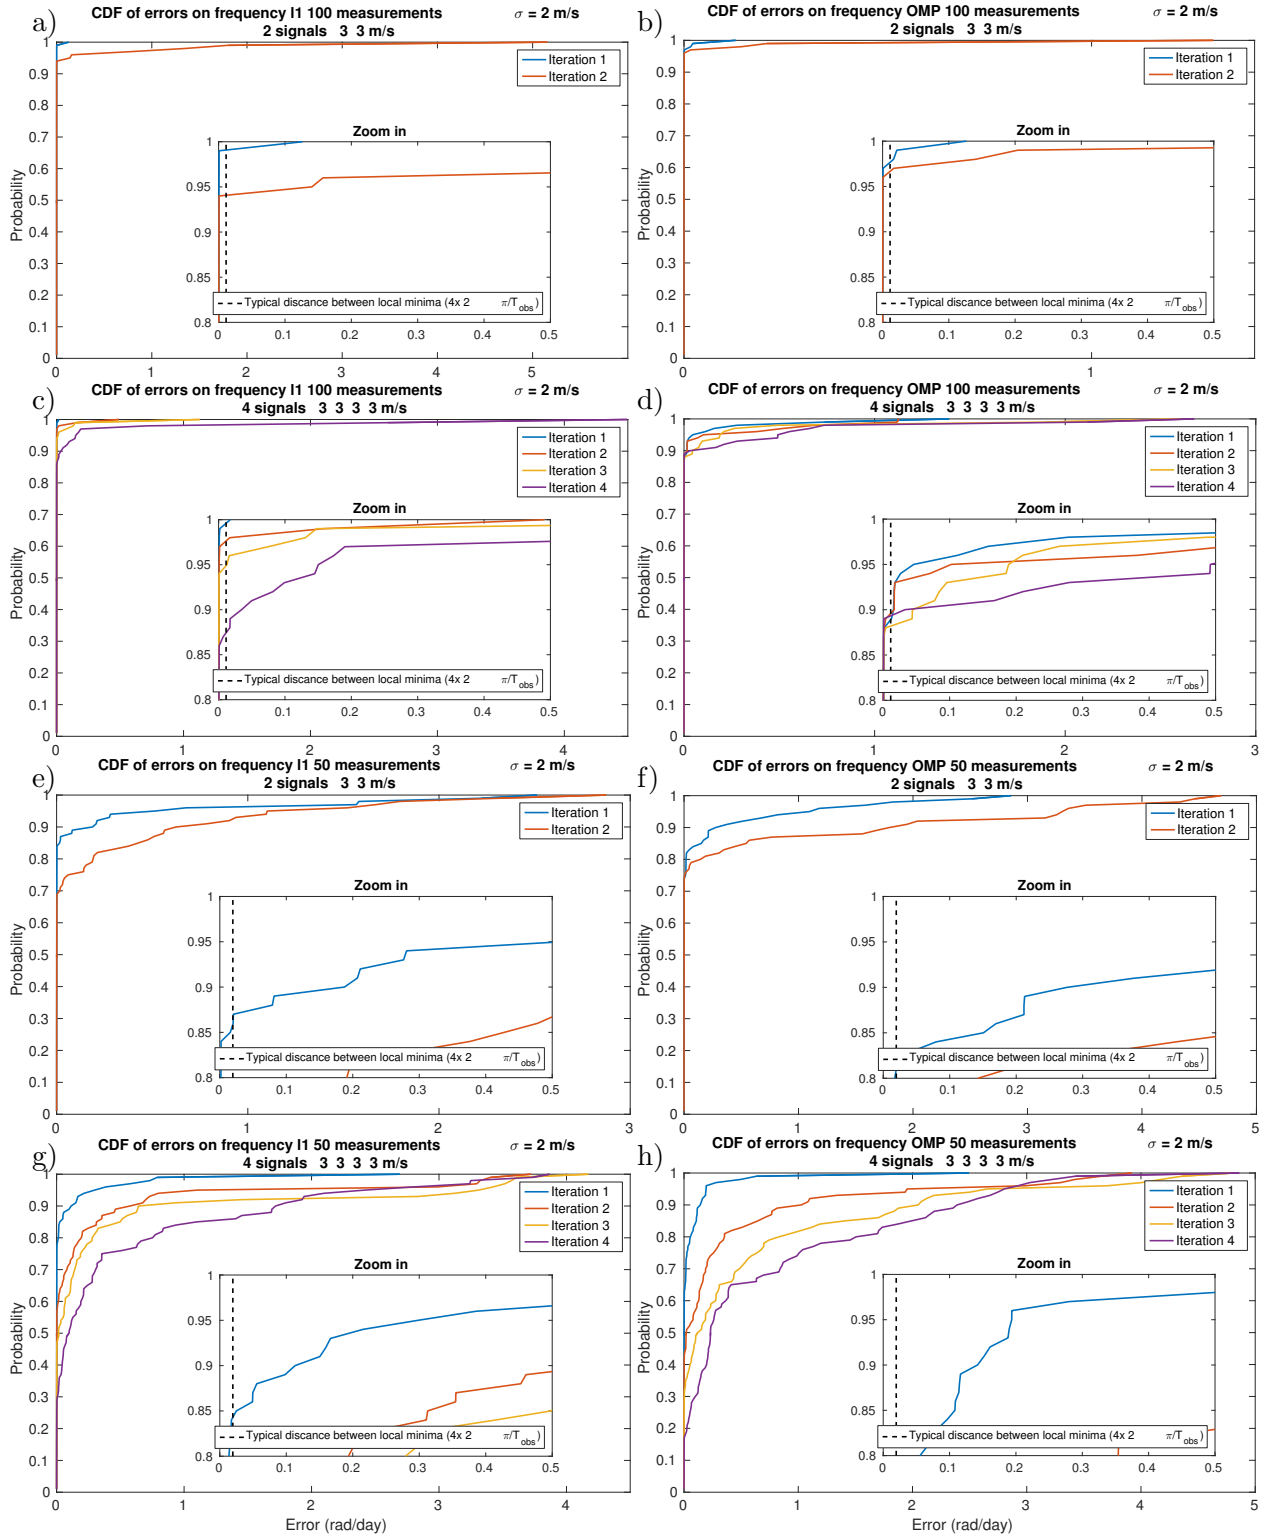

Figure 2.18: Cumulative distribution functions of errors, FAP threshold = 1.

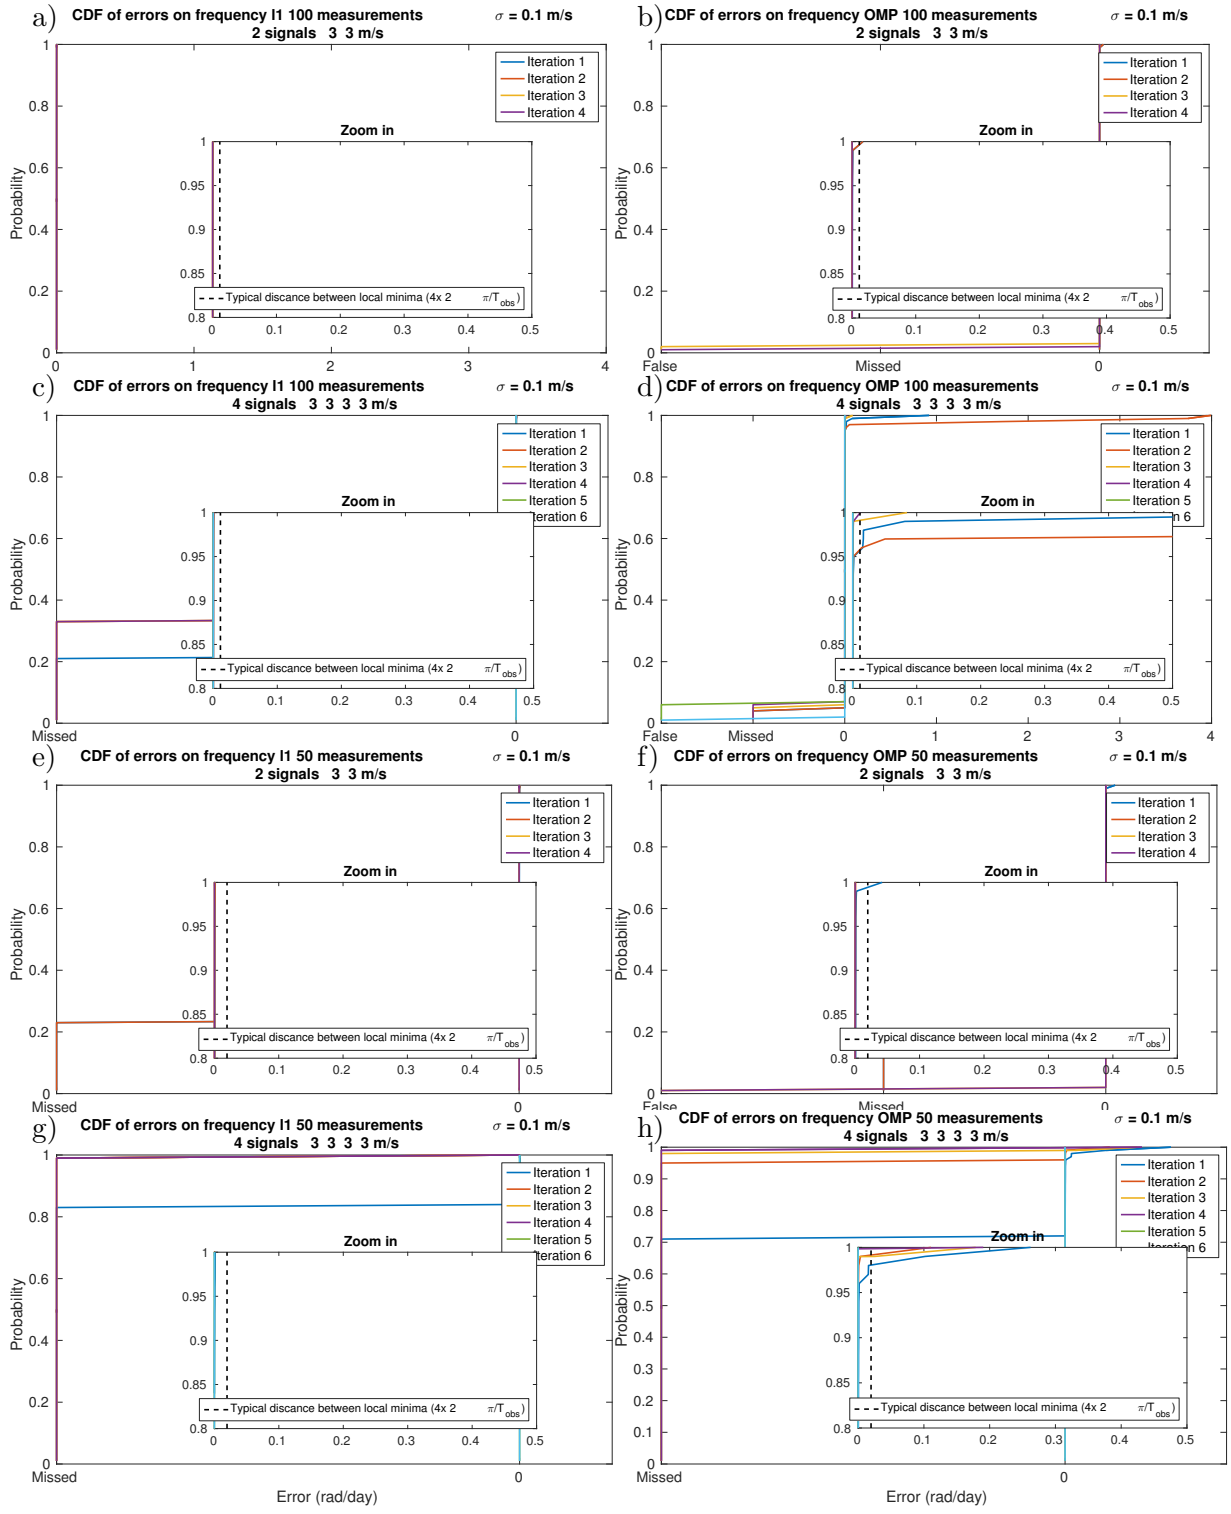

Figure 2.19: Cumulative distribution functions of errors, FAP threshold = 10^{-3}.

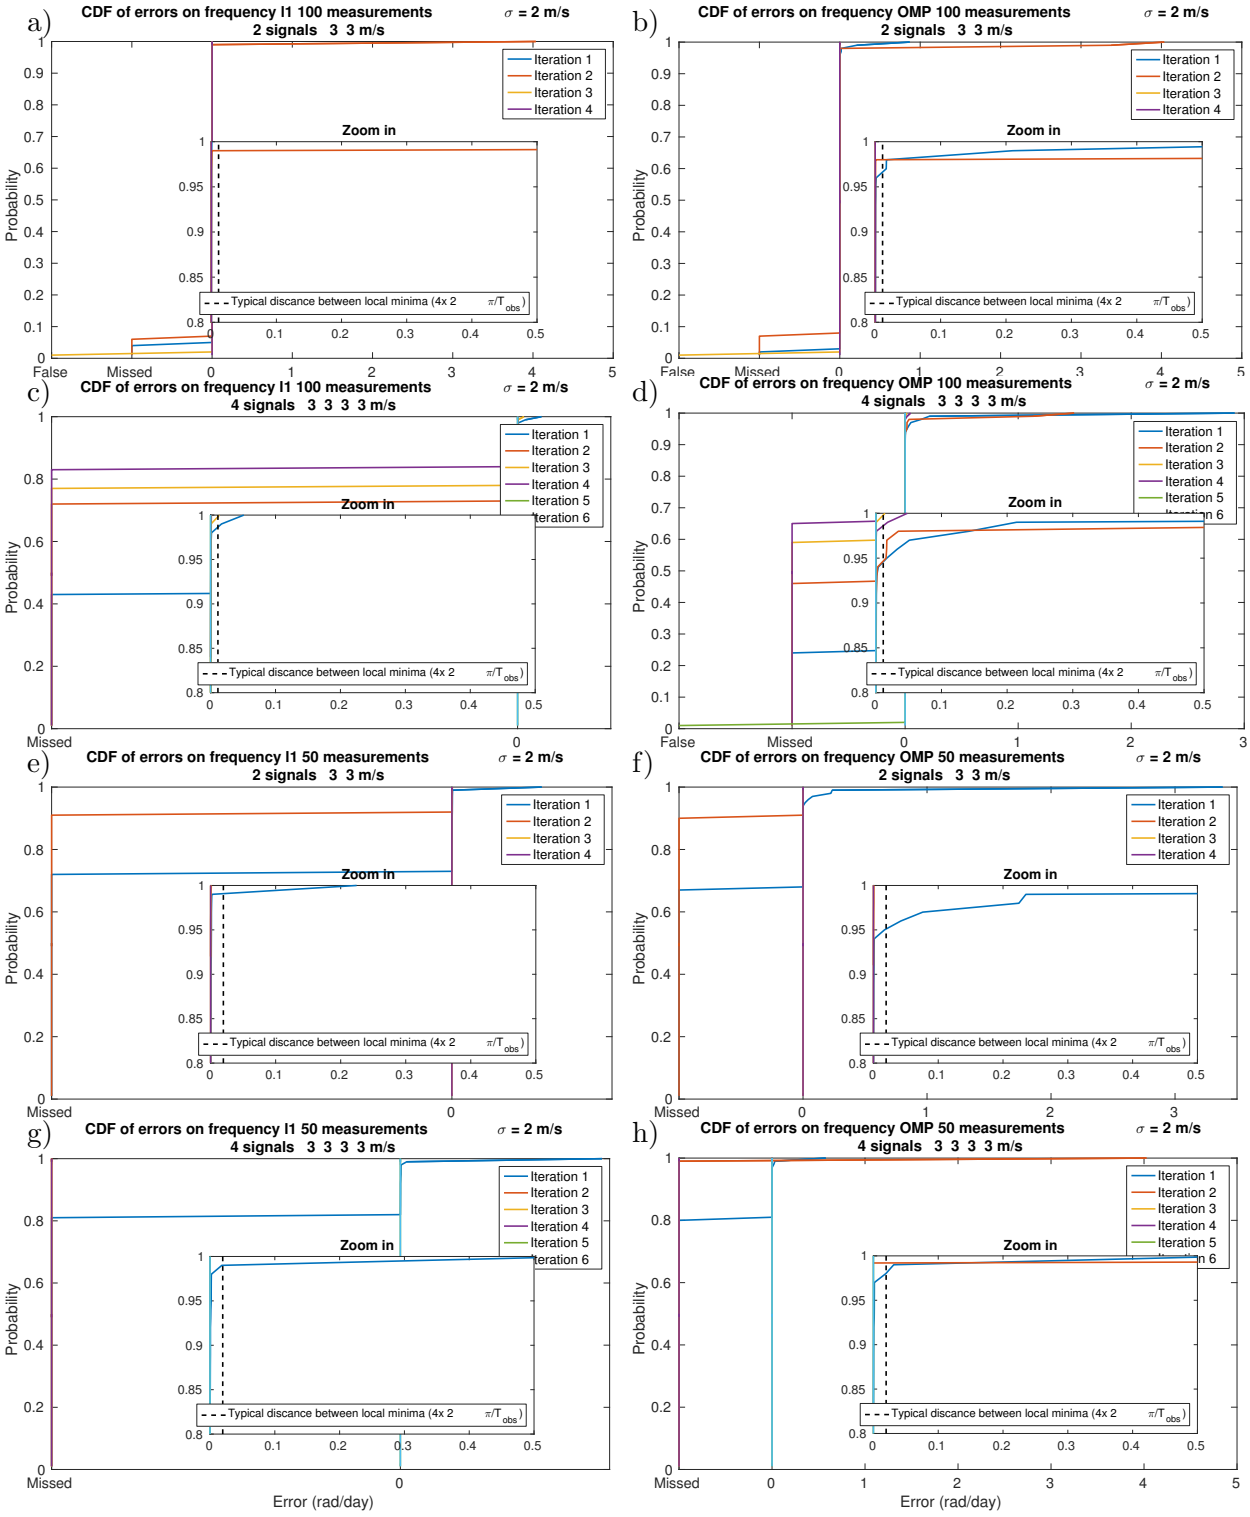

Figure 2.20: Cumulative distribution functions of errors, FAP threshold = 10^{-3}.

projected the FWHM, bisector span and $\log R'_{HK}$ onto those subspaces, obtaining $3c$ vectors $(\boldsymbol{a}_j)_{j=1..3c}$. In our study we chose $c = 5$, so the dictionary was made of $n = 2n_\omega + 3c$ entries. In what follows, $\mathbf{A}$ denotes the $m \times n$ matrix whose columns are $\mathrm{e}^{-\mathrm{i}\omega_\mathrm{k}\mathrm{t}}$ and $\mathrm{e}^{\mathrm{i}\omega_\mathrm{k}\mathrm{t}}$ for $k = 1..n_\omega$ and $(\boldsymbol{a}_j)_{j=1..3c}$.

Figure 2.21 shows on the first system of the RV fitting challenge what the FWHM projected on the C_j families of polynomials looks like. The FWHM time series is represented in black on the top figure, along with its projection on polynomials of degrees 0 to 14. The middle figure shows the FWHM projected on higher order polynomials. With the notation of the previous paragraph, $j_1 = 14$, $j_2 = 59$, $j_3 = 149$, $j_4 = 299$ and $j_5 = 492$, which is the total number of measurements. One can see that the higher orders, the finer structures are seen. We finally plot the Generalized Lomb-Scargle periodogram of each of the 5 filters. Their signal is indeed localized in frequency.

In the version used for the RV Challenge, additional steps were also performed: the minimization was applied to $\boldsymbol{y}(\boldsymbol{t}) = \mathrm{FWHM}(\boldsymbol{t})$ and $\boldsymbol{y}(\boldsymbol{t}) = \mathrm{BIS}(\boldsymbol{t})$, where BIS stands for bisector span. We denote by $\mathrm{FWHM}(\boldsymbol{t})$ the vector of observations of FWHM and adopt the same convention for BIS and $\log R'_{HK}$ to avoid an overuse of bold fonts) In each case, the entries of the dictionary corresponding to the signal considered were removed (respectively FWHM or bisector span), and the filtered time series of the radial velocity were put instead. In other words, denoting by s_1,..s_c a signal s filtered respectively by the c class of polynomials $C_1, ...C_c$, when computing the ℓ_1-periodogram of the radial velocity time series one takes a dictionary made of $\mathrm{FWHM}_1(\boldsymbol{t}), ..., \mathrm{FWHM}_c(\boldsymbol{t})$, $\mathrm{BIS}_1(\boldsymbol{t}), ..., \mathrm{BIS}_c(\boldsymbol{t})$, $\log R'_{HK1}(\boldsymbol{t}), ..., \log R'_{HKc}(\boldsymbol{t})$ and $\mathrm{e}^{-\mathrm{i}\omega_k \boldsymbol{t}}$ and $\mathrm{e}^{\mathrm{i}\omega_k \boldsymbol{t}}$ for $k = 1..n_\omega$. When computing the ℓ_1-periodogram of the FWHM time series, the dictionary was made of $\mathrm{RV}_1(\boldsymbol{t}), ..., \mathrm{RV}_c(\boldsymbol{t})$, $\mathrm{BIS}_1(\boldsymbol{t}), ..., \mathrm{BIS}_c(\boldsymbol{t})$, $\log R'_{HK1}(\boldsymbol{t}), ..., \log R'_{HKc}(\boldsymbol{t})$ and $\mathrm{e}^{-\mathrm{i}\omega_k \boldsymbol{t}}$ and $\mathrm{e}^{\mathrm{i}\omega_k \boldsymbol{t}}$ for $k = 1..n_\omega$ where RV stands for radial velocity.

This was thought to prevent us from finding peaks due to a difference of frequency of the same activity signal measured by the radial velocity, FWHM, or bisector span. The result of such an approach are represented in figure 2.22 for the system 2 of the radial velocity fitting challenge. The planets to be found are represented by the red stems. The $\log R'_{HK}$ periodogram is scaled to fit on the figure with an arbitrary coefficient. The only peak that appears clearly on the radial velocity ℓ_1-periodogram and not on the other ℓ_1-periodograms is 3.77 days, which was then the only one reported.

2.4.3 Second approach

After the challenge, we kept on working on the method. We made a small change which radically modified the aspect of the figures, whose efficiency could be evaluated as the result became public. For instance, we could retrieve figure 2.23.b instead of figure 2.22. The first change consisted in fitting the ancillary indicators: FWHM, Bisector span and $\log R'_{HK}$ along with a constant and a trend before running the SPGL1 solver. The second change was the normalization and averaging operations presented section 2.2.3 and 2.2.3.

We checked that fitting a linear model with the ancillary indicators did not changed drastically the spectral content of the signal (appendix D of Hara et al. (2017), reproduced page 211) and discussed the pros and cons of fitting a model itself noisy (the indicators).

The method so defined combined with a search with a residual periodogram at low amplitude is extremely efficient. It retrieves planets with a semi-amplitude of ≈ 34 cm/s buried in stellar activity of several meters per second. For instance for system 2, the signal at 20.16

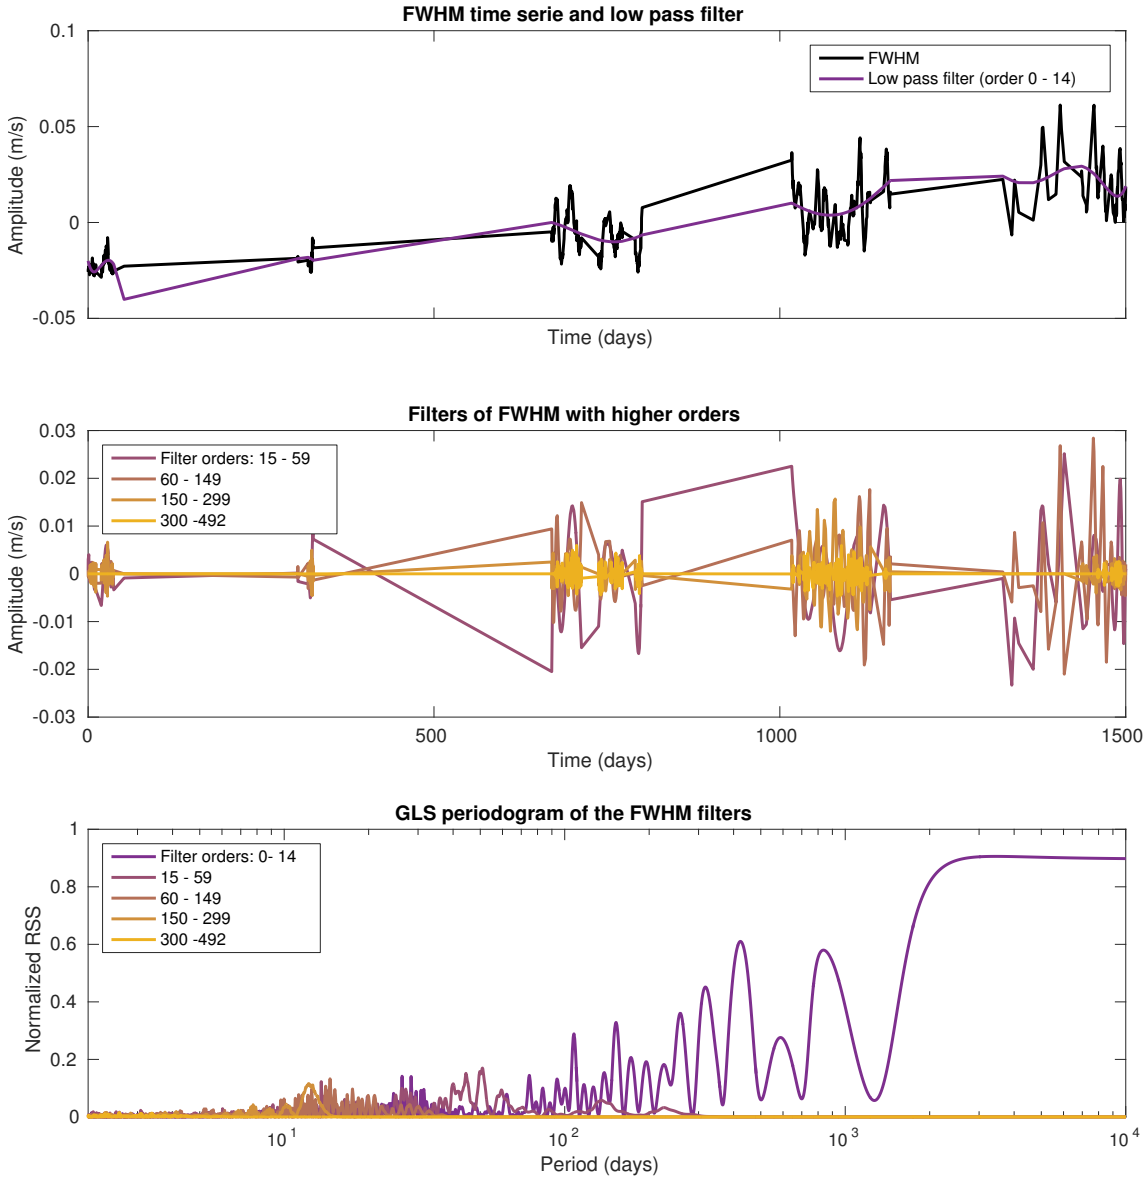

Figure 2.21: Filtered signals that were included in the dictionary in the old version of the ℓ_1-periodogram, used for the Radial Velocity Fitting Challenge.

days appears on the periodogram and is statistically significant. Its K/N as defined in Dumusque et al. (2016) is 2.26, that is for a semi-amplitude of the planet K, a radial velocity root mean square $\mathrm{RV}_{\mathrm{rms}}$ and m measurements,

$$K/N = \frac{K}{\mathrm{RV}_{\mathrm{rms}}\sqrt{m}}. \tag{2.57}$$

No team found signals with SNR lower than 5 and most bona-fide detections were made at $K/N \geqslant 7.5$. The ℓ_1-periodograms of the five first radial velocity signals are shown figure 2.23.a, b, c and 2.24.a, b. Even more remarkably, no red noise model was used to obtain

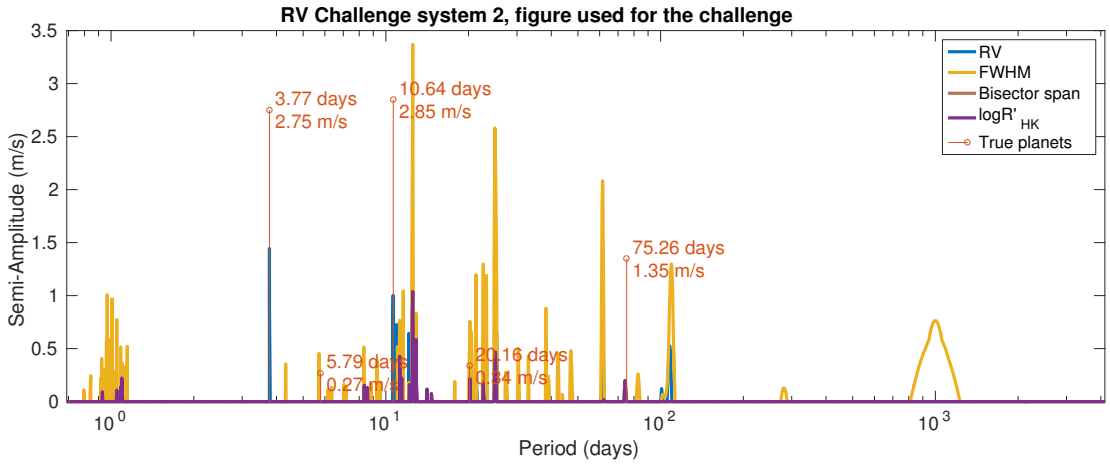

Figure 2.22: Figures computed for the radial velocity challenge.

the figures.

It is tempting to then claim the superiority of the ℓ_1 minimization over other methods but we believe that more care is needed for two reasons. First of all, the efficient method was found after the results were made public. The other teams possibly modified their method to retrieve small amplitudes as well. More importantly, it suggests that the modelling of the activity used for the challenge is too well-behaved. The fact that a linear combination of the ancillary indicators gives an extremely good account of the activity seems unrealistic. This suggestion is also supported by the fact that Mikko Tuomi's team, who performed best, is the only one that fitted an activity model of the form $a\text{FWHM} + b\text{BIS} + c \log R'_{HK}$. Also Phil Gregory used a term proportional to $\log R'_{HK}$ in its model and obtained good results (he analyzed only the first five systems). As a conclusion, either the problem of activity can be efficiently mitigated with enough measurements, or the simulation of the activity by SOAP 2.0 is not completely realistic.

2.5 HD 169830

The method has been applied to four radial velocity signals in Hara et al. (2017). We here treat one more case which illustrates an aspect treated more in-depth next chapter: the sensitivity of the data analysis to error estimation. The system in question is HD169830, around which two Jupiters are known to orbit (with best fit periods 225.62 and 2102 days). The data set contains measurements from CORALIE and HARPS, which have very different precisions.

The periodogram of the raw data is displayed i figure 2.25.a. It exhibits an odd structure with short period oscillations modulated by larger ones. In fact, the larger peaks correspond to $\approx 4 \times 2\pi/T_{\text{obs,HARPS}}$, which means these are due to the spectral window of HARPS. Since the nominal errors on the measurements of HARPS data are much smaller than CORALIE's but there might be an offset between the two data sets, spurious periods are created.

When an additional error is added in quatrature to the nominal one, that are $\sigma'_{RV} = \sqrt{\sigma_{RV}^2 + \sigma^2}$, the fact that we have overly precise measurement but poor accuracy becomes less

a)

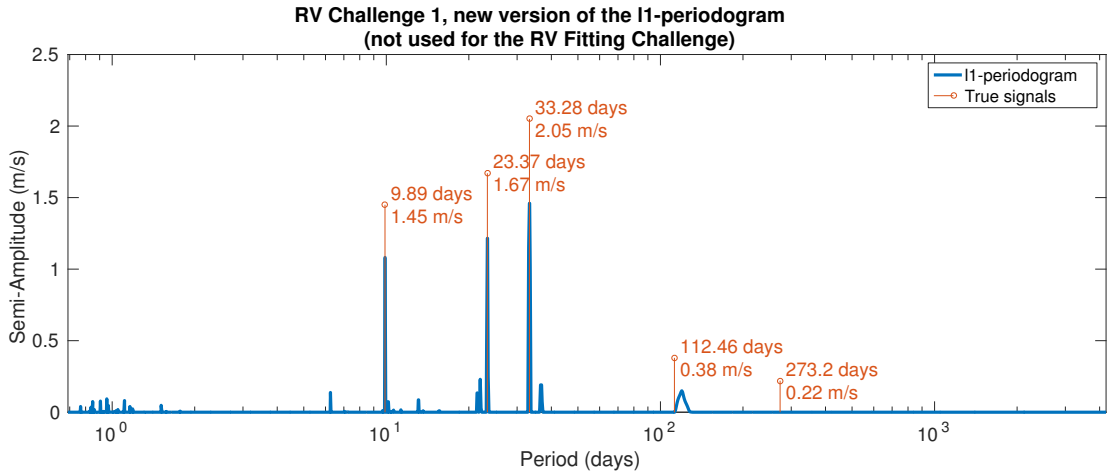

b)

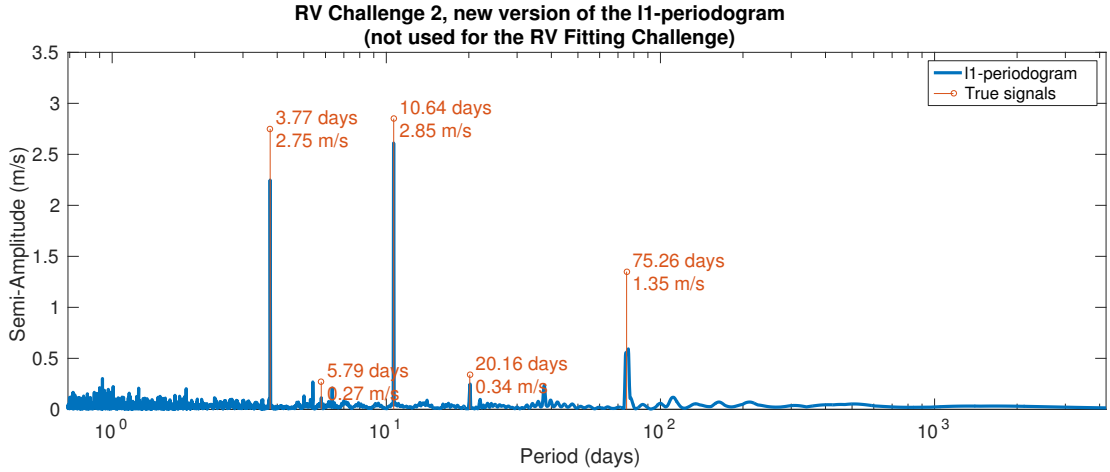

c)

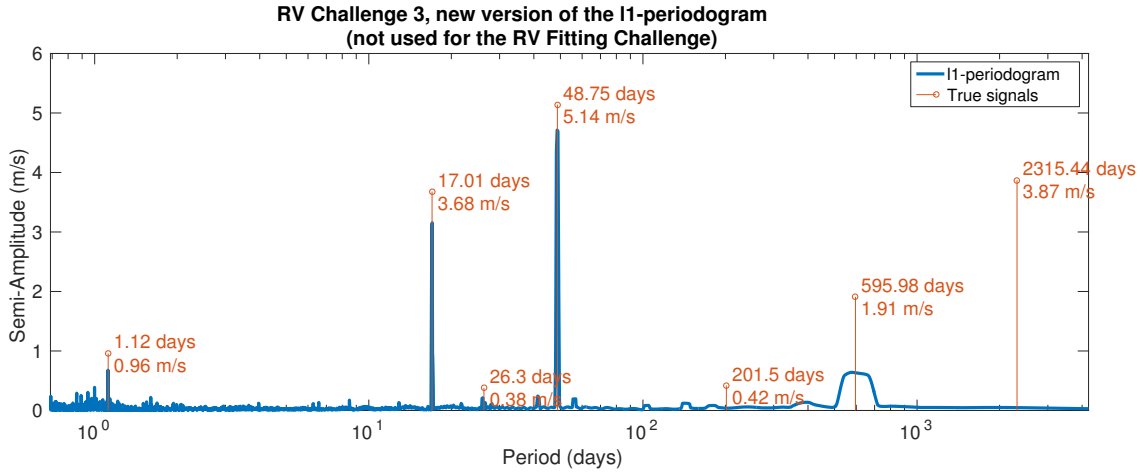

Figure 2.23: RV fitting challenge system 1, 2, 3 (respectively a), b), c))

a)

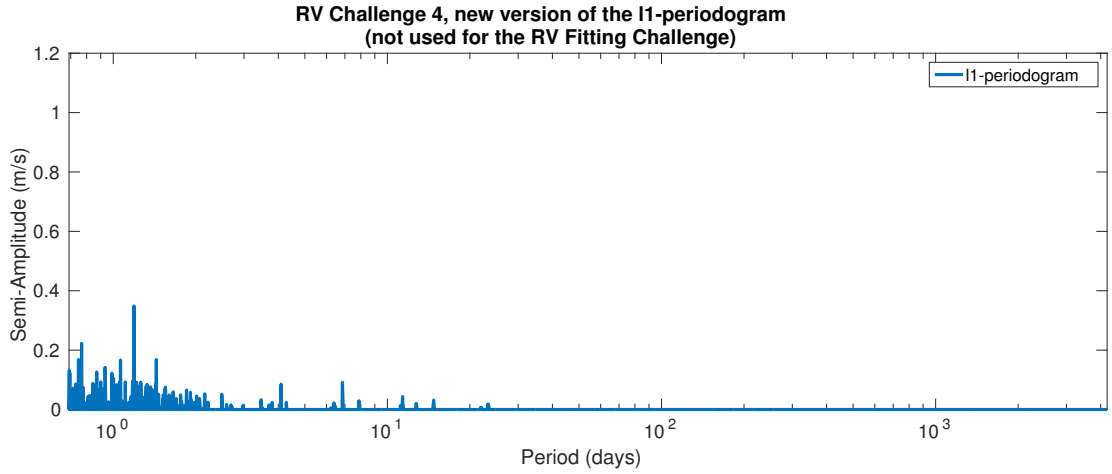

b)

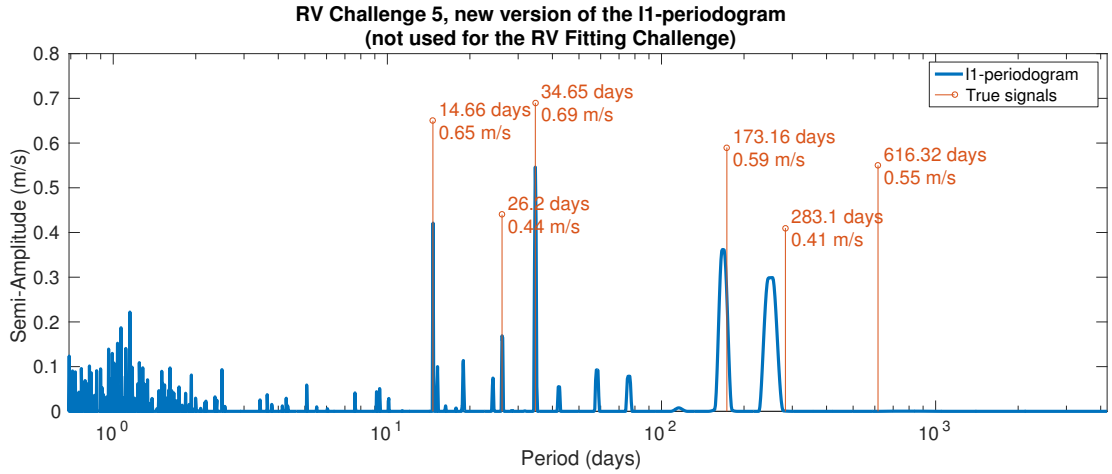

Figure 2.24: RV fitting challenge system 4, 5 (respectively a), b))

of a problem. We represent the periodogram for $\sigma = 1$ and 2 m.s^{-1} respectively figure 2.25.b and c. HARPS data does not constrain the data as much as before and the long period oscillations of the periodogram disappear, a strong peak at 225 days appears.

When computing the ℓ_1 periodogram with nominal error bars the two known periods clearly appear (see figure 2.25.a). Fitting them along with the offsets of the data sets allows to have an estimate of the offsets. When these are corrected and a 1 m.s^{-1} error is added in quadrature to the nominal one, one obtains figure 2.25.b. A signal appears very clearly at $122.5 = 225/2$ days, which indicates an eccentric orbit of the planet orbiting at 225 days, and clarifies what signals are in the system.

a)

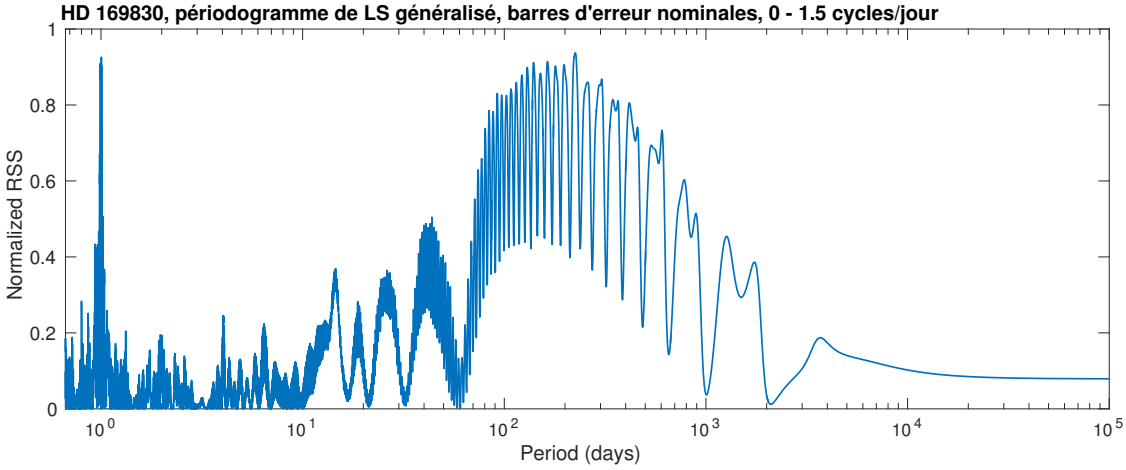

b)

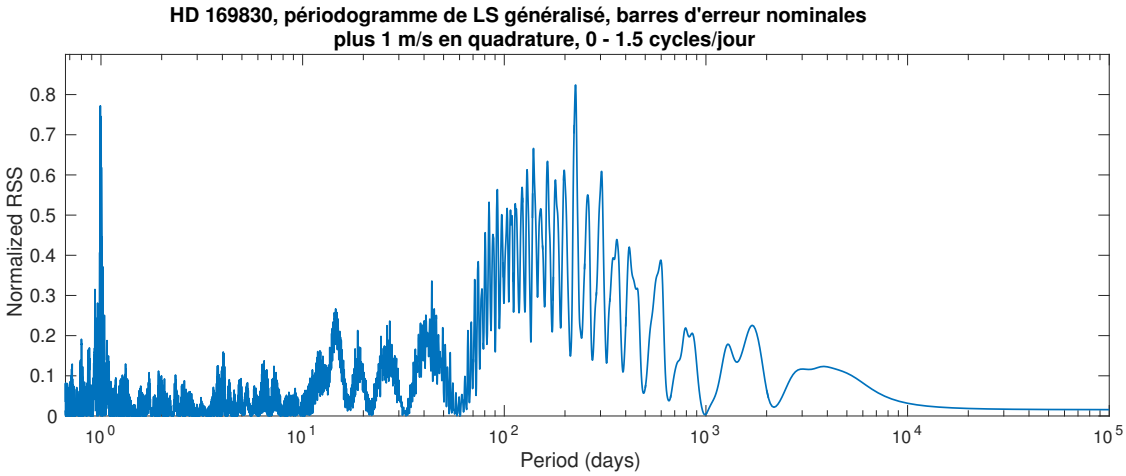

c)

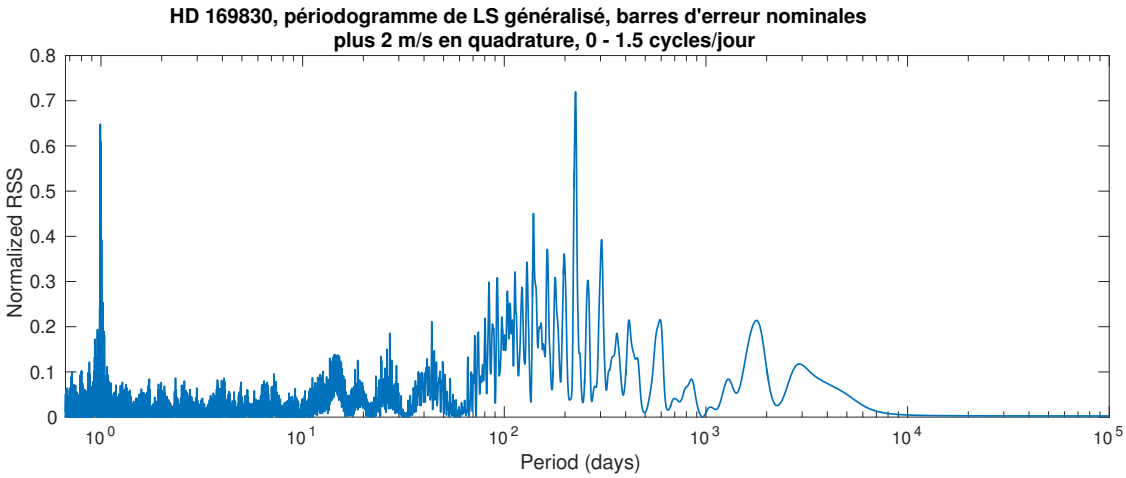

Figure 2.25: RV fitting challenge system 1, 2, 3 (respectively a), b), c))

a)

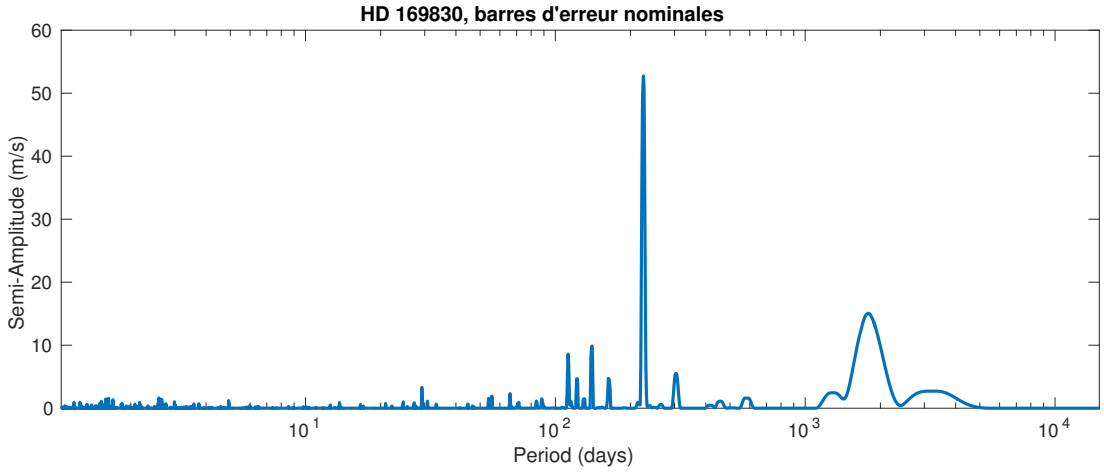

b)

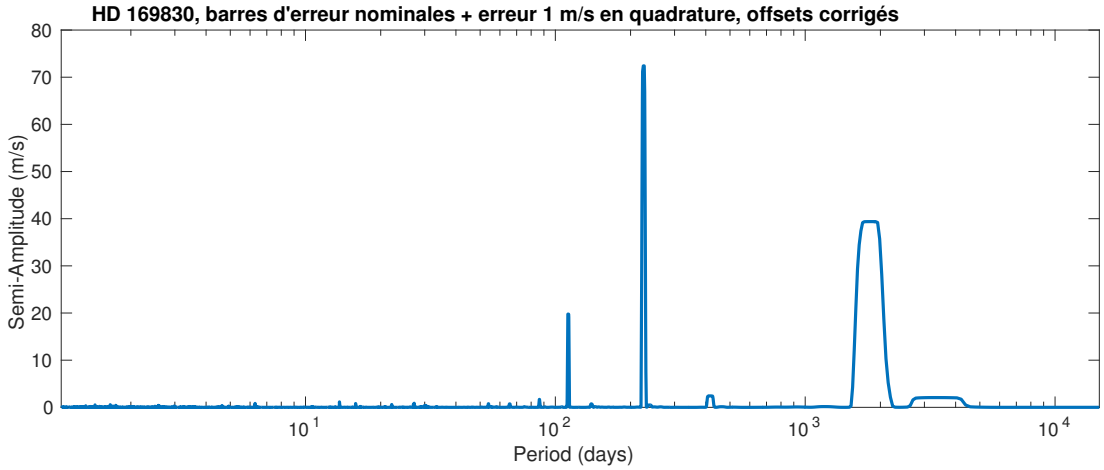

Figure 2.26: RV fitting challenge system 1, 2, 3 (respectively a), b), c))

2.6 Discussion

The goal of this chapter was to present theoretical and practical aspects that could not be developed in Hara et al. (2017). On the theoretical sides, we highlighted the difficulty of obtaining results that apply to continuous dictionaries. In the meantime, we consider the columns of the dictionary $\mathbf{A}$ not like single entities, but group them when those represent the same periodogram peak by summing up contributions of neighbouring frequencies. Motivated by a theoretical result we introduced the use of the restricted isometry constant as a generalization of the spectral window to study correlation between signals, but more work is yet to be done to link precisely this quantity to actual chances of success of sparse methods on realistic models of radial velocity data.

On the practical side, we highlighted the difficulty of securing small signals in noisy data. In particular the structure of the ℓ_1 periodogram might change depending on whether the grid

includes or not frequencies of 1 cycle per day. Due to the daily repetition of measurements, the correlation between a signal at frequency ω_0 and its alias at $\omega_0 + 2\pi$ is above 0.9 in general and as expected from theoretical considerations, it lowers the chance of success of sparse methods. To circumvent this issue, we argue that recursive/residual periodograms are efficient ways to dig in the noise. The ℓ_1 minimization is useful to prevent from being mislead by combinations of higher SNR signals that are strongly correlated. Furthermore, ℓ_1 methods offer a point of comparison for residual periodograms at low amplitude signals. One can check that the signals that appear on residual periodograms are not aliases of the low-amplitude peaks that appear on the ℓ_1 periodogram.

In the era of parallel computing, the need for efficient algorithm decreases. A brutal force search of two or three frequencies on a grid is manageable. Also, the algorithm FREDEC (Baluev 2013a) has an efficient approach. A certain number N of candidate frequencies are selected from the periodogram and all their $2^N - 1$ combinations are fitted non-linearly. The algorithm allows the fitting of ≈ 25 frequencies and an analytical approximation of the false alarm probability of such a procedure is available (Baluev 2013). These approaches could prevent most situations where the aliasing misleads the analyst by enhancing spuriously the power at certain frequencies.

Furthermore, the detection of planets is secured by computing the Bayesian posterior of models with k planets. The posterior distribution of the orbital elements answers the "right question", which is the probability of the model knowing the data. As we discussed section 1.3, as soon as we look at the data the properties of inference are changed. So why should the analysis be biased by looking at other quantities, here the ℓ_1 periodogram or residual periodograms?

We believe there is one fundamental reason why these might still be interesting tools, which is that the model is not always correct. Brutal force search or Bayesian evidence does not allow to "see" the data, and a trained observer knows how to recognize signatures of outliers, wrong noise models, etc. that appears on certain figures. Besides the present work, residual periodograms (Baluev 2008, 2009, 2013b; Anglada-Escudé & Tuomi 2012; Baluev 2013, 2015b) Other tools such as the stacked periodograms (Mortier & Collier Cameron 2017), that compute a series of periodograms by adding one measurement at a time or AGATHA (Feng et al. 2017) that can compute RV from different part of the spectra and on a moving timespan, allow to quickly look at the data from different angles. If finally finding a signal with one of those method is not reliable to confirm a signal, since the more hypotheses are explored, the more likely it is to find something, if a signal appears whatever the way the data is being looked at, this reinforce confidence in it. Secondly, the codes to compute the evidence are not (yet) perfectly stable, and cross-validating a result can improve confidence in it. Finally, the search for candidate interesting signals, further confirmed by additional information, is an important part of data analysis.

2.6.1 Other RV analyses with ℓ_1 minimization

Before moving to another topic, let us briefly present Hara et al. (2017) (see appendix G) and highlight the parts that were not thoroughly discussed above. We refer to section "G.i" as the section i of the article for $i = 1..6$. Sections "G.7" to "G.10" refer to the appendices of the article (A to D).

In section G.1, we present some sparse recovery methods for spectral line estimation, and

show that similar ideas emerged contemporaneously in the statistical and astronomical communities. Section G.2 briefly presents the Basis Pursuit and Atomic norm De-Noising algorithms. In section G.3, we describe the algorithm tuning. Both subjects were treated into more depth respectively in sections 2.1.3 and 2.2. In section G.4, we analyse in detail four systems: HD69830 (Lovis et al. 2006), HD10180 (Lovis et al. 2011b), 55 Cnc (Butler et al. 1997; Marcy et al. 2002; McArthur et al. 2004; Fischer et al. 2008; Endl et al. 2012) and GJ 876 (Marcy et al. 1998; Delfosse et al. 1998; Rivera et al. 2005; Rivera et al. 2010; Correia et al. 2010). Note that the cited papers are the ones reporting measurements, previous data analyses are considered in the article. Results are discussed in section G.5, in particular we emphasize the role of aliases as threats to the quality of the recovery. We suggest to use the spectral window to spot aliases, as classically done with the Lomb-Scargle periodogram. Finally, conclusions are drawn in section G.6. Adding Keplerian atoms to the dictionary in section 2.3 is a first step in one of the direction for future work outlined in that section.

In the appendices of the article, the following topics are addressed. In section G.7, we compute the grid spacing as in 2.2.3, paragraph "Grid spacing and tolerance, first approach". Appendix B (section G.8) discusses the occurrence of spurious tallest peak of the periodogram with the HD 69830 measurement times.

Let us highlight that appendices C and D (sections G.9 and G.10) treat subjects that are not considered in detail in the present chapter. In appendix C (section G.9), we show that using the appropriate weight matrix $\mathbf{W}$ can help bringing out small signals in correlated noise. Perhaps more importantly, in appendix D (section G.10), we discuss the soundness of fitting the ancillary measurements (FWHM, BIS, $\log R'_{HK}$) to the RV times series. The content of these two appendices, along with the application to real data (section G.4) are complementary to the present chapter.

Chapter 3

Bias and robustness of eccentricity estimates

Not having all the information you need is never a satisfactory excuse for not starting the analysis.

9^{th} Akin's law of Spacecraft Design

3.1 Notation

In this section we will need to use mean motions, usually denoted by n. In what follows, the number of measurements will be denoted by N and the number of parameters of a model is denoted by p.

3.2 Introduction

3.2.1 Content

It has been known for long that parameter estimates are in general biased when the model is non linear. In the context of exoplanet searches, this applies in particular to inclination, projected mass and eccentricity obtained respectively from astrometry, radial velocity and both. Eccentricity is of particular interest because it is an informative indicator of the past of the planetary system, and is also of greater complexity since it suffers detection bias on top of an estimation one. Our aim is to improve the robustness of the inference of eccentricities, and orbital parameters in general. We provide a geometric interpretation of the eccentricity detection and estimation biases, which also explains biases on inclination and projected mass. We then quantify analytically the amplitude of the biases. We pay particular attention to mismodelling and show that not only the uncertainty increases as the assumed model moves away from the real one, but also so does the bias, while uncertainties might be underestimated. We then derive consequences on the analysis of individual systems as well as populations of planets. The solution we advocate for individual systems is close to common practice, which is to privilege interval estimates rather than point estimates, scaling the error so that the χ^2 of the residuals equals one, to use the most complex model that seems reasonable, to check

that the residuals are well-behaved and to envision worst-case scenarios. This solution is discussed in the Bayesian and frequentist settings on simulated data sets as well as on the CoRoT-9 data. We also suggest a few methods to correct the eccentricity distributions, and to design observation strategies to ensure that the inferred distribution will have a certain resolution.

3.2.2 Introduction

The nearly coplanar and circular orbits of the Solar system have long been an argument in favour of Laplace and Kant's theory of formation of planets in a disk (Swedenborg 1734; Kant 1755; Laplace 1796). The first observations of exoplanets soon suggested that such low eccentricities were rather the exception than the rule. The "eccentricity problem", along with the need to envision migration scenario for hot Jupiters, triggered several theoretical studies which explored migration scenarios after the disk has dissipated. For instance Jurić & Tremaine (2008) and Petrovich et al. (2014) (and reference therein) evaluate the likelihood of formation scenarios of hot and warm Jupiters through their agreement with observed distributions. Since high eccentricities allow to constrain the history of planetary systems, it is key to estimate it as precisely as possible. However, it is known since Lucy & Sweeney (1971) that eccentricities are over estimated.

This issue is known among the exoplanet community, and several works addressed it. The amplitude of the bias has been examined through numerical simulations in Shen & Turner (2008) and Zakamska et al. (2011). It was found to depend on the signal to noise ratio as well as on the time span of the observations, and for the latter the phase coverage. In addition, Zakamska et al. (2011) updated the Lucy & Sweeney (1971) null hypothesis test to determine if the hypothesis of a null eccentricity can be rejected or not. Pont et al. (2011), Husnoo et al. (2011) and Husnoo et al. (2012) used Bayesian Information Criterion to confirm non-zero eccentricities. More recently, Bonomo et al. (2017b,a) used a Bayes factor to determine if a non null eccentricity is significant or not with non informative prior. A Bayesian test with a physically motivated prior on eccentricity was also devised by Lucy (2013). For the correction of the eccentricity distributions, Hogg et al. (2010) searches for the distribution that maximises a certain posterior likelihood, possibly parametrized by a beta function as advocated by Kipping (2014). Furthermore, it has been noted by Cumming (2004) that eccentricities $\gtrsim 0.6$ are more difficult to detect.

In the works mentioned above, the amplitude of the bias and the robustness of the inference is tackled. There are also some insights on the origin of the bias, for instance that since eccentricity cannot be negative, the variation of the estimate due to noise can only result in a strictly non zero eccentricity with probability one. Also, it is noted that the estimates of $k = e\cos\omega$ and $h = e\sin\omega$ have a Gaussian distribution, therefore $e = \sqrt{k^2 + h^2}$ has a Rice distribution and is biased upwards (Shen & Turner 2008). The initial motivation of the present work was to build upon these results and understand precisely the properties of the eccentricity estimates. More precisely, supposing the noise is Gaussian (possibly correlated), we make the following contributions

- We devise several ways to visualise the geometry of Keplerian models in the space of observations, which is basically a five dimensional manifold embedded in a N dimensional space, N being the number of observations.

- With a model linearised in k and h one can compute a good analytical approximation of the bias that generalizes the analysis of Lucy & Sweeney (1971) and sheds some light on the observations of Shen & Turner (2008) and Zakamska et al. (2011).
- When the signal is modelled incorrectly (there is an unseen noise, unseen outlier, a non-detected companion, the model is too complex...) the average bias is higher. The uncertainties are in general underestimated if the model is too simplistic or obviously when the noise is assumed to be too low.

From these observations we make recommendations and design tools to improve the robustness of the analysis of orbital parameters.

- We suggest a simple weighting of the posterior distributions to assess dependence of the result on the prior.
- We compute an analytical approximation of the distribution of the maximum likelihood ratio to obtain a precise confidence interval on eccentricities.
- We show that normalizing to one the reduced χ^2 of the residuals, which is common practice, indeed allows in most cases to draw correct conclusions even if the additional noise is non Gaussian. We also show in which (unlikely) situations this normalization could lead to spurious results.
- The maximum likelihood highlights relative merits: it selects the best model in some sense, but possibly among models that are all very bad. We suggest a method to analyse the residuals of a fit to check whether a selected model is acceptable in an absolute sense.
- We provide a simple adaptation of Hogg et al. (2010) to take detection and estimation bias in the corrections of distributions;
- We propose a method to design observation strategies as a function of the desired accuracy of the population retrieval.

The chapter is organized as follows. In section 3.3 we first define precisely the problems we are considering and our notations. Then we develop a geometrical view and some tools to understand the origin of the bias. Finally we draw consequences on the reliability of the estimates of the eccentricity of a planet and the distributions of eccentricities. In section 3.4, we present the results listed above. These ones are applied to real cases in section 3.5. In particular we confirm that CoRot-9b has a non null eccentricity, as claimed in Bonomo et al. (2017b). We discuss the advantages and limits of our analyses in section 3.6. In order to facilitate the reading, we use as few equations as possible in the core of the text, precise statements as well as proofs and mathematical details are given in the appendices.

3.3 Origin of the eccentricity bias

3.3.1 Problem statement

In this section we state the problem under study mathematically. Let us consider a series of N observations, modelled as a vector $\boldsymbol{y} = (y(t_k))_{k=1..N}$, such that

$$\boldsymbol{y}(\boldsymbol{t}) = f(\boldsymbol{t}, \boldsymbol{\theta}) + \boldsymbol{\epsilon} \tag{3.1}$$

where $\boldsymbol{t} = (t_k)_{k=1..N}$ is the vector of measurement times, f is a deterministic model depending on parameters $\boldsymbol{\theta} \in \mathbb{R}^p$ and $\boldsymbol{\epsilon}$ is a random variable modelling the noise. We wish to determine the model parameters $\boldsymbol{\theta}$. An estimator of $\boldsymbol{\theta}$ is a function δ of $\boldsymbol{y}(\boldsymbol{t})$ whose output is wanted to be close to $\boldsymbol{\theta}$ in a sense chosen by the data analyst. Since $\boldsymbol{\epsilon}$ is a random variable, $\boldsymbol{y}$ and $\delta(\boldsymbol{y})$ are too. If the mathematical expectancy of $\delta(\boldsymbol{y})$, denoted by $\mathbb{E}\{\delta(\boldsymbol{y})\}$ is not equal to $\boldsymbol{\theta}$, the estimator is said to be biased and

$$b_\delta(\boldsymbol{\theta}) = \mathbb{E}\{\delta(\boldsymbol{y})\} - \boldsymbol{\theta} \tag{3.2}$$

is called the bias of the estimator δ in $\boldsymbol{\theta}$. A common estimator is the maximum likelihood, which reduces to the least-square estimator when the errors are Gaussian (Pelat 2013). If $(\sigma_k)_{k=1..N}$ are the measurement uncertainties,

$$\delta_{LS}(\boldsymbol{y}) = \arg\min_{\boldsymbol{\theta}\in\Theta} \sum_{k=1}^{N} \frac{(y(t_k) - f(t_k, \boldsymbol{\theta}))^2}{\sigma_k^2}. \tag{3.3}$$

When the model is linear the observations are modelled by $\boldsymbol{y} = \mathbf{A}\boldsymbol{\theta} + \boldsymbol{\epsilon}$ where $\mathbf{A}$ is a matrix. If $\mathbf{A}$ is known, as long as $\boldsymbol{\epsilon}$ has zero mean and covariance matrix $\mathbf{V} = \mathbf{W}^{-1}$, the ordinary least square estimate

$$\delta_{LS}(\boldsymbol{y}) = (\mathbf{A}^T\mathbf{W}\mathbf{A})^{-1}\mathbf{A}^T\mathbf{W}\boldsymbol{y} \tag{3.4}$$

is unbiased. However, when the model is of the form $\boldsymbol{y} = \boldsymbol{f}(\boldsymbol{\theta}) + \boldsymbol{\epsilon}$ where $\boldsymbol{f}$ is a non linear function, the least square estimate can be biased. This has been noted for instance by Hurwicz (1950) and discussed in Hartley (1964); Bates & Watts (1980); Cook & Witmer (1985); Firth (1993).

We will be concerned with the Keplerian model, which is non linear. For clarity, we recall the four constitutive equations of the model as they were given in section 1.3.3,

$$y(t, e, K, P, \omega, M_0) = K(\cos(\omega + \nu(t, e, P, \omega, M_0)) + e\cos\omega) \tag{3.5}$$

$$\cos\nu = \frac{\cos E - e}{1 - \cos E} \tag{3.6}$$

$$\sin\nu = \frac{\sqrt{1-e^2}\sin E}{1 - \cos E} \tag{3.7}$$

$$E - e\sin E = M_0 + \frac{2\pi}{P}t \tag{3.8}$$

Where t, e, K, P, ω, M_0 designate respectively the time, eccentricity, semi-amplitude, period, argument of periastron and mean anomaly at $t = 0$. The sine and cosine of the true anomaly can be computed from equation (3.6), (3.7) and the Kepler equation (3.8). In the above notation, $\boldsymbol{\theta} = (e, K, P, \omega, M_0)$.

Equation (3.3) defines an estimator, which in principle allows us to compute the distribution followed by the estimator value, given a certain prior on the underlying distribution of $\boldsymbol{\theta}$. However, the definition of the maximum likelihood estimator is implicit, making such calculations difficult in most cases. In this section we will show that the bias can be understood by qualitative properties of the model and more precisely by its geometry.

3.3.2 Notations

In this chapter we use the formalism of probability theory. To facilitate the reading, we define the conventions taken to denote the usual operators. For a vector of parameters $\boldsymbol{\theta}$, we denote by $p(\boldsymbol{\theta})$ its probability density, that is the function such that $\Pr\{\boldsymbol{\theta} \in \Theta\} = \int_\Theta p(\boldsymbol{\theta}) d\boldsymbol{\theta}$ for some measurable set Θ. The assertion "The random variable $\boldsymbol{\theta}$ has a distribution $p(\boldsymbol{\theta})$" is denoted by $\boldsymbol{\theta} \sim p(\boldsymbol{\theta})$. The probability of $\boldsymbol{\theta}$ knowing an event ω is denoted by $p(\boldsymbol{\theta}|\omega)$ and its mathematical expectancy is denoted by $\mathbb{E}\{\boldsymbol{\theta}\} = \int_\Theta \boldsymbol{\theta} p(\boldsymbol{\theta}) d\boldsymbol{\theta}$. The true value of a parameter is denoted by the subscript t (here $\boldsymbol{\theta}_t$) and its maximum likelihood estimate is $\widehat{\boldsymbol{\theta}}$. We denote a Keplerian model sampled at times $\boldsymbol{t} = (t_k)_{k=1..N}$ and of parameters $\boldsymbol{\theta}$ (orbital parameters plus possibly offset, trend...) by $\boldsymbol{y}(\boldsymbol{t}, \boldsymbol{\theta})$.

3.3.3 A geometrical interpretation

Set up

We consider the model $\boldsymbol{y}(\boldsymbol{t}, \boldsymbol{\theta})$ as the set of points $\mathcal{M} = \{\boldsymbol{y}(\boldsymbol{t}, \boldsymbol{\theta}) \in \mathbb{R}^N, \boldsymbol{\theta} \in \Theta\}$ where $\Theta = \{(e, K, P, \omega, M_0)\} = [0, 1] \times \mathbb{R}_+ \times \mathbb{R}_+ \times [0, 2\pi[\times [0, 2\pi[$ is the set of possible values of the parameters. $\mathcal{M}$ is a manifold with boundaries in the space of observations. The least square problem then consists in finding the element of this manifold that is closest to the observations in the sense of the usual distance (or Euclidian norm). We denote by $\mathcal{M}_e$ the set of models that have an eccentricity e, all other parameters free (in other words we consider a foliation of $\mathcal{M}$). The eccentricity estimated is e if the model closest to the observation is in $\mathcal{M}_e$. We have found several properties that account for detection and estimation biases of eccentricities.

1. The relative positions of $\mathcal{M}_e$, $e \in [0, 1]$ is such that there are "more models" with high eccentricity.
2. The eccentricity does not parametrize uniformly the models.
3. The manifold $\mathcal{M}_e$ has a shape that is more and more complicated as e grows. More specifically, it explores more and more dimensions of the space of observations.
4. The above three properties have an effect on the bias regardless of the implementation of the data analysis method. Besides, an imprecise numerical scheme can lead to even more incorrect estimates.

The precise description of these properties is the object of the following paragraphs.

More models with high eccentricity

Let us first consider the case of low eccentricities, so that the Keplerian model (1.2) can be approximated by its development at first order in eccentricity. The approximated model is

$$y(\lambda, K, P, e, \omega) = K(\cos(\lambda) + e\cos(2\lambda - \omega)) \tag{3.9}$$

where $\lambda = \omega + M_0 + nt = \lambda_0 + 2\pi t/P$ is the mean longitude, λ_0 being its value at $t = 0$. Denoting by $n = 2\pi/P$ the mean motion, the above expression can be re-written

$$\boldsymbol{y}(\boldsymbol{t}, A, B, C, D, n) = A\cos n\boldsymbol{t} + B\sin n\boldsymbol{t} + C\cos 2n\boldsymbol{t} + D\sin 2n\boldsymbol{t} \tag{3.10}$$

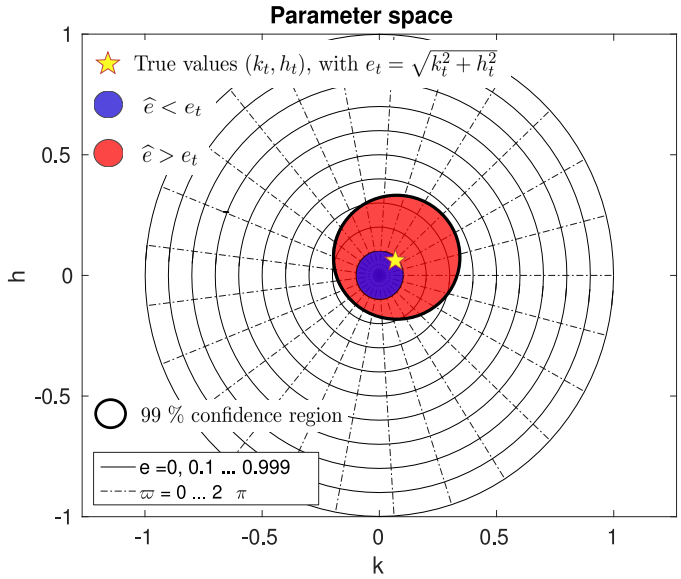

Figure 3.1: Representation of the estimation of $\sqrt{k^2 + h^2}$ when k and h are Gaussian.

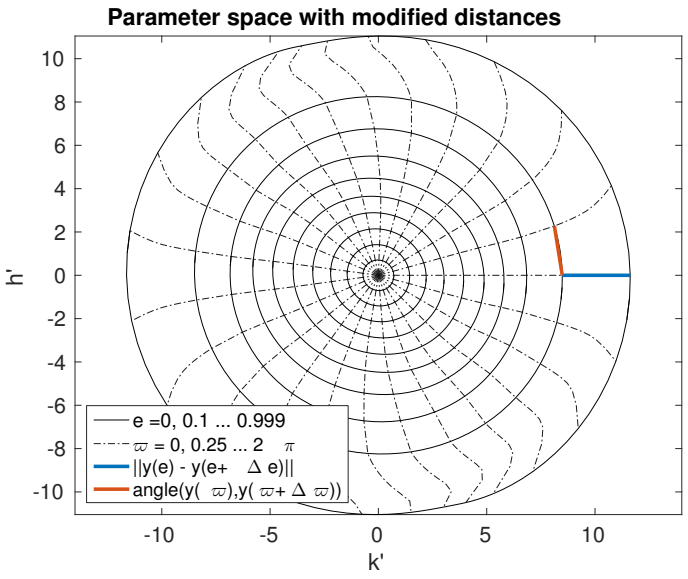

Figure 3.2: Representation of the phase space with subsequent points spaced proportionally to the distance between the corresponding models.

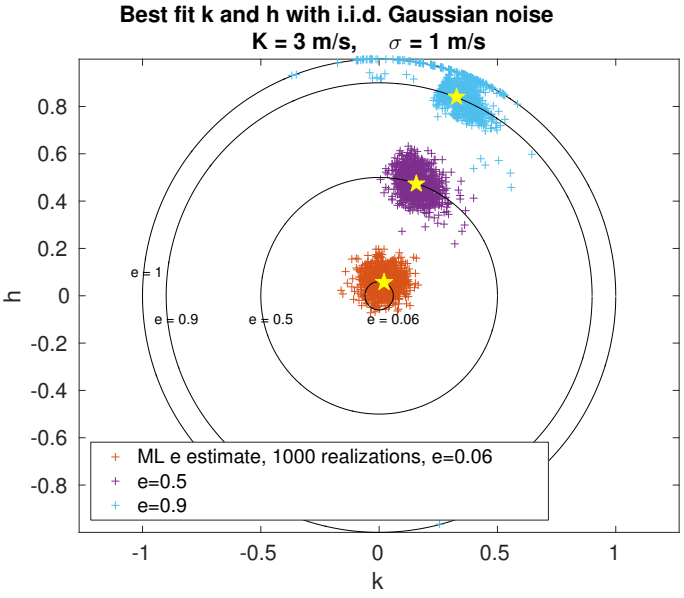

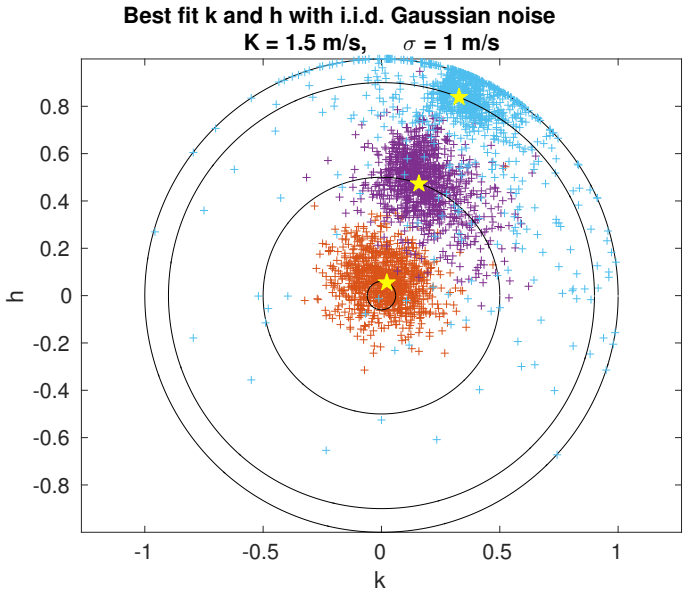

Figure 3.3: Estimates of k and h for true eccentricities e=0.06, 0.5 and 0.9 when the noise standard deviation is σ=1 m.s^{-1} and K=3m/s (left) or K=1.5 m.s^{-1} (right).

For the sake of simplicity we here suppose that n is fixed at the value given by the periodogram. The result is extended to the case where it is a free parameter in appendix C. With this assumption, in a matrix form, $\boldsymbol{y} = \mathbf{M}\boldsymbol{x}$. The estimate of $\boldsymbol{x}$, denoted by $\widehat{\boldsymbol{x}}$ verifies $\widehat{\boldsymbol{x}} = (\mathbf{M}^T\mathbf{W}\mathbf{M})^{-1}\mathbf{M}^T\mathbf{W}\boldsymbol{y}$ where $\mathbf{W}$ is the inverse of the noise covariance matrix. In this approximation, the semi-amplitude estimate is $\widehat{K} = \sqrt{\widehat{A^2} + \widehat{B^2}}$. We show in appendix C that provided e is small enough ($\leqslant 0.2$) and the number of observation is sufficient, $\widehat{k} = \widehat{C}/\widehat{K}$ and $\widehat{h} = -\widehat{D}/\widehat{K}$ both follow a Gaussian law of same variance.

In Fig. 3.1, we represent schematically the situation. The pair k_t, h_t is represented by a yellow star. Because of the noise, the values of $\widehat{k}, \widehat{h}$ are isotropically distributed around k_t, h_t. The bold black line represents the 99% confidence region. When the estimate falls in the blue-coloured region, the eccentricity is under-estimated. When it falls in the red-coloured region, the eccentricity is over-estimated. Since there is a bigger volume of higher eccentricity model $\mathcal{M}_e$ in the vicinity of k_t, h_t, the eccentricity is on average over-estimated, which is exactly saying it is biased upwards.

Since we made an approximation, one can wonder how accurate that explanation is. To check that, we generate Keplerian signals of eccentricity 0.06, 0.5 and 0.9 and in each case 1000 noise realizations. The array of time $\boldsymbol{t}$ is the 30 first measurements of GJ 876 (Correia et al. 2010). In both figures 3.3, the noise standard deviation is 1 m.s^{-1}. The amplitude of the true signal is respectively 3 and 1.5 m.s^{-1} (resp left and right). When $K = 3$ m.s^{-1}, the distributions of $\widehat{k}, \widehat{h}$ is fairly isotropic for $e_t = 0.06$ and $e_t = 0.5$. For $e = 0.9$ there seems to be more complicated phenomenon at work. The distribution has no circular symmetry and some estimates have $\widehat{e} = 1$. In the lower signal to noise ratio setting, the distribution spreads out. There is a slight deformation of the distribution at $e_t = 0.5$ towards $h \approx 0$ and $k \approx 0.5$. The $e_t = 0.9$ case is even more complicated with several dozens of estimate completely off.

Before moving to other effects that bring partial explanations to the properties of estimates of high eccentricity orbits, we can obtain an explicit expression of the bias at low eccentricity with formula (3.10). We argue in appendix C that $\widehat{k}$ and $\widehat{h}$ have a standard deviation $\sigma_k = \sigma_h = \frac{\sigma}{K_t}\sqrt{\frac{2}{N-p+1}}$, where p is the number of adjusted parameters. Then $\widehat{e} = \sqrt{\widehat{k^2} + \widehat{h^2}}$, which is a modulus of a random two-dimensional vector, follows a Rice distribution. Denoting by $S = 1/\sigma_k$ the signal to noise ratio and e_t the true eccentricity, in the linear approximation the bias reads

$$b(e_t, S) = \frac{1}{S}\sqrt{\frac{\pi}{2}}L_{1/2}\left(\frac{S^2 e_t^2}{2}\right) - e_t. \tag{3.11}$$

where $L_{1/2}$ is the Laguerre polynomial of order 1/2. The fact that $\widehat{k}$ and $\widehat{h}$ follow approximately a Gaussian distribution was noted in Ford (2005) and Shen & Turner (2008). In the latter it is explicitely mentioned that the eccentricity then follows a Rice distribution. Our discussion is fully compatible, and provides analytical expressions for the parameters of the distribution. Equation (3.11) reduces to a very simple expression when $e_t = 0$. In that case, the eccentricity follows a Rayleigh distribution as computed by Lucy & Sweeney (1971) and the bias becomes

$$b(0, S) = \sqrt{\frac{\pi}{2}}\frac{1}{S} = \frac{\sigma}{K_t}\sqrt{\frac{\pi}{N-p+1}}. \tag{3.12}$$

This expression is identical to equation (18) of Lucy & Sweeney (1971) except that we approximated the effect of the correlation of matrix $\mathbf{M}$ on the signal to noise through the term $-p+1$. [1]

Formula (3.12) is useful to see a few trends: the bias is proportional to the uncertainty on k and h, which is proportional to the inverse of K_t and $\sqrt{N-p+1}$. As a consequence, the bias increases as the signal to noise ratio decreases, i.e. as σ increases or as K_t or N decrease. This is fully compatible with the simulations of Shen & Turner (2008). We add that increasing the number of fitted parameters, p, increases the bias (at least at low e_t). We show examples in C where it is clear that adding the correction due to the number of parameters does give a better approximation of the bias.

The situation is in fact more complex, and formula (3.11) should be used as a lower bound. There are particular cases where the correlations between parameters increase drastically the uncertainties on k and h and therefore the bias of low e_t. Depending on the period, the bias can be as high as three times the value given by equation (3.11). Let us illustrate this with an example. We take the 28 measurement times of CoRoT-9 (Bonomo et al. 2017b) and generate a signal of a circular orbit with $K/\sigma = 3.5$. We then consider an array of equispaced mean motions. At a given value of the mean motion, we generate realisations of noise, and for each of them we compute an estimate of the eccentricity with the linearised model and the usual Keplerian model. We take 4000 realisations for the linear model and 100 for the non-linear one, due to computation time constraints. The 100 respectively 4000 values are averaged so to obtain an estimation of the bias per period. We then compute the average bias over all periods for both the linearised and full model. Fig. 3.4 shows the results of these computations, the bias as a function of period is in blue and gray for the linearised and Keplerian model respectively. The bias clearly depends on the period, reaching up to four times its average value. Formula (3.12), whose value is represented in purple, approximates well the average of the linear model (in red), but the bias of the Keplerian model (in black) is slightly greater. Below, in light blue, we represent the condition number (see definition 3 page 69) of the matrix $\mathbf{M}(P)^T\mathbf{M}(P)$ as a function of the period P. We see a correlation with the peaks of bias. This is easily interpreted: as the condition number increases, the inverse is more and more ill-defined. The diagonal terms of the covariance matrix, which is $\mathbf{M}(P)^T\mathbf{M}(P)$, increase and so does the uncertainty on $\boldsymbol{x}$ and in turn e. When considering correlated noise, the covariance matrix of the parameters is $(\mathbf{M}(P)^T\mathbf{W}\mathbf{M}(P))^{-1}$. It might turn out that the conditioning of $\mathbf{M}(P)^T\mathbf{W}\mathbf{M}(P)$ is worse than $\mathbf{M}(P)^T\mathbf{M}(P)$, yielding greater uncertainties and increased bias.

[1] Interestingly enough, eccentricities following a Rice distribution appear in another context, which is the chaotic diffusion in the Solar System. Laskar (2008) integrates over 5 Gyr the secular equations of motion of the Solar systems, as given in Laskar (1990). The integration is made for 1001 initial conditions, which are obtained by small perturbations of $k = e\cos\varpi$ and $h = e\sin\varpi$ in the initial conditions used in Laskar (1986), the other orbital parameters being unchanged. The integration step is 200 years, so one obtains 1001 series of the orbital elements every 200 years. For each planet, the 1001 series of eccentricities are grouped by 250 Myr time intervals. On each of these intervals, the experimental probability density functions of eccentricities are computed. For the 4 inner planets and each 250 Myr interval, the densities so obtained are very well fitted by a Rice distribution. The variance of these distributions provides an estimate of their spreading. For the 4 inner planets, the variances of the eccentricity distributions are found to increase linearly with time to a good approximation, and thus quantify the diffusion of the eccentricities. Note that the eccentricities of the outer planets have a very different CDF, which rules out the possibility that the good fit by a Rice distribution comes from the choice of the initial values of k and h.

From a geometrical point of view, $\mathbf{M}(P)^T\mathbf{W}\mathbf{M}(P)$ is the Fisher information matrix, which is also the matrix of the metric induced by $\mathbb{R}^N$ on $\mathcal{M}$. If it is ill conditioned, it means that there is a direction of the parameter space where an increase results in almost no change in the models: different values of the parameter give almost the same models and are therefore hard to distinguish from each other. The likelihood that the matrix becomes more ill defined grows with the period, since the terms $\cos n\boldsymbol{t}$ and $\sin n\boldsymbol{t}$ become increasingly close to the constant. Also, as the period of the planet P_t increases, $n = 2\pi/P_t$ and $2n$ might end up in the same local minimum of frequency and are therefore hard to disentangle. Furthermore, if the phase coverage is poor, $\cos n\boldsymbol{t}$ and $\sin n\boldsymbol{t}$ are not very well separated, again the condition becomes poorer. In the worst case, when measurement times are taken with a frequency n, both $\cos n\boldsymbol{t}$ and $\sin n\boldsymbol{t}$ are constant vectors. This is compatible with the observations of Shen & Turner (2008) and Zakamska et al. (2011), who noticed the increase of bias as the period increases and as the phase coverage degrades. For more details on this issue, we refer the reader to Appendix C.

One could also expect that since the model is in fact non-linear, the bias on the Keplerian model is higher not only because of poor condition number, but also because of the curvature of the Keplerian model. In fact, the bias does not automatically increase with the model curvature and depends on the geometry of the model. This one seems to be responsible for extra peaks of bias that are not spotted by the condition number, but are uneasy to predict at a given period.

Finally, we make two remarks. First, we now see clearly why when a Keplerian fit is initialized at a maximum of the periodogram, the eccentricity can go to one. Indeed, we are in a case where "the observations fall outside the target". In that case, the closest model will have eccentricity one. Secondly, this analysis also applies to the projected mass $m \sin i$ obtained via radial velocity measurements and the inclination obtained from astrometry. Indeed, these ones play the role of the radius of the target: the space occupied by the models that have a fixed $m \sin i$ or i is more and more voluminous as these parameters increase.

A non uniform parametrization

Another interesting fact is that the eccentricity does not parametrize uniformly the models. To see that qualitatively, we take the same grid as in figure 3.1 but two subsequent nodes are now spaced by the distance between their corresponding models. That is, if a node is at (e, ω), the node $(e + \Delta e, \omega)$ is represented at a distance $\|\boldsymbol{y}(\boldsymbol{t}, e + \Delta e, K, P, \omega, M_0) - \boldsymbol{y}(\boldsymbol{t}, e, K, P, \omega, M_0)\|$. The resulting plot is shown in figure 3.2. One can see that a constant increase in eccentricity does not represent equally spaced models: the more the eccentricity increases, the wider is the gap between the corresponding models.

We now highlight two consequences of this property of the parametrization by eccentricity. The first one is that this effect tends to make eccentricities biased downwards. Informally, this is due to the fact that since the models are farther away as eccentricity grows, when the true eccentricity is above a certain threshold the eccentricity is more often overestimated, but the "mean distance" between overestimated eccentricity and true eccentricity is lower than the "mean distance" between underestimated eccentricities and the true one. One can

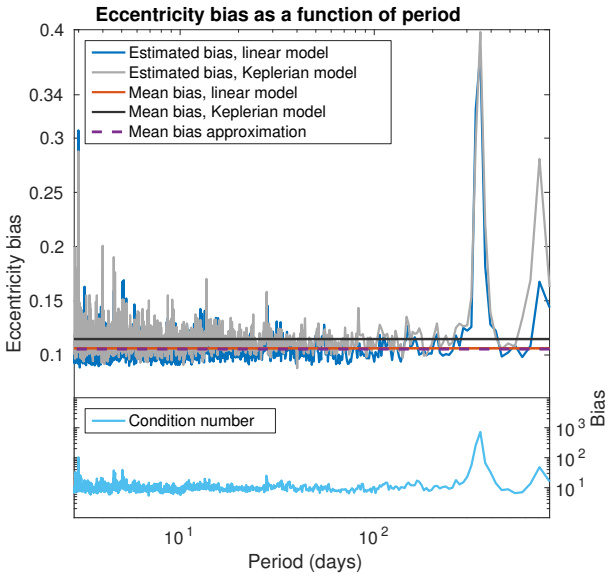

Figure 3.4: Bias as a function of the period for a system with 28 measurements, $K/\sigma = 3.5$ for the linearised model (blue) and the Keplerian model(gray) and their respective average over the frequency of the orbit (red and black). The values of "bias approximation', in purple, is $\frac{\sigma}{K_t}\sqrt{\frac{\pi}{N-p+1}}$.

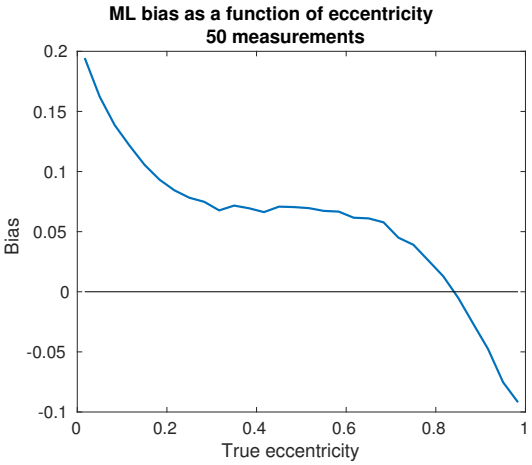

Figure 3.5: Bias of eccentricity estimate as a function of the true eccentricity

re-write the expectation of eccentricity knowing the true model to make this effect appear,

$$\mathbb{E}\{\widehat{e}|\boldsymbol{y}_t\} = \int_0^1 \widehat{e}\,\Pr(\widehat{e}|\boldsymbol{y}_t)d\widehat{e} \tag{3.13}$$

$$= \int_0^1 \widehat{e}d\widehat{e}\underbrace{\int_{\mathcal{M}_{\widehat{e}}} dm}_{\text{size}}\underbrace{\int_{\mathcal{Y}_m} \Pr(\boldsymbol{y}|\boldsymbol{y}_t)d\boldsymbol{y}}_{\text{Mean distance}}. \tag{3.14}$$

where $\mathcal{Y}_m$ is the set of all observations $\boldsymbol{y}$ for which the maximum likelihood model is m. The manifold $\mathcal{M}$ is foliated in submanifolds $\mathcal{M}_{\widehat{e}}$. Informally, for each point m of $\mathcal{M}_{\widehat{e}}$ the right hand side integral (with the "Mean distance" label) counts "how often" realizations of the noise give a maximum likelihood model estimate m, assuming the true signal is $\boldsymbol{y}_t$ and the noise is Gaussian with a given covariance. There is a trade-off between the size of the manifold $\mathcal{M}_{\widehat{e}}$ that gives a certain estimate $\widehat{e}$ of the eccentricity and on average, how far are the models $\boldsymbol{y} \in M_{\widehat{e}}$ from the true signal $\boldsymbol{y}_t$, $\mathcal{Y}_m$. To illustrate this, we generate 200 000 systems with eccentricity uniformly chosen between 0 and 1, uniform angles M_0 and ω and $K = 2$ m.s^{-1}. The systems are simulated on the first 50 measurement dates of GJ 876 (Correia et al. 2010). One sees in Fig. 3.5 that from $e_t \approx 0.82$, the maximum likelihood estimate of the eccentricity — which here is also the maximum *a posteriori* — is biased downwards. The downward bias is amplified towards the edge of the model, at $e = 1$.

This section illustrates that choosing a uniform prior on a parameter is not neutral. A uniform prior on e does not translate to a uniform prior on the models. Due to the fact that the Fisher information matrix is also a Euclidian metric when the noise is Gaussian, Jeffreys' prior Jeffreys (1946) can be seen as uniform on the models. Two set of models with equal volumes have an equal prior probability. However, this prior is not necessarily better. In our case, Jeffreys' prior would be proportional to e close to $e = 0$ and would penalize low eccentricity even more than the maximum likelihood. As a remark, one easily sees that a uniform prior on e is not uniform in $\sqrt{1-e^2}$ nor in e^2, while these are acceptable parametrizations.

Finally, Fig. 3.2 shows that high eccentricity models are several σ away from the null eccentricity one. This corroborates the well known fact that high eccentricities are difficult to detect when searching for circular orbits, with a Lomb-Scargle periodogram for instance. The difficulty of detecting very eccentric signals with the Lomb-Scargle periodogram was already mentioned in section 1.3.3. It therefore seems to make sense to use tools such as the Keplerian periodogram (Gregory 2007b; O'Toole et al. 2009; Zechmeister & Kürster 2009b; Baluev 2015b) to spot high eccentricity signals.

A complicated shape

The two previous sections are concerned respectively with the volume occupied by the set of models that have a fixed eccentricity and how that volume grows with eccentricity. The question of the shape of the set of models in $\mathbb{R}^N$ when e is high has not yet been treated. As we will see, it allows to better understand the difficulty of exploring the high eccentricity domain. In particular, we will see that it is likely to find a local minimum in this region that corresponds to nothing physical. Before detailing our analysis, let us consider the implication on estimation of the model shape.

Finding the best fitting model with a Gaussian noise amounts to finding the closest model to the observation in a geometrical sense. If the space described by the models with high eccentricity is bigger in the sense of a greater volume, one could think that there is a greater chance to find spurious planets in this region. However, it could turn out that though the space is bigger, it is compacted and folded in such a way that the set of observations that are closer to it is smaller than the set of observations closer to low eccentricity models, which is less voluminous but better spread. It is not the case, but rather the opposite. Not only the high eccentricity models occupy a wider space but they also are not confined to low dimension spaces.

The set of models $\boldsymbol{y}(\boldsymbol{t}, e, K, P, \omega, M_0)$ describe a 5-dimensional manifold in $\mathbb{R}^N$ if there are N observations ($\boldsymbol{t} \in \mathbb{R}^N$). Obviously, one cannot visualise directly the shape of a manifold at least 5-dimensional in a N-dimensional space. To circumvent this limitation, we project the models onto lower dimensional spaces. But first, we can already get rid of the semi amplitude parameter K. Indeed, instead of searching the model $\boldsymbol{y}(\boldsymbol{t}, \boldsymbol{\theta})$ that minimizes the distance to observations, we can search the model $\boldsymbol{u}(\boldsymbol{\theta}') = \boldsymbol{y}(\boldsymbol{t}, \boldsymbol{\theta})/\|\boldsymbol{y}(\boldsymbol{t}, \boldsymbol{\theta})\|$, where $\boldsymbol{\theta} = (K, \boldsymbol{\theta}')$, that maximizes the correlation with the observations. Indeed,

$$\boldsymbol{\theta}'^{\star} = \arg\min_{\boldsymbol{\theta}' \in \Theta'} \left(\arg\min_{K \in \mathbb{R}} \sum_{k=1}^{N} (y_k - y(t_k, \boldsymbol{\theta}))^2 \right) = \arg\max_{\boldsymbol{\theta}' \in \Theta'} \left| \left\langle \frac{\boldsymbol{y}(\boldsymbol{t}, K, \boldsymbol{\theta}')}{\|\boldsymbol{y}(\boldsymbol{t}, K, \boldsymbol{\theta}')\|}, \boldsymbol{y} \right\rangle \right|. \tag{3.15}$$

Here $\boldsymbol{\theta}'$ designates the vector of parameters without K. This remark allows us to limit our analysis to a 4-dimensional manifold, namely the one described by $\boldsymbol{u}(\boldsymbol{\theta}')$ as ω and M_0 describe $[0, 2\pi]$, P and e respectively describe the positive real numbers and $[0, 1]$. We can further reduce the dimension since we are in the estimation setting. In that case, a periodicity has already been secured by a periodogram analysis so that we can assume that P is fixed. We are left with three parameters, e, ω, M_0. For the sake of simplicity, we fix M_0 to 0, which is equivalent to shift the time scale and does not impact the conclusion of the following analysis.

We now have a normalized model that depends on two parameters, $\boldsymbol{u}(e, \omega)$, and we are interested in the shapes in $\mathbb{R}^N$ it draws as ω describes $[0, 2\pi]$ for each $e \in [0, 1]$. It is proven in Appendix C that for small e, one can think of the model space as concentric circles of radius e all lying in a two dimensional space generated by the vectors $\cos(2n\boldsymbol{t})$ and $\sin(2n\boldsymbol{t})$. We test if the manifolds $\mathcal{U}(e) = \{\boldsymbol{u}(e, \omega),\ \omega \in [0, 2\pi])\}$ all lie in a two dimensional space by searching for each $e = 0..1$ what is the m dimensional linear subspace of $\mathbb{R}^N$ which best approximates $\mathcal{U}(e)$. More precisely, given e and a linear space of dimension m, $\mathcal{S}_m$, for each ω, we compute the distance between $\boldsymbol{u}(e, \omega)$ and its projection on $\mathcal{S}_m$. This distance is integrated on ω, and we obtain a measure of the total distance between $\mathcal{U}(e)$ and $\mathcal{S}_m$. We then aim at finding the space $\mathcal{S}^\star_m$ that minimizes this total distance.

A solution to this problem is given by the Karhunen-Loève theorem and a discretization, as detailed in Appendix D. Here we present our procedure. Its first step is to discretize the values taken by ω, ω: $\omega_k = 2\pi k/q,\ k = 0 : q - 1$ for some integer q. Then, the models at a given eccentricity are concatenated so to form a $n \times q$ matrix, $\mathbf{M}_q = [\boldsymbol{u}(e, \omega_0), ..., \boldsymbol{u}(e, \omega_q)]$[2]. One can then perform a singular value decomposition,

$$\mathbf{M}_q = \mathbf{PSV}^T \tag{3.16}$$

where $\mathbf{P}$ and $\mathbf{V}$ are orthonormal matrices and $\mathbf{S}$ is a diagonal matrix whose entries are called singular values, and are the eigenvalues of $\mathbf{M}_q\mathbf{M}_q^T$. The subspace $\mathcal{S}^\star_m$ is generated by the vectors corresponding to the m greatest singular values.

If the models of given eccentricity $\mathcal{U}(e)$ lie in a subspace of dimension 2, the diagonal of the matrix $\mathbf{S}$ should have only two non null values (see equation (3.16)). The value of the diagonal elements is also informative: a non-null singular value indicates that a dimension is explored by the model but the smaller it is, the lower is the component of $\mathcal{U}(e)$ along the corresponding direction.

Let us apply this methodology to an example. We consider the 138 measurement times of GJ3998 (Affer et al. 2016). We fix a period of 20 days and $M_0 = 0$. We then take $\omega_k = 2\pi k/q,\ k = 0 : q-1\,; q = 1000$ and apply the procedure to $\mathcal{U}(e)$ for $e = 0.1, 0.2...0.9$ and $e = 0.999$. The singular values are then ordered and plotted Fig. 3.6 (left). We see that as the eccentricity grows, the number of explored dimensions grows, indicating that the model shape is more and more complicated. Another way of visualizing this complexity is to concatenate normalized models $\boldsymbol{u}(e_j, \omega_k)$ for ω_k previously defined and $e_j = e_{\max} j/q,\ j = 0 : q-1\ q = 100$ with $e_{\max} = 0.6$. We then obtain a set of principal directions, of which we extract the three corresponding to the highest singular values: x, y, z. We then plot the coordinates of the vectors of $\mathcal{U}(e)$ on x, y, z (see Fig. 3.6, right). We see that the projection up to $e \approx 0.8$ is regular, but at higher eccentricity it becomes much more erratic.

[2]The same reasoning applies to a matrix $\mathbf{M}'_{2q}$ that would concatenate $\dot{U}$ and $\dot{V}$, as defined in equation (1.26). Doing so would isolate the non-linear behaviour.

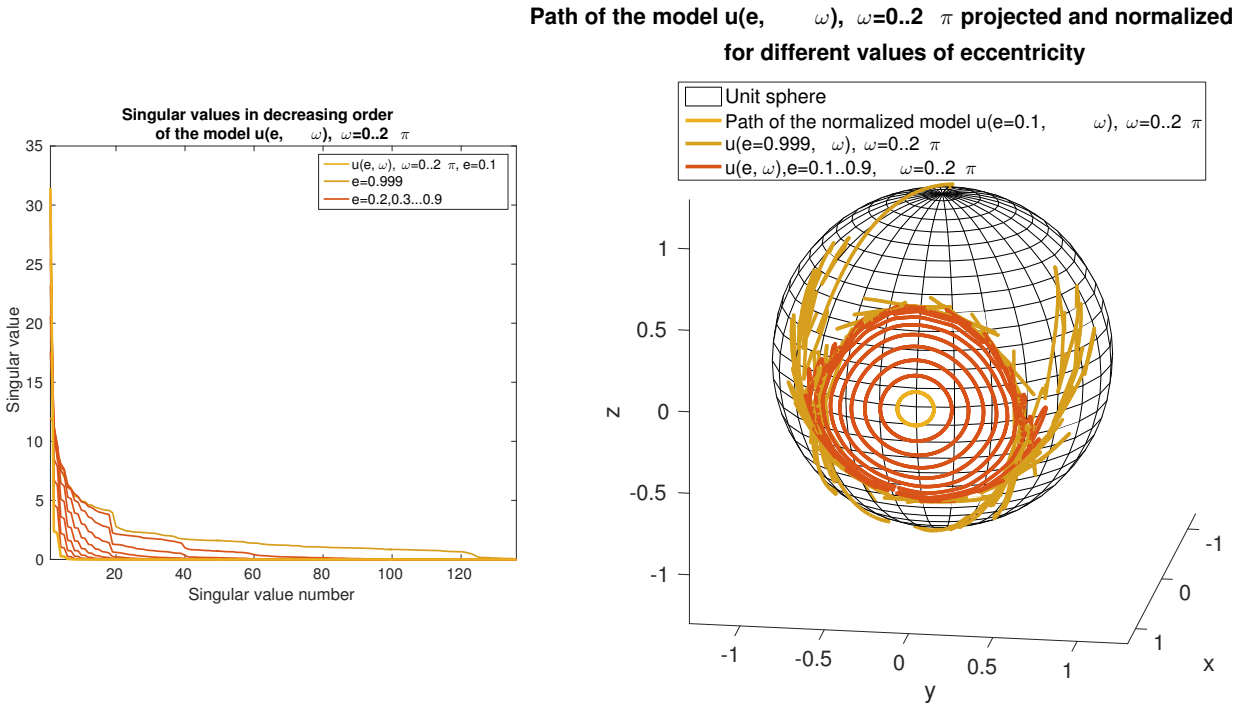

Figure 3.6: Projections of paths of $(u(t,e,\omega))_{\omega=0..2\pi}$ for different eccentricities onto the set $S_3^\star((u(t,e,\omega))_{e=0.1,0.2...0.7})$ for P=20 days, $M_0 = 0$

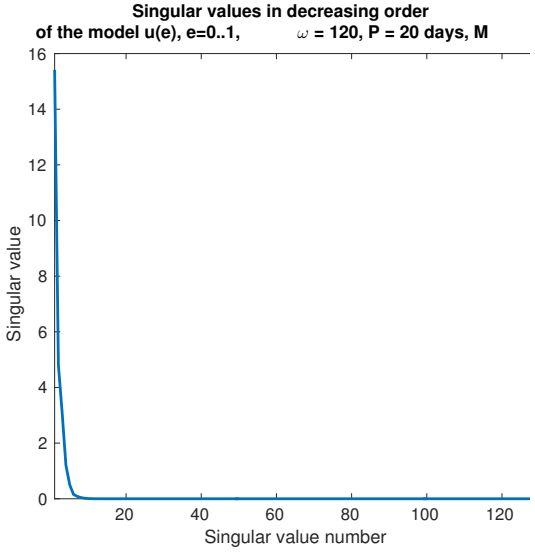

Figure 3.7: Singular values of the model as e varies.

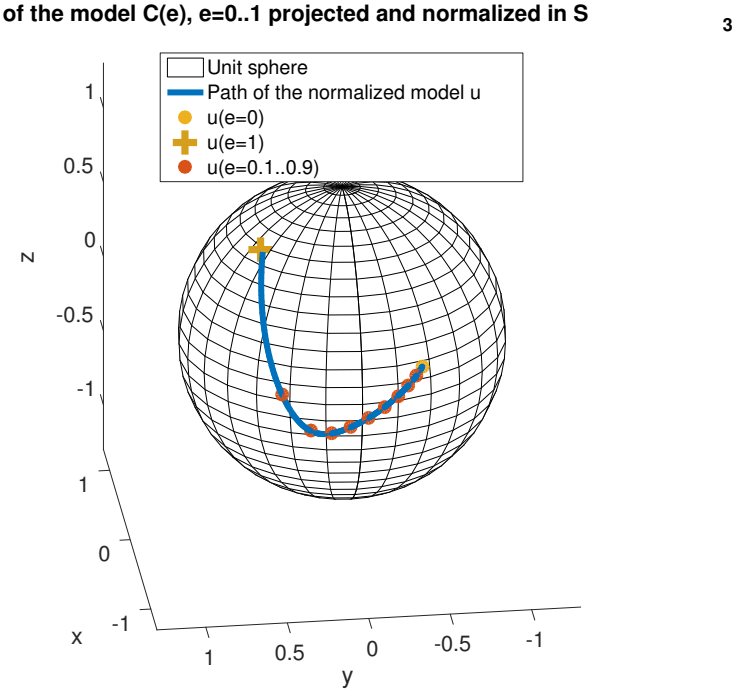

Figure 3.8: Projection of $(u(t,e))_{e=0..1}$ onto the space $S_3^\star((u(t,e))_{e=0..1})$ for P=20 days, $M_0 = 0$ and $\omega = 120°$

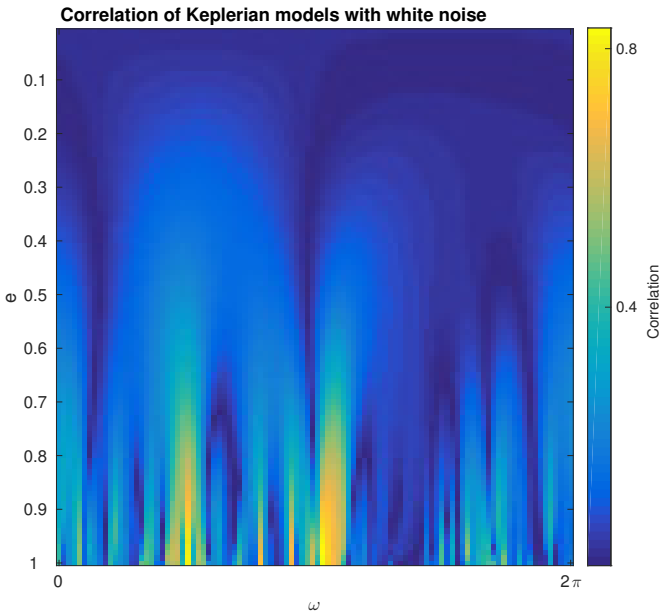

Figure 3.9: Correlation between a Keplerian model sampled on 30 point and white Gaussian noise. More precisely, we represent that is $|\langle \boldsymbol{u}(e,\omega,M_0,P), \boldsymbol{y}/\|\boldsymbol{y}\|\rangle|$ for fixed period P and initial mean anomaly M_0 and a grid of values of e and ω, where $\boldsymbol{u}(e,\omega,M_0,P,K)$ is the Keplerian model divided by its Euclidian norm and $\boldsymbol{y}$ is a white Gaussian noise realization. The light yellow corresponds to a high correlation, which appears at high eccentricities, dark blue corresponds to a low correlation.

Finally, we can apply the same procedure but here on a discretization of the values of the eccentricity, the argument of periastron being fixed. Again, we concatenate the vectors, compute the singular values and represent them in Fig. 3.7. This shows that most of the variation is contained in a three-dimensional space, with small components on two additional dimensions. We therefore can represent accurately the set of models by plotting the coordinates of the vectors of $\omega = 0$, e varying from 0 to 1 on the three principal components (see Fig. 3.8). The models corresponding to $e = 0, 0.1, ..., 0.9, 0.999$ are highlighted with markers. Again we see that from $e \approx 0.8$, the models explore another dimension and as seen in the previous section, e does not parametrize uniformly the models.

For the sake of brevity we will not present other examples but just mention that we observed very similar features on other measurement times. The complexity of the model shape increases with eccentricity and especially for $e \gtrsim 0.8$. This is fully compatible with the observation of Baluev (2015b): the number of local minima increases significantly when one searches for high eccentricities. We represent the correlation of Keplerian models with e and ω varying (30 measurement) with a generated white noise Fig. 3.9. It shows that for high eccentricities there are numerous maxima of correlations, corresponding to minima of χ^2. Interestingly enough, it looks very much like a time-frequency representation, e playing here a role similar to the time-scale. Indeed, when e grows, the signal is non-zero on smaller and

smaller support, just like a wavelet of more and more precise time localisation. The argument of periastron cannot be seen as a frequency though, but rather as a phase.

Numerical effects

On top of the biases generated by geometrical features, an incorrect implementation of the numerical schemes can worsen the situation. We here give three examples concerning the estimation of eccentricity from radial velocity data, some of which were noted in previous works: an increased bias due to an error in the MCMC implementation, the possible failure of the Kepler equation solver at high eccentricities, and the fact that eccentricity might be spuriously stuck to zero.

Indeed, as noted by Eastman et al. (2013), there is a specific error in the implementation of the Metropolis-Hastings algorithm that worsens the bias when the true eccentricity is zero and when the parameter space is parametrized by (e, ω) instead of (k, h). Such an algorithm is based on Markov chains that explore the parameter space, sometimes called walkers that give a sequence of parameters $(\boldsymbol{\theta}_k)_{k=1..}$, whose distribution approximates a certain posterior distribution. The generation of the sequence proceeds as follow: from $\boldsymbol{\theta}_k$, $\widetilde{\boldsymbol{\theta}}_{k+1}$ is generated according to a proposition distribution. Depending on whether $\widetilde{\boldsymbol{\theta}}_{k+1}$ is more likely than $\boldsymbol{\theta}_k$, it will be accepted with a certain probability. The error consists in not assigning $\boldsymbol{\theta}_{k+1} \leftarrow \boldsymbol{\theta}_k$ if $\widetilde{\boldsymbol{\theta}}_{k+1}$ is rejected. Close to zero eccentricity, there will be a higher rejection rate due to proposals that verify $e < 0$, and the density of parameters in this region will decrease. This problem comes from the fact that the model has a parameter boundary (which is not a boundary in the geometrical sense). The same remark can be done on the boundary $e = 1$, to see that the incorrect implementation increases the downward bias in that region. A possible workaround to avoid the numerical effect near zero is to use k and h instead of e and ω to parametrize the model. Near $e = 1$, the boundary comes from the model geometry, therefore the numerical effect cannot be avoided by a re-parametrization.

There is another potential threat in the latter region: the validity of the Kepler solver. Indeed, some numerical schemes aiming at solving Kepler equations lead to spurious results at high eccentricity. If this is the case, there are some signals that cannot be fitted simply because they cannot be generated.

Our last example is well-known to observers: sometimes, when fitting a Keplerian curve initialized at a certain period, the eccentricity is stuck to zero. This situation is for instance encountered in Brown (2017). Our analysis of the geometry of the models shows that there is no discontinuity of the model manifold at $e = 0$. Informally, the "target" (Fig. 3.2) has an edge in $e = 1$ but is continuous at $e = 0$, only the parametrization has a singularity. Therefore, we should not have an eccentricity exactly null after a local minimization unless there is a numerical problem. Indeed in the examples we could find of this phenomenon, the eccentricity is equal to zero at every step of the minimization, so that the outputs are equal to the inputs. The reason of this behaviour is that the Newton steps computed by the solver are too wide, so that the next step is systematically rejected. The weight of that step is reduced iteratively, still without success, until it falls below the threshold that is supposed to indicate convergence. This happens in particular when adding a trend to the model without translating the times of measurements from BJD to zero.

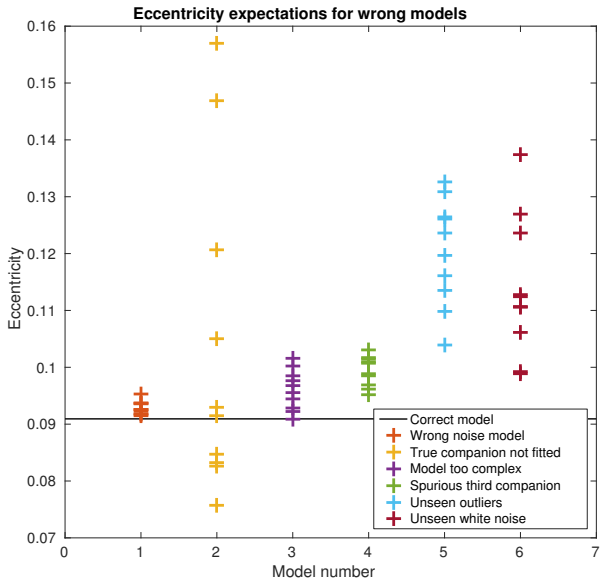

Figure 3.10: Eccentricity estimates for different realizations of modelling errors

3.3.4 Conclusion for individual systems

The description of the bias done in previous sections clarifies some features of the inference on eccentricities. First, let us note that the bias increases with the noise amplitude. For instance, when the true eccentricity is zero, the bias is inversely proportional to the signal-to-noise ratio (see equation (3.12)). A similar result was known since Lucy & Sweeney (1971), but we add that any additional noise or model misspecification can worsen the bias. First, as the noise variance grows the bias gets higher approximately proportionally to the noise standard deviation in equation (3.12)). Secondly, as the model complexity increases, the uncertainty on the parameter increases and so does the bias. This can be seen in formula (3.12), as the number of parameters p grows the bias increases. Finally, when the model is too simplistic, since there are "more" models with high eccentricity, there are more chances that the maximum likelihood model has a higher eccentricity — though in some cases, too simple a model might decrease the bias at some values of e_t.

To illustrate this, we generate 100 systems with two planets plus white and red noise (model $\mathcal{M}_t$). For each realization we find the least square eccentricity with models $(\mathcal{M}_k)_{k=1..7}$ that are wrong for different reasons ((1) wrong noise model, (2) models too complex ...). For each subclass of wrong models we adjust some of their parameters randomly (for instance the period of an unseen planet) and compute the mean eccentricity retrieved on the 100 trials. Figure 3.10 shows the result of the experiment and reveals that eccentricities are over-estimated on average when the model is incorrect. We can sum up the situation for low eccentricities by: the more uncertain, the higher the bias. In the uncertainty we include standard noise but also the uncertainty on the model, that is the error we do by picking a wrong model for the data.

The fact that the bias is stronger than we might think is especially problematic if the error bars are incorrect or misleading. If the noise is not fitted, for example one only considers the nominal uncertainties of the observations, then the error bars will be under-estimated.

If the noise is fitted with a realistic model, then the error bars will be more appropriate but still subject to caution since they depend on assumptions, such as analytical approximations, absence of prior distributions for confidence intervals, or the prior distribution in Bayesian analysis. In the latter case, if one computes credible intervals centred on the *a posteriori* maximum, since this maximum is biased, one would have a systematic tendency to exclude $e = 0$ from this interval. This does not mean that the interval is incorrect, but that other intervals containing $e = 0$ might also have a high probability.

In summary, there is a systematic effect (the bias of eccentricity) that is underestimated, and on which the uncertainty is also underestimated. In section 3.4 we consider several options to address this issue.

3.3.5 Conclusion for population analysis

The fact that eccentricity is biased at low and high eccentricity has another consequence: the distribution of best fit (or maximum likelihood) eccentricity is unreliable in itself. To illustrate it, we generate 200,000 systems with an eccentricity distribution D_1, we initialize a non-linear least-square fit at the correct period and draw the distribution of estimated eccentricity D_2. Figure 3.11 (left) represents the result of such an experiment with D_1 being uniform. The angles ω and M_0 are chosen uniformly over $[0,\ 2\pi]$, $K = 2$ m.s^{-1} and $P = 10$ days are fixed, the measurement dates are the 50 first measurements of GJ 876 (Correia et al. 2010), the noise standard deviation is 1 m.s^{-1}. We do the same experiment but inject systems whose eccentricity is null with a 30% chance and follows an exponential law otherwise (see Fig. 3.12, left). In both cases, we plot the density of the two-dimensional variable (e measured, true e). For each interval of measured eccentricities, we compute the mean, median and mode of the distribution of true eccentricities (Fig. 3.11 and 3.12, right).

The first striking feature is that the true and measured distributions are different. In the case of Fig. 3.12.a, there is a clear deficit of measured eccentricities close to zero. Now if we do the same simulation but increase the noise by a factor two we obtain a distribution that is farther away from the true one. The purple curve results from fitting Keplerian curves to pure noise. This one is fully consistent with our geometrical analysis: since the space occupied by high eccentricity models grows with eccentricity, a vector whose direction is uniformly chosen (which is the case of Gaussian noise) should on average be closer to a high eccentricity model. Approximately 40% of the fits have an eccentricity stuck to one. This is also easily understandable geometrically. Indeed, the model has an edge at $e = 1$, just like a target as we noted in section 3.3.3. If due to noise, an observation is so far from the considered models that it falls "outside the target", the closest model will always be on its edge. The fraction of models outside the edge is smaller in our first simulations but is non null. This sheds some light on a situation well known to observers, sometimes when fitting the eccentricity this one goes to one. The suspicion that the result is spurious is justified, this is probably due to noise.

There is another interesting feature revealed by Fig. 3.11 and Fig. 3.12 (right). When finding an eccentricity close to 0 - 0.3 in a case of non negligible noise, one could be tempted to test systematically if the zero eccentricity can be rejected by a null hypothesis test and if not, fix the eccentricity to 0, as done in Lucy & Sweeney (1971) and Zakamska et al. (2011). The accuracy of this procedure depends on the true distribution of eccentricity. If there is indeed a significant number of true eccentricities close to 0, then Fig. 3.12 (right) shows that

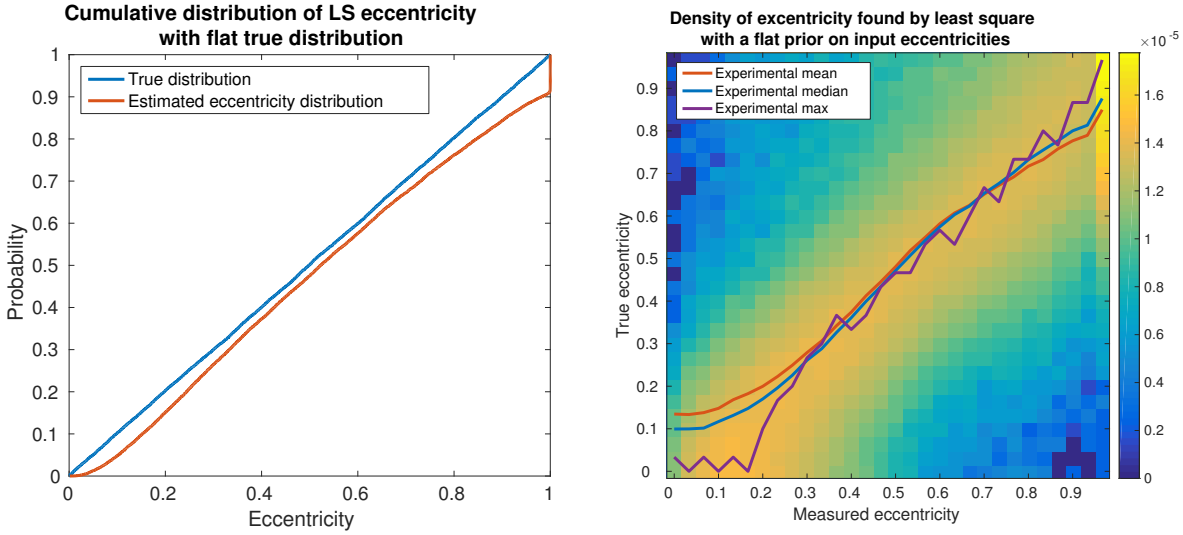

Figure 3.11: Left: Cumulative distribution function of the input distribution (blue) and retrieved distribution (red). Right: distribution of the eccentricity measured and the input eccentricity (in color code). The curves represent the mean, median, and mode of the distribution of true eccentricity given the measured eccentricity.

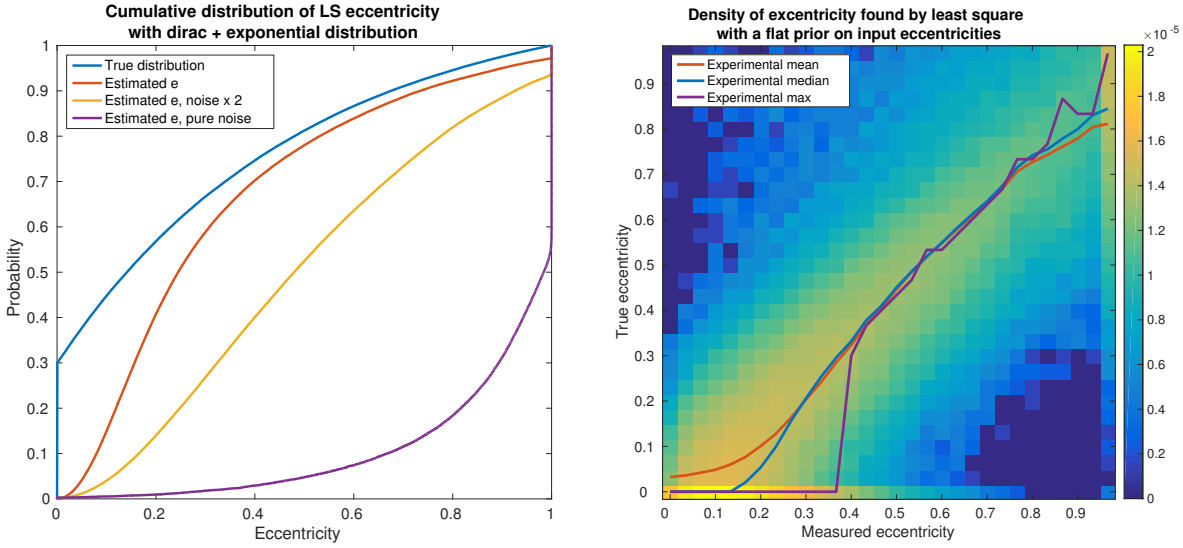

Figure 3.12: Left: Cumulative distribution function of the input distribution (blue) and retrieved distribution for different level of noise. (red: 1m/s, yellow: 2 $m.s^{-1}$, purple: the input signals are pure independent Gaussian noise realizations). Right: distribution of the eccentricity measured and the input eccentricity (in color code). The curves represent the mean, median, and mode of the distribution of true eccentricity given the measured eccentricity.

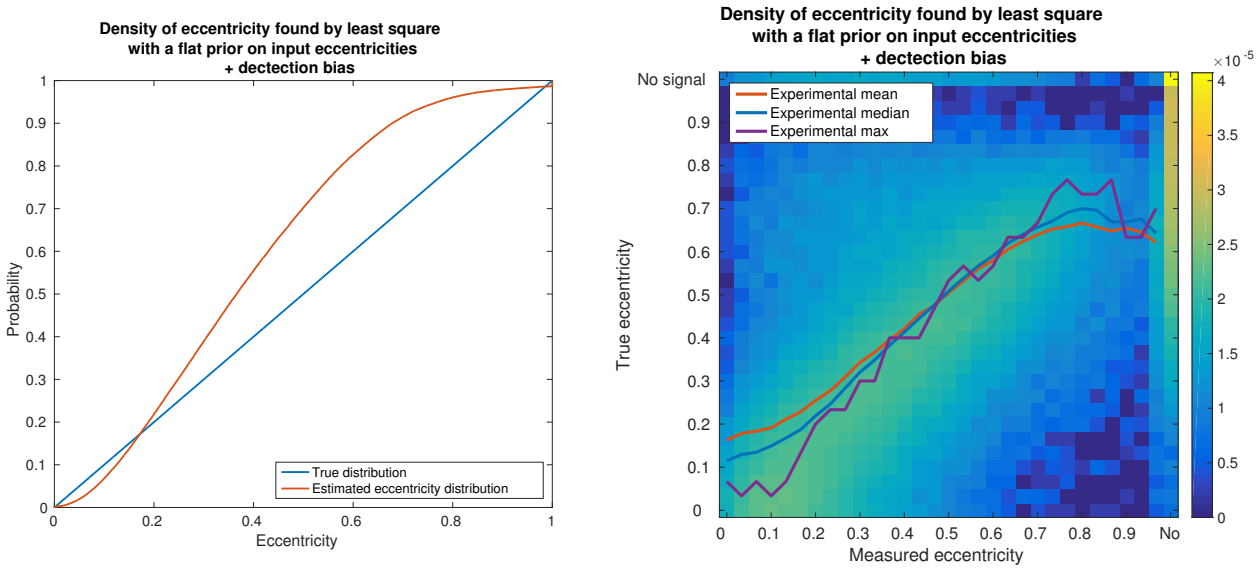

Figure 3.13: Left: Cumulative distribution function of the input distribution (blue) and retrieved distribution for different level of noise taking into account the detection bias. Right: distribution of the eccentricity measured and the input eccentricity (in color code). The curves represent the mean, median, and mode of the distribution of true eccentricity given the measured eccentricity.

when the measured eccentricity is between 0 and ≈ 0.37, the most likely situation is that $e_t = 0$ (purple curve). However, for the same range of measured eccentricities, the median and mean true eccentricities reach values as high as 0.25 (red and blue curves). This issue is worse in the case of a flat true distribution: even when the fitted eccentricity is close to 0, on average the true eccentricity is ≈ 0.13 and the median is ≈ 0.1. Obviously a flat distribution is very unlikely to actually happen but these examples show conceptually that accepting a zero eccentricity if it cannot be rejected is not optimal. The problem of isolating the case $e = 0$ prevents us from seeing the possible subtleties of the eccentricity distributions and by forcing to look at a specific case, it might lead us to indeed see what we want (or expect) to see. Some more accurate alternatives are outlined in section 3.4.3.

The simulations made so far do not take into account the detection bias. To evaluate the effect of a more realistic situation, we do the same experiment as above but with a small modification. Once the signal is generated, a periodogram is computed between 0 and 1.9 cycles per day and if a peak is higher than the 1% FAP threshold, a Keplerian fit is initialized at the corresponding frequency and the best fit eccentricity is reported. An important difference with previous simulations is that 40% of the signals generated are pure noise. Figure 3.13 (left) shows the cumulative distribution function of the best fit eccentricities. This one has a slope greater than one between ≈ 0.05 and 0.6 which indicates an overestimated density in that region, and underestimated density elsewhere. In figure 3.13 (right) we add a line labelled "no signal" which accounts for the case where the signal is absent. An extra column is also added to account for non detections. Globally, the effect of detection

bias is to shrink the distribution between 0.1 and 0.7, since very low eccentricities are detected but overestimated and high eccentricities are under-detected and not fairly represented in the estimated distribution. Though realistic, the detection criterion we adopt is not optimal. A Keplerian periodogram (Cumming 2004; Zechmeister & Kürster 2009a; Baluev 2015b) would be more efficient.

The present discussion considers planets that all have the same semi-amplitude. Shen & Turner (2008) consider planets of equal masses. In that case, the semi-amplitude of the signal is stronger for high eccentric orbits, since $K \propto \frac{m \sin i}{\sqrt{1-e^2}}$. The increase in semi-amplitude reduces the detection bias and increases the signal-to-noise, reducing the estimation bias. They find that in the constant mass case, the bias is much less of a problem when estimating distributions. We simply point out that they assume that the observations are uniformly distributed over two periods, which is likely an overly optimistic assumptions. Uniformly distributed measurements are in fact extremely efficient measurements, the daily repetition should lower the detection efficiency. A comprehensive simulation of the measurement process might show that stars with highly eccentric planets are given up after a few measurements if they are sampled far from their periastron passages. Such a simulation is beyond the scope of this work. Our point on a global correction of distributions instead of assigning $e = 0$ or $e = \widehat{e}$ is still valid.

3.4 Improving the inference robustness

3.4.1 Prospects

Now that some cause for concern has been raised, several approaches can be taken to improve the situation:

- Having criteria to ensure that there are enough measurements to exclude a spurious eccentricity taking into account possible modelling errors.
- Studying past observations: Spotting already observed systems whose eccentricity might be very ill constrained.
- Trying to obtain better estimates on each system, by choosing another estimator than the maximum likelihood, or using bias correction methods such as the Jacknife (Quenouille 1949).
- Trying to correct the eccentricity distributions, but not systems one by one.
- Defining observation strategies based on requirements on population retrieval.
- Improving models, especially stellar noise models, so they have the correct amount of complexity.

The latest item being a subject in its own, it will not be considered here. Regarding the other options, one must bear in mind that obtaining orbital elements that are useful for drawing scientific conclusions is not uniquely a problem of minimizing bias. In general, decreasing the bias of an estimator increases its variance, and a trade-off is to be found (see Lehmann & Casella (1998) chapter 2 for example). One could minimize the squared

bias plus the variance, that is the mean squared error but this metric depends on the prior distribution of the parameter, which is unknown. Secondly, this metric is reasonable but not universal, one could privilege minimizing the absolute mean error $\mathbb{E}\{|\widehat{e} - e|\}$ or any other risk. However, the estimators minimizing a given risk are not always easy to express or to compute.

A notion of risk that seems relevant for our purposes could be stated informally as: "is the scientific conclusion drawn from the observations correct?". Most statistical methods address this issue, but here we will design specific methods for orbits estimation. The conclusions of such fits concern mostly formation and dynamics. We would not want an utterly wrong estimate that would lead us on to the wrong migration path for a particular system, nor having claims on populations based on observational features that are due to statistical biases.

At this point it seems reasonable to try to de-bias observations. However, we must keep in mind that the properties of any correction methods are established within a certain model, and that model is never completely true. Secondly, the correction methods are based on assumptions that might lead us to "see what we want to see" in the data, and miss out interesting features. For instance, Lucy & Sweeney (1971) and Zakamska et al. (2011) systematically set to zero the eccentricities that are not significant.

Thus, our general philosophy is not to try at all cost to reduce the error on the point estimates, but to devise tools that can help to draw robust conclusions from the data. More precisely, we advocate to use interval estimates. We also suggest a method to design observations based on requirements on the population retrieval. The next sections are devoted to detailing these options and when relevant applying them to known data sets.

3.4.2 Individual systems

Using hypothesis rejection instead of point estimates

When analysing a given system, one often wants to determine if a planet is in the habitable zone, what are the possible formation scenarios etc. In both cases, the claims will be different if the eccentricity of a planet under scrutiny is below or above 0.1, and also different if it is greater than 0.3, say. Let us know suppose that we have strong suspicion that the eccentricity is between 0.1 and 0.3 which is our hypothesis $\mathcal{H}$. As said in section 1.3.1 the methodology we suggest follows the Popperian tradition, and aims at excluding all other explanation for the signal: we want to quantify the confidence that the complementary hypothesis $\mathcal{H}^c$ that eccentricity is anywhere but between 0.1 and 0.3 is incorrect. In the spirit of the previous section, we first reason assuming the observation model is correct. As previously, we distinguish two cases: testing this hypothesis assuming the model is true, and assessing whether this claim is endangered by a model misspecification.

Before moving forward, let us consider a classical remark on the point estimates (e.g. MacKay 2003). In the Bayesian framework one computes the posterior density of a vector of parameters $\boldsymbol{\theta}$,

$$p(\boldsymbol{\theta}|\boldsymbol{y}) = \frac{p(\boldsymbol{y}|\boldsymbol{\theta})p(\boldsymbol{\theta})}{p(\boldsymbol{y})}. \tag{3.17}$$

To estimate the eccentricity, one can compute the maximum of this posterior probability or its expectancy. These two first options are point estimate which depend not only on the

prior, but also on the parametrization of the prior. The median and the probability that the parameters lie in a certain domain do not. For instance it means that if we parametrize the model with e, or $\eta = \sqrt{1-e^2}$, the posterior maximum and expectancy will give two orbital solutions. On the contrary the probability that e lies in $[e_1, e_2]$ is equal to the probability that η lies in $[\sqrt{1-e_2^2}, \sqrt{1-e_1^2}]$.

Bayesian methodology

From an abstract point of view, we are in the following situation: we have a model that can include a parameter or not (eccentricity). We wish to determine if this one should be included in the model and if yes, what is the range of values it can reach. We will denote the parameter by $e \in I$ but this applies to any other one. The first problem is classically addressed by computing the ratio of posterior likelihoods as

$$R = \frac{\Pr\{e \neq 0|\boldsymbol{y}\}}{\Pr\{e = 0|\boldsymbol{y}\}} = \frac{\Pr\{\boldsymbol{y}|e \neq 0\}}{\Pr\{\boldsymbol{y}|e = 0\}} \frac{\Pr\{e \neq 0\}}{\Pr\{e = 0\}} \tag{3.18}$$

where $\Pr\{\boldsymbol{y}|e \neq 0\} = \int_{\boldsymbol{\theta}\in\Theta} p(\boldsymbol{y}|\boldsymbol{\theta})p(\boldsymbol{\theta})d\boldsymbol{\theta}$. If this ratio is superior to a certain value, then one favours $e \neq 0$ over $e = 0$. As the number of samples N tends to infinity, this ratio is equivalent to the Bayesian Information Criterion (BIC, Schwarz 1978). Furthermore, one can compute a credible set, that is an ensemble of e, denoted by C such that

$$\Pr\{e \in C|\boldsymbol{y}\} = \int_C p(e|\boldsymbol{y})de = \frac{\int_C p(\boldsymbol{y}|e)p(e)de}{\int_I p(\boldsymbol{y}|e)p(e)de} = \alpha \tag{3.19}$$

where $\alpha \in [0, 1]$ is a probability. The set C is in general taken as an interval but this need not be the case. Let us note that this approach is also a ratio of posterior likelihood as in Eq (3.18). If $e = 0$ is given a non null probability, then the prior probability takes the form $p(e) = p'(0)\delta(0) + p'(e)$ where δ is the Dirac function. The probability that e lies outside C is $1 - \alpha$, so computing a credible interval might be seen as evaluating the probability that e lies elsewhere than in the interval under consideration and possibly reject this hypothesis above a certain α. If 0 is in the rejected set, circular orbits are rejected.

The robustness of these approaches are endangered by two sources of uncertainty: the prior and the model. Indeed, for different choices of these, one might obtain different values of the quantities (3.18) and (3.19), and therefore take different decisions about the data. For instance, it might turn out that when the model is given more degrees of freedom to fit the star activity, then a credible interval might broaden and include zero.

The most straightforward way to explore the dependency of the Bayes factor on the prior is to recompute it with another prior distribution. However, this might be lengthy. MCMC analysis takes typically 20 minutes and full Bayes factor calcultion up to several days, as pointed out in section 1.3.6. An estimate of the dependency can be obtained with the so-called Laplace formula, given in equation 14 of Kass & Raftery (1995). We here propose an alternative, which consists in multiplying the prior by a constant on a subset of its domain of definition, and to scale it elsewhere. In so doing, the output of the posterior sampler can be used straightforwardly without doing any sampling. This can be used for instance to enhance the prior at low eccentricities, and evaluate how credible intervals of eccentricity evolves. More precisely, we proceed as follows.

We consider a measurable subset D of $[0,1]$. For instance an interval $[0,e_0]$ for some arbitrary $e_0 \in (0,1]$. We define a new prior p' such that for $e \in D$, $p'(e) = xp(e)$. To ensure that $\int_0^1 p'(e)de = 1$, we take $x \in [0,(1-a)/a]$ and $p'(e) = (1-ax)/(1-a)p(e)$ for $e \in \bar{D}$, where $\bar{D}$ is D complementary set in $[0,1]$ and $a = \int_D p(e)$. We now want to compute the probability that $e \in C$ for a prior distribution p'. This one is given by replacing p by p' in equation (3.19), giving

$$\Pr\{e \in C|\boldsymbol{y}\} = \int_C p'(e|\boldsymbol{y})de = \frac{x\int_{C\cap D} p(\boldsymbol{y}|e)p(e)de + \frac{1-ax}{1-a}\int_{C\cap\bar{D}} p(\boldsymbol{y}|e)p(e)de}{x\int_D p(\boldsymbol{y}|e)p(e)de + \frac{1-ax}{1-a}\int_{\bar{D}} p(\boldsymbol{y}|e)p(e)de} \tag{3.20}$$

Where $p'(e|\boldsymbol{y})$ is the posterior distribution when the prior is $p'(e)$ and $p(\boldsymbol{y}|e)$ is the likelihood marginalized on all parameters but eccentricity. When $x = 0$, all the prior probability goes to the complementary of D and $\Pr\{e \in D|\boldsymbol{y}\}$ goes to 0. If $x = (1-a)/a$, e is certainly in D therefore $\Pr\{e \in D|\boldsymbol{y}\} = 1$. The advantage of this calculation is that the integrals can be computed from the posterior samples. The reasoning can be extended straightforwardly to credible regions D and C in the parameter space, and to prior region subdivisions in $D_1...D_q$ with disjoint $(D_i)_{i=1..q}$ whose union is the whole parameter space. Since the integrals to be evaluated from posterior samples are random variables, it must be ensured that they have a controlled uncertainty. When breaking the posterior in many domains $D_1...D_q$, the procedure we outline may become unreliable. Further investigation is left for future work.

The uncertainty on the model including the likelihood can be computed via the Bayes model averaging technique, wich consists in computing

$$\Pr\{e \in C|\boldsymbol{y}\} = \sum_{k=1}^{M} \Pr\{e \in C|\boldsymbol{y},\mathcal{M}_k\}\Pr\{\mathcal{M}_k\} \tag{3.21}$$

for a family of M candidate models $(\mathcal{M}_k)_{k=1..M}$ which were assigned a probability. Several techniques are available for Bayesian variable selection (see for instance Hoeting et al. 1999; O'Hara & Sillanpää 2009).

Our aim here is not to enter in the details of these methods, but to outline solutions for assessing the dependency of a Bayesian inference on the prior. Bayesian model averaging could be overly complicated in the case of orbital fits. We will motivate a simpler approach for handling uncertainties in paragraphs "Minimizing the impact of model errors" and "Detecting model errors". This one consists in selecting among the "reasonable" model the most complex one, scaling the errors to obtain a χ^2 equal to one, then checking that the residuals behave roughly as Gaussian independent variables.

Frequentist methodologies

Testing if the eccentricity is 0 can also be performed in a frequentist framework. This one has the advantage to limit the use of *a priori* assumptions, but does not allow to ask the correct question. We will see though that correctly calibrated, confidence intervals are useful tools to make inferences on eccentricities.

So far, (Lucy & Sweeney 1971) and Husnoo et al. (2012) respectively used p-values and Bayesian Information Criterion (BIC) to test the hypothesis that eccentricity is non zero. More precisely, Lucy & Sweeney (1971) compute the probability distribution of the eccentricity estimate under the hypothesis that the eccentricity is null and find a Rayleigh distribution

whose variance depends on the signal-to-noise ratio (which we also obtain as a special case of our analysis section 3.3.3 with $p = 0$). For a given measured eccentricity $\widehat{e}$, they measure the probability that the Rayleigh distribution is higher than $\widehat{e}$ and report an eccentric orbit if this probability is lower than a certain threshold.

Husnoo et al. (2012) computes

$$\mathrm{BIC}(\mathcal{M}) = \chi^2_{\min}(\mathcal{M}) + k \ln N + \ln(2\pi|\mathbf{V}|) \tag{3.22}$$

where $\chi^2_{\min}$ is the minimum χ^2 obtained when minimising the distance between the data and model $\mathcal{M}$, k is the number of degrees of freedom of $\mathcal{M}$ (three for a sine model and five for a Keplerian one) and $|\mathbf{V}|$ is the determinant of the correlation matrix. The orbit is said to be eccentric if $\mathrm{BIC}(\mathcal{M}_{\mathrm{ecc}}) \geqslant \mathrm{BIC}(\mathcal{M}_{\mathrm{circ}})$ where $\mathcal{M}_{\mathrm{ecc}}$ and $\mathcal{M}_{\mathrm{circ}}$ are respectively eccentric and circular models.

Though reasonable, these techniques can be improved. First, they both consider the alternative e is zero or non-zero, and do not allow to test if a given value of eccentricity is compatible with the data or not. Secondly, the analytical approximation of the eccentricity distribution is not always accurate. Indeed, Fig. 3.4 shows that the mean value of estimated eccentricity can be several times superior to the analytical approximation. Finally, as we saw in section 3.3.3, it does not seem fair to penalize equally all parameters, including the eccentricity in the model is far from adding a linear parameter. The Bayesian information criterion (3.22) gives equal weight to all parameters, only their number k appears. This approximation is valid in the limit of a large number of observations, condition which is not always fulfilled.

Our aim is to overcome as much as possible these limitations. It turns out that the procedure to construct confidence intervals outlined in Casella & Berger (2001), chapter 9, will allow us to test the hypothesis that the true eccentricity is equal to a certain value e for all e. While it is based on analytical computations, it does not suffer from the same drawbacks as the p-values test. The idea is to reject the hypothesis that eccentricity is equal to e if all models with eccentricity e have a likelihood lower than a fraction of the maximum likelihood. The following criterion is computed in Appendix F. We reject the hypothesis that the eccentricity has a certain value e with a confidence level α if

$$LR := \frac{\max\limits_{\boldsymbol{\theta}\in\Theta_e} f(\boldsymbol{y}|\boldsymbol{\theta})}{\max\limits_{\boldsymbol{\theta}\in\Theta} f(\boldsymbol{y}|\boldsymbol{\theta})} \leqslant \frac{1}{\sqrt{2\pi}^N |V|}\, \mathrm{e}^{-\frac{1}{2}\beta} \tag{3.23}$$

$$\beta = F^{-1}_{\chi^2_\nu}(1-\alpha) \tag{3.24}$$

$$\nu = 2 + 2S'^2\frac{e^2}{1+e^2} - \frac{\pi e}{1+e^2} L_{\frac{1}{2}}\left(-\frac{S'^2}{2}\right) L_{\frac{1}{2}}\left(-\frac{e^2 S'^2}{2}\right). \tag{3.25}$$

where Θ_e is the set of parameters that have all eccentricity e, $f(\boldsymbol{y}|\boldsymbol{\theta})$ is the likelihood, $F^{-1}_{\chi^2_\nu}$ is the inverse cumulative distribution function of a χ^2 law with ν degrees of freedom , $S' = (\sigma/K_t)\sqrt{2/N}$ and $L_{\frac{1}{2}}$ is the Laguerre polynomial of order 1/2. The quantity (3.23) is simply the ratio of the maximum likelihood obtained by restriction to the models with fixed eccentricity divided by the maximum likelihood on all models. The condition states that if all models that have eccentricity e have too low a likelihood, then e is rejected. The following equations give the value of that threshold, which is obtained by calculating the law followed

by the random variable LR under the hypothesis that the true eccentricity is e ($LR|(e_t = e)$). It is in fact easier to compute the law followed by the logarithm of LR, to obtain a χ^2 law whose degree depends on a definition of the signal to noise ratio S and on the eccentricity under study, but is always smaller than 2. Our computations, detailed in Appendix F, also make use of simplifying assumptions, but these are checked to give satisfactory results on simulated signals.

One of the problems of that expression is that it depends on the true value of the semi amplitude, K_t, which is unknown. There are two ways to circumvent this issue: either by assuming that $\nu = 2$ for all e, which is the maximum value ν can take, of K_t, or by approximating K_t by the semi amplitude of a circular orbit fitted at the period of the signal. The first option can be used to obtain conservative intervals to ensure that e is non zero. The second one will give a more realistic criterion to reject an eccentricity if no extra care is needed. Let us note that $\nu = 2$ is obtained for $e = 0$. This has a simple interpretation: the model can be approximated by a linear one in $k = e\cos\omega$ and $h = e\sin\omega$. When $e = 0$, both k and h are set to zero, which blocks two degrees of freedom. Denoting by $\boldsymbol{y}_e$ the model with fixed eccentricity e that has maximum likelihood and $\boldsymbol{y}^\star$ the model with maximum likelihood, all parameters free,

$$0.5\ln(LR) = \|\mathbf{W}(\boldsymbol{y} - \boldsymbol{y_e})\|^2 - \|\mathbf{W}(\boldsymbol{y} - \boldsymbol{y}^\star)\|^2. \tag{3.26}$$

behaves then as a χ^2 law with two degrees of freedom.

To offer a point of comparison of the confidence interval computed equation (3.25), we test two other metrics. First, we simply consider

$$F = \frac{N-p}{\nu}\frac{\|\mathbf{W}(\boldsymbol{y} - \boldsymbol{y_e})\|^2 - \|\mathbf{W}(\boldsymbol{y} - \boldsymbol{y}^\star)\|^2}{\|\mathbf{W}(\boldsymbol{y} - \boldsymbol{y}^\star)\|^2} \tag{3.27}$$

which is basically equation (3.26), normalized by $\|\mathbf{W}(\boldsymbol{y} - \boldsymbol{y}^\star)\|^2$ so that it depends less on the noise level assumption. The quantity (3.27) is assumed to follow a F distribution with ν and $N - p$ degrees of freedom.

Secondly, we generalize the test suggested by Lucy & Sweeney. Let us denote by $e^\star$ the estimate of eccentricity obtained by maximum likelihood when all parameters are free. For eccentricity e, we fit a Keplerian model that has an eccentricity fixed at e. We then compute the probability

$$\Pr\{|\widehat{e} - e| > |e^\star - e||e, \mathbf{V}, \widehat{e} \sim \text{Rice}(e, \eta^2)\} \tag{3.28}$$

that is the probability that an eccentricity estimated by maximum likelihood $\widehat{e}$ is at least as far from its assumed value e than the distance between e and the best fit actually observed, assuming the noise model is Gaussian with known covariance matrix $\mathbf{V}$. We also assume that e follows a Rice distribution as in section 3.3.3. A Rice distribution can be seen as the modulus of a vector with two independent Gaussian variables that have the same variance, $X \sim G(a, \eta^2)$ and $Y \sim G(b, \eta^2)$ where k and h are the means of these variables. To specify the distribution, we need therefore two scalars: the variance of both random variables η^2) and the modulus of the mean of these two variables, $r = \sqrt{a^2 + b^2}$. Here $a = k$ and $b = h$, so $r = e$. Then η^2 is the variance of the estimates of k or h, which under the hypotheses of section 3.3.3 have the same variance $\eta^2 = (\sigma^2_{\text{RV}}/K_e^2)(\pi/(N - p))$. Then the quantity (3.28)

can easily be evaluated by the cumulative distribution function of the Rice distribution, which is a Marcum Q-function.

Computing (3.26), (3.27) or (3.28) necessitates to compute the minimum distance between the observations and a model with fixed eccentricity. To do so, we exploit the fact that Keplerian models are partly linear, $\boldsymbol{y}(\boldsymbol{t},\boldsymbol{\theta}) = A\dot{\boldsymbol{X}}(P,e,\omega) + B\dot{\boldsymbol{Y}}(P,e,\omega) + C$ where $\dot{\boldsymbol{X}}$ and $\dot{\boldsymbol{Y}}$ are the components of the velocity on the orbital plane. For each couple e,ω, we can minimize $\|\boldsymbol{y} - \boldsymbol{y}(\boldsymbol{t},\boldsymbol{\theta})\|$ on A, B, C and P, which are respectively three linear parameters and one non-linear parameter. Such a problem is fast to solve with, for instance, a Levenberg-Marquardt algorithm (Levenberg 1944; Marquardt 1963). If the period is already known (which is supposed here), obtaining an array of χ^2 on a fine grid of e and ω (60 values each) takes only up to one minute. Let us finally note that the idea of restricting the global χ^2 minimization to a grid of non linear parameters is not new (Hartkopf et al. 1989; Lucy 2014). There are even further resemblances with Lucy & Sweeney (1971), however they use a degree of freedom $\nu = 1$ for all parameters, which is correct only if the model is linear in all the parameters or approximately linear in the vicinity of the best fit and unimodal.

The formula (3.24) and (3.25) have been derived with simplifying assumptions. To test and compare them to other options, we proceed as follows:

1. Generate a population of exoplanets according to a certain prior density on the orbital elements $p(K,e,P,M_0,\omega)$. The measurement times are taken from existing data sets. The noise generated according to a Gaussian density of covariance matrix V.

2. For each system, we compute the set of eccentricity that are not rejected, we check that the true eccentricity belongs to this set and compute the measure of its complement in $[0,1]$, that is the measure of the set of rejected eccentricities.

3. The results are summarized in two plots. First, the fraction of cases where the true eccentricity is not in the acceptable interval as a function of α. Second, the curve drawn when α goes from 0 to 1 by a point whose ordinate is measure of the complement of the set of acceptable eccentricity and whose abscissa is the fraction of cases where the true eccentricity is not in the acceptable interval.

Such tests were carried out with the following inputs: the measurement times are those of CoRoT-9 (Bonomo et al. 2017b). The angles ω and M_0 are chosen uniformly, e follows an uniform distribution. In Fig. 3.14 we plot the result of the experiment for a period fixed at 95 days the semi amplitude is fixed to $K = 3.5\sigma$ where σ is the RMS of the errors. These are the parameters of the detected Jupiter in the system. In Fig. 3.15, we let the period vary uniformly in $\log P$ and compute the same quantities.

The left plot of Fig. 3.14 and. 3.15 is labelled "ROC - like" curve as a reference to receiver-operator characteristic. These curves represent the fraction of false positives as a function of the false negatives for binary tests. In our continuous case, these are interpreted as follows. For a given rate of true eccentricity rejected (false negatives), the y axis gives the precision on the estimate. The more eccentricities are rejected, the more precise the estimate. The closer such a curve is to the upper left corner the better: regardless of the value of α, the fraction of true e rejected is zero (no false negatives) and almost all other eccentricities are rejected: the estimation is very precise.

Interestingly enough, the ROC curve (left) is very similar for all the metric considered with a slight advantage for the F-ratio and the likelihood ratio tests (formula (3.26) and (3.27)),

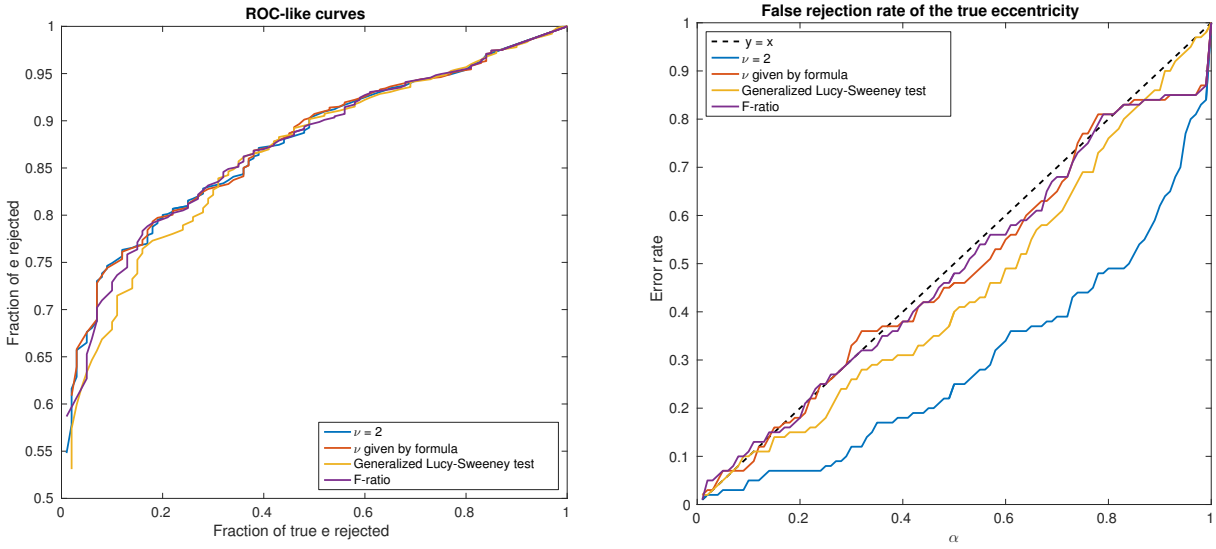

Figure 3.14: period fixed

which have a better precision (more eccentricities rejected) when the fraction of true e rejected is low. We now need to set the level of true e rejected. As expected, the curve obtained for $\nu = 2$ gives an overestimated error rate for a given α. For the three other tests, the correspondance seems appropriate. Overall, the F-ratio and the likelihood ratio tests seem to perform best. In Fig. 3.15, the plots are unchanged and lead to the same conclusions.

The main advantage of the method is that the graphs are very similar when the population set in input varies: it is weakly sensitive to the prior. Furthermore, we can set a parameter, α that allows us to directly control the confidence intervals meaning. Let us mention that we observed some peculiar behaviour of the estimates for some periods where the matrix of the linearized model is ill-conditioned, that we wish to investigate into more depth in future studies. In those cases, the hypotheses allowing to compute formula (3.25) are not verified and Bayesian analysis or more sophisticated formula would be required. On the other hand, it seems unlikely that someone would want to prove a non zero eccentricity of a planet particularly poorly sampled.

Minimizing the impact of model errors: normalizing the χ^2

When fitting orbits and noise models to radial velocity, one often encounters the situation where the reduced χ^2 of the data is greater than one. A common practice is to add a constant error σ^2 to the nominal ones, $(\sigma_k)_{k=1..N}$ in quadrature to obtain a reduced χ^2

$$\frac{1}{N-p}\chi^2 = \frac{1}{N-p}\sum_{k=1}^{N}\frac{(y_k - f(\boldsymbol{\theta}^\star, t_k))^2}{\sigma_k^2 + \sigma^2} = 1 \tag{3.29}$$

where $\boldsymbol{\theta}^\star$ are the parameters giving the best fit. We refer to that procedure as "normalizing the χ^2". The goal of this section is to show that this practice is efficient. More precisely, by doing so, the decisions made from the interval estimates are in general correct. We will show

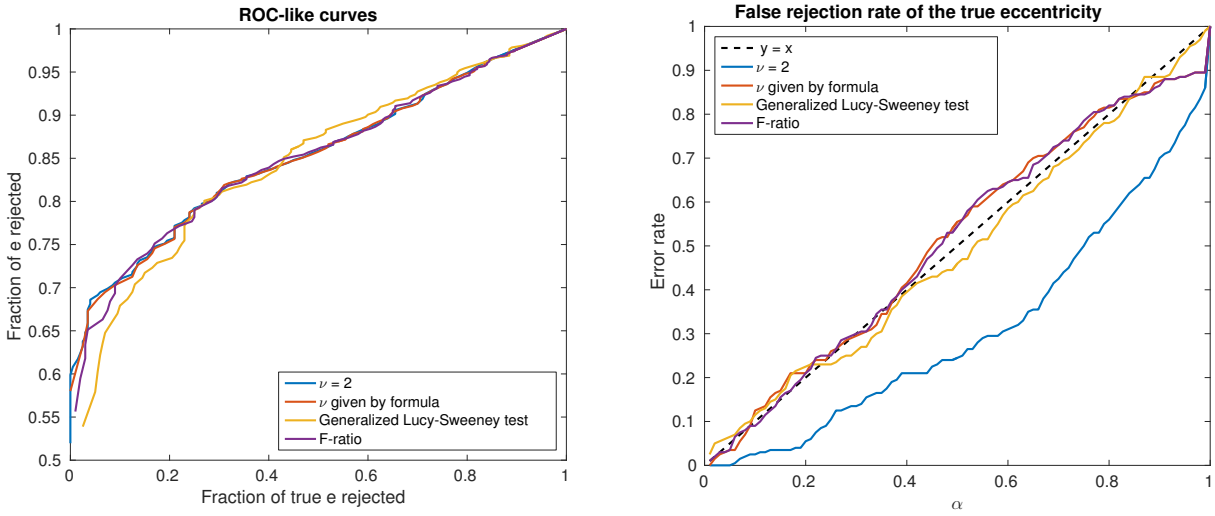

Figure 3.15: All parameters vary

different examples of spurious models where the normalization leads to statistical properties similar to when the model is perfectly known.

The first experiment we carry is to generate a circular orbit ($e = 0$) of semi amplitude 3 m.s^{-1}, period chosen uniformly in $\log P$ and phase also selected uniformly in $[0, 2\pi]$. The nominal noise is 1 m.s^{-1}. In addition, we generate several other signals: unseen white, red noises, outliers, unseen companion or mixes of them. In each case, the χ^2 is normalized. Then, for each realization, we find the best fit eccentricity $e^\star$ and compute numerically the probability $p_{e^\star}$ to have an estimated eccentricity fitted at least as high as $e^\star$, knowing the eccentricity is equal to 0. In summary, for each system generated, we obtain a p-value $p_{e^\star}$, and we can decide that if it is less than 0.05, the null eccentricity is rejected. Since it is a p-value test and the null hypothesis is true, the distributions of $p_{e^\star}$ should be uniform. Fig. 3.16 shows that if there is no correction, this is far from true, there will be up to 65% spurious rejection of $e = 0$. However, in all cases, when normalizing the χ^2, the p-values are distributed uniformly. Therefore the false alarm rate chosen is effectively observed.

Let us note however that normalizing the χ^2 does not improve the accuracy of the least square estimates. Fig. 3.17 shows that the distributions of estimated eccentricities with and without corrections are similar.

The fact that too high a reduced χ^2 leads to underestimated error bars is also true in the Bayesian setting. To illustrate it, we took the 74 measurement dates of HD69830 (Lovis et al. 2006) and generated a circular orbit of semi amplitude 2 m.s^{-1} and period 31.56 days. We then generated white noise with the measurements nominal uncertainties (whose average is 0.7 mm/s) plus a noise of 2 m.s^{-1}. After checking that the detection was made with high probability, the posterior densities of orbital elements were calculated with adjustment of the noise level in two ways: with a narrow Gaussian prior on the extra noise variance centred on 1 m.s^{-1} and with a free level initialized so that the normalized χ^2 equals one. In both cases we calculate the ratio of probability $\Pr\{e > 0.05\}/\Pr\{e < 0.05\}$. In the first case we obtained 61.21, and in the second only 10.62. Even in Bayesian analysis, when the noise is fitted, it must be ensured that the reduced χ^2 of the residuals is not unrealistic, regardless of

its origin.

Finally, let us briefly address the question: when does the χ^2 scaling leads to spurious error bars? We do not claim to list all possible cases, but will limit ourselves to one important one: when the noise can be fitted by the model. To see that, let us first consider the simplest case. Let us suppose we have observations $\boldsymbol{y} = \mathbf{A}\boldsymbol{x}_t + \boldsymbol{\epsilon}$ where $\boldsymbol{\epsilon}$ is an unknown vector and the columns of $\mathbf{A}$, denoted by $\boldsymbol{a}_k$, are orthonormal. Then the least square estimate of component k of $\boldsymbol{x}$ is $\hat{x_k} = \boldsymbol{a}_k^T\boldsymbol{y} = \boldsymbol{x}_{t,k} + \boldsymbol{a}_k^T\boldsymbol{\epsilon}$. A property of the Gaussian noise is that the law followed by $\boldsymbol{a}^T\boldsymbol{\epsilon}$ does not depend on $\boldsymbol{a}$. It is Gaussian of mean 0 and variance σ^2 if $\boldsymbol{\epsilon}$ is identically and independently distributed Gaussian variables. When scaling the error, we simply assume that the noise, spread roughly equally over the vectors of the model, has an additional variance σ'^2 and we expect the error on parameter x_k to behave Gaussianly with a variance $\sigma^2 + \sigma'^2$. On the other hand, if the unknown noise turns out to be very correlated with one of the $\boldsymbol{a}_k$, in the limit case where they are collinear, $\hat{x_k} = x_{t,k} + \|\boldsymbol{\epsilon}\|$. Also, in that case the residuals will be unchanged, since $\boldsymbol{\epsilon}$ has completely been absorbed in the fit. Between this limit case, and $\boldsymbol{\epsilon}$ being distributed uniformly in all directions there is a range of situation where the residuals might undergo very small change, therefore very small scaling, but the error on one of the parameters might high. Now if the columns of $\mathbf{A}$ are not orthonormal, the reasoning is the same: we can end up with some of the noise absorbed in one parameter and unseen in the residuals. The situation is in general even worse since the condition number of $\mathbf{A}$ will in general be higher than one. In that case, we would not know that in fact the noise is much higher than expected but was absorbed in the fit. Then even scaled, the errors are under-estimated and small eccentricities are spuriously rejected. There is indeed a situation where this might happen: when there is an inner orbital companion in a 2:1 mean motion resonance. As Anglada-Escudé et al. (2010) noted, it can be hard to disentangle from an eccentric orbit. To a lesser extent, all inner resonant planets might be a source of increased eccentricity since Keplerian signals are periodic, therefore admit a Fourier expansion whose coefficients might absorb signal at the multiples of the frequency of the planet.

In the non linear case, the dimension of the space the model explores might be greater than in the linear one, as demonstrated section 3.3.3. It means that the fitting process would absorb most of the noise energy and lead to spurious error bars. In case of doubt, it seems good practice to try to fit only noise with the model evaluated at the measurement times and initialized at the period of the planet candidate. If the squared norm of the residuals are systematically lower than what would be expected from a linear least square fit, that is $N - p$, it means that the parameters of a planet will be ill-constrained. This check is useful even when doing a Bayesian analysis.

Detecting model errors: residual analysis

We have seen that normalizing the reduced χ^2 improves the robustness of the conclusions. However, though it is an efficient method, it hides a deeper problem. Indeed, if the reduced χ^2 is far too high to be a likely realization of a Gaussian noise it means that even the best fit model is extremely unlikely. In other words, we selected the best option among models that are all very bad. A good model $\boldsymbol{y}^\star$ of the data ($\boldsymbol{y}$) should be such that obtaining $\boldsymbol{y}$ knowing $\boldsymbol{y}^\star$ is not completely unlikely. Adapting the error so that the χ^2 of the best fit residuals, $\chi^\star$ ensures that $\Pr\{\chi^2 \geqslant \chi^\star|\boldsymbol{y}^\star\}$ is above a certain rejection rate. But this adaptation, besides being *ad hoc*, guarantees only one property that a Gaussian residual should have. There are

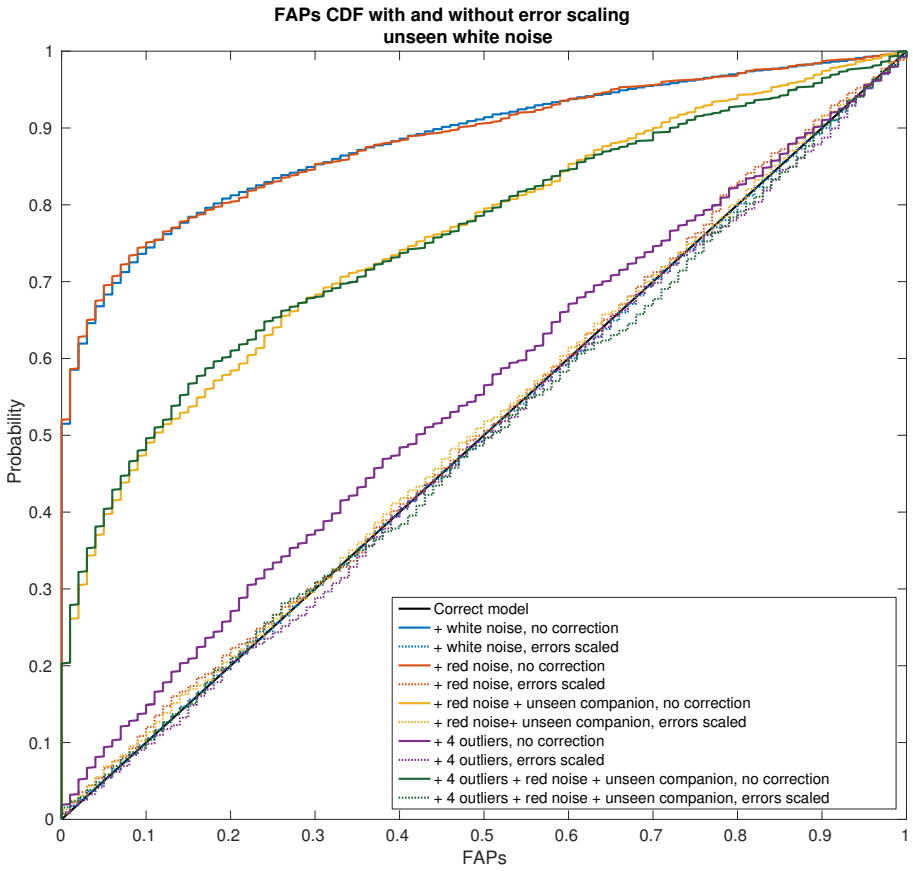

Figure 3.16: Cumulative distribution function of false alarm probabilities for various spurious models. CDF of the p-values without error scaling are represented in solid lines. CDF of p values with corrections are in dotted lines, and almost indistinguishable from the $x = y$ line, and means these are uniformly distributed, which is what is desired.

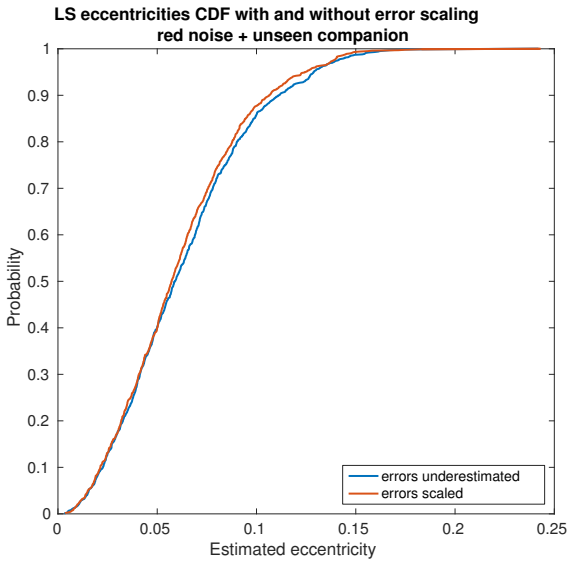

Figure 3.17: Cumulative distribution functions of false alarm probabilities for various spurious models.

many other tests that can help us guaranteeing that the residual is consistent with the model in an absolute sense. This problem is sometimes referred to as the goodness-of-fit problem, and is in general not trivial Lehmann & Romano (2005). The use of Likelihood ratio or Bayes factor is largely due to the difficulty to prove that a given model is absolutely consistent with the data, and allows to chose models that are significantly better than other ones. But even if one does not want to search for an absolute explanation of the signal, it is still interesting to know if there is something that is not understood in the signal, possibly for future researches.

In summary, our goal is to ensure that there is at least one acceptable model of which the data are a likely realization. There is a classical method to do it, the Kolmogorov-Smirnov test Kolmogorov (1933); Lehmann & Romano (2005), which consist in computing a test statistic which measures the probability that a data set is a realization of a parent distribution. However, this test would be correctly used if the hypothetical parent distribution is chosen independently of the fitting process. For the sake of speed we would like to test models that are close to the best candidates, which are chosen after looking at the data. Also this test measures the agreement between distributions in a binary manner and does not give hints on the origin of a possible discrepancy. Instead, we will reason as follow: if the set of model is appropriate to describe the data, then the residual of the best fit must verify certain properties. If they do not, then we will reject the hypothesis that there is one acceptable model that explains entirely the data set.

As in section 3.4.2, we will first examine the linear case and show that the results are still helpful in the non linear setting. Let us suppose we have a linear model $\boldsymbol{y} = \mathbf{A}\boldsymbol{\theta} + \boldsymbol{\epsilon}$ where $\mathbf{A}$ is a $N \times p$ matrix and $\boldsymbol{\epsilon}$ is a Gaussian noise of covariance matrix $\mathbf{V} =: \mathbf{W}^{-1}$. Let us denote by $\widehat{\boldsymbol{y}}$ the least square fit model, and suppose the model $(\mathbf{A}, \mathbf{V})$ is known. Then the weighted residual

$$\boldsymbol{r}_W := \mathbf{W}^{1/2}(\boldsymbol{y} - \widehat{\boldsymbol{y}}) \tag{3.30}$$

is a vector of N random variables that are approximately independent, Gaussian of null mean and variance one. To obtain a weighted residual that is a vector of independent Gaussian variables, let us define $\mathbf{Q}$, the matrix such that $\mathbf{J} = \mathbf{Q}^T(\mathbf{I}_N - \mathbf{W}^{1/2}\mathbf{A}^T(\mathbf{A}^T\mathbf{W}\mathbf{A})^{-1}\mathbf{A}^T\mathbf{W}^{1/2})\mathbf{Q}$ is diagonal (it exists). Then the re-weighted residual $\boldsymbol{r}'_{QW} = \mathbf{Q}^T\mathbf{W}^{1/2}(\boldsymbol{y} - \widehat{\boldsymbol{y}})$ has p null components. The $N - p$ others are Gaussian variables of mean 0 and variance 1. In what follows, we will denote by $\boldsymbol{r}_{QW}$ the vector made of the $N - p$ components of $\boldsymbol{r}'_{QW}$. These two results are proven in Appendix E.

In practice, $\mathbf{A}$ and $\mathbf{V}$ are unknown, and we choose models $\mathbf{A}'$ and $\mathbf{V}'$. The two above properties can be used to test if $(\mathbf{A}, \mathbf{V}) = (\mathbf{A}', \mathbf{V}')$ because if so, then the weighted residuals $\boldsymbol{r}_W$ and $\boldsymbol{r}_{QW}$ have a known law. There are numerous techniques to study correlations in time series sampled at equispaced times, but for uneven samples, to our knowledge there are no generic tools for such situations. We devised two tests adapted to our problem,

1. Plotting $d(t_i, t_j) = \boldsymbol{r}_W(t_i) - \boldsymbol{r}_W(t_j)$ as a function of $t_i - t_j$ for $t_i > t_j$. If $\boldsymbol{r}_W$ is indeed independent and Gaussian, $d(t_i, t_j)$ should not depend on the time interval.

2. Ordering $\boldsymbol{r}_{QW}$ by increasing values $(\tilde{\mathbf{r}}_{QW})$ and plotting the vector $\boldsymbol{z} = (k/(N - p))_{k=1..N-p}$ as a function of $\tilde{\boldsymbol{r}}_{QW}$. In other words, we compute an experimental cumulative distribution function (CDF). If our model is correct, then it should follow closely the CDF of a Gaussian variable of mean zero and variance one. One can potentially

test the hypothesis that $\mathbf{r}_{QW}$ is a realization of such a law with a Kolmogorov-Smirnov test or other metrics such as Anderson & Darling (1954), Shapiro & Wilk (1965), etc.

Let us now show how it can be used in practice. We take the 214 measurement times of Proxima b Anglada-Escudé et al. (2016). $\mathbf{A}$ is made of six columns as defined in Appendix C and fix $\boldsymbol{x}_t$. We then generate three series of a thousand realisation of $\boldsymbol{y} = \mathbf{A}\boldsymbol{x}_t + \boldsymbol{\epsilon}$. The covariance matrix of the noise has a kernel $\mathrm{e}^{|\Delta t|/\tau}$ where Δt is the duration between two samples. The three series are generated with a noise time-scale $\tau = 0$, 10 and 100 days. For each of the 3×1000 signals generated, we compute the least square fit with the correct matrix $\mathbf{A}$, but with a weight matrix $\mathbf{W}$ equal to identity, so our model will be entirely correct only in the first case. First, we pick randomly one realization among the 1000 available in each series, and perform the first test whose result is plotted Fig. 3.18. One clearly see a pattern: the higher the correlation, the smaller is the difference between residuals. Fig. 3.19 shows the 1000 empirical CDFs in the three cases. The deviation from a normal distribution grows with the correlation.

The plots 3.18 and 3.19 are useful indicators of remnant correlation in the residuals. However, they do not constitute metrics with known statistical properties. As said earlier, classical goodness-of-fit tests provide such metrics. Here we test the Anderson-Darling metric, which proceeds as follows: for each of the 3000 residuals obtained, we compute the Anderson-Darling statistic of the couple: residual/normal distribution, as defined in Anderson & Darling (1954). The higher the value of the statistic, the more unlikely it is that the residual follows an independent Gaussian distribution. We then obtain three series of a thousand numbers, whose CDF are plotted Fig. 3.20. The blue curve corresponds to a correct model. If we set as a condition that the null hypothesis (the model is correct) is rejected in 5% of the cases knowing it is true, it means that we will reject approximately 100 - 67 = 33% of the realizations knowing that the model with $\tau = 100$ days is correct, which is not a very satisfactory rate. The visual inspection, though less quantifiable, seems more accurate.

Study of correlations in RV residuals has already been undertaken for instance by Baluev (2011, 2013) with a smoothed residual periodogram (Baluev 2009). We leave the systematic comparison and testing of statistics for future work.

Reducing bias

The method presented above seem to be useful tools for a sound scientific analysis, which we believe is based on hypothesis rejection. However, it is also possible to use our understanding of the bias to aim at reducing it. Section 3.3 showed that lower eccentricities models occupy a smaller volume of the phase space than models with higher eccentricities, especially those with $e \gtrsim 0.8$. In some sense we would like to penalize the higher eccentricities as a function of the volume occupied by the corresponding models.

This situation is very similar to model selection, where extra credit is given to simpler models. Viewing the best estimate of a single parameter as model selection should help us reducing the bias.

3.4.3 Populations

A way to correct a distribution is to consider its constitutive systems individually. For each eccentricity estimate, the probability to have an eccentricity greater or equal to the measured

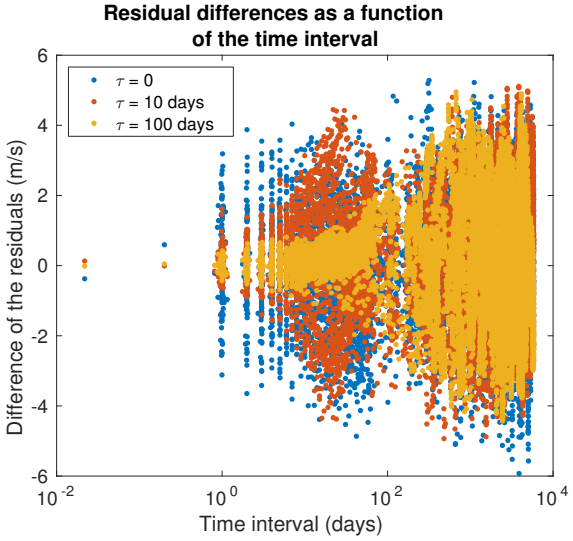

Figure 3.18: Difference between the residuals at two different time as a function of the time interval between them in three cases: when the noise has a time-scale of 0, 10 and 100 days.

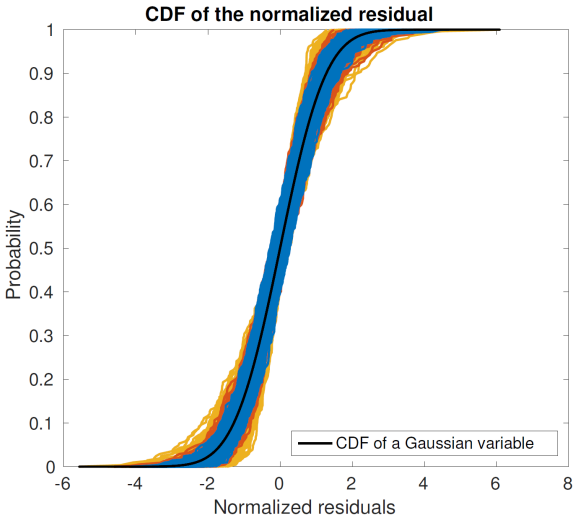

Figure 3.19: A thousand realizations of the cumulative distribution functions of the normalized residual in three cases: when the noise has a time-scale of 0, 10 and 100 days.

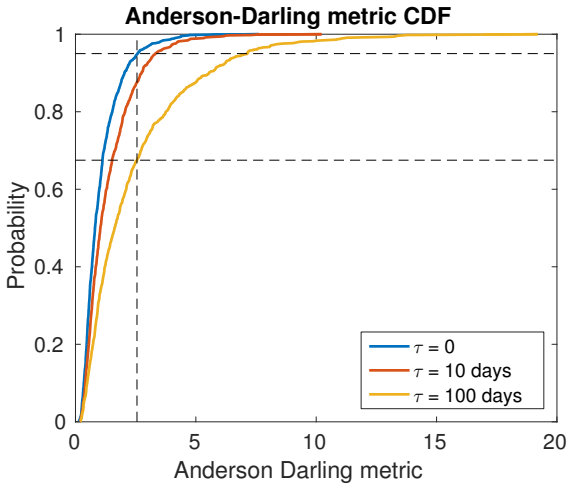

Figure 3.20: Anderson-Darling test

one, knowing the true eccentricity is zero (Lucy & Sweeney 1971; Zakamska et al. 2011) (this probability is often called a p-value). If this probability is below a certain threshold, the eccentricity is considered to be zero. This procedure implies that all correction should concern zero eccentricities, and its accuracy depends on the true proportion of zero eccentricities, which is unknown. Figure 3.11 shows the distribution of true eccentricity as a function of the estimated one for a flat true distribution. Paradoxically enough, when the estimated eccentricity is between 0 and ≈ 0.2, one could think that on average the true eccentricity was lower because of the bias but it is not the case. On average, the true eccentricity is greater

than the one measured. Also, by having such a particular treatment of zero eccentricity the risk of "seeing what we expect to see" is greater. Indeed, zero eccentricity is privileged because orbits are expected to be circularised by tidal forces. One cannot in turn make strong claims about the dynamics from a corrected distribution.

Two more sophisticated correction methods were suggested by Hogg et al. (2010). They both rely on considering the prior distribution itself as the unknown, and see observed distributions as the true one convolved with uncertainties. More precisely, the two options are the following. First, one can consider that the observed eccentricity distribution, $p(e)$ is a realization of the true one $p_t(e)$ convolved with the uncertainties, that is each observed eccentricity $\widehat{e}$ follows the distribution p,

$$\widehat{e} \sim p(\widehat{e}) = \int_0^1 p(\widehat{e}|e)p_t(e)de. \tag{3.31}$$

In practice, one would take a sequence of increasing eccentricities $(e_j)_{j=1..m}$ and consider

$$\Pr\{\widehat{e} \in [e_j, e_{j+1}]\} = \sum_{k=0}^{m} \Pr\{\widehat{e} \in [e_j, e_{j+1}]|e_t \in [e_k, e_{k+1}]\}\Pr\{e_t \in [e_k, e_{k+1}]\}\ . \tag{3.32}$$

Now if we denote by $\widehat{p}_j := \Pr\{\widehat{e} \in [e_j, e_{j+1}]\}$, $P_{jk} = \Pr\{\widehat{e} \in [e_j, e_{j+1}]\,|e_t \in [e_k, e_{k+1}]\}$ and $p_{t,k} := \Pr\{e_t \in [e_k, e_{k+1}]\}$, we have

$$\begin{bmatrix} \widehat{p}_0 \\ \widehat{p}_1 \\ \\ \vdots \\ \widehat{p}_m \end{bmatrix} = \begin{bmatrix} P_{11} & \dots & P_{1m} \\ \vdots & & \\ & & \\ P_{m1} & & P_{mm} \end{bmatrix} \begin{bmatrix} p_{t,0} \\ p_{t,1} \\ \\ \vdots \\ p_{t,m} \end{bmatrix} \tag{3.33}$$

This can be written in a more compact form, $\widehat{\boldsymbol{p}} = \mathbf{P}\boldsymbol{p}_t$. In the real setting, we would not observe $\widehat{\boldsymbol{p}} = (\widehat{p}_j)_{j=0..m}$ but some estimates of eccentricity from radial velocity measurements. If we assume that each estimate of eccentricity is an independent realization of the random variable $\widehat{e}$, then the number of eccentricities that fall in the interval $[e_j, e_{j+1}]$, denoted by $(n_j)_{j=1..m}$, follows a multinomial law of parameters $\widehat{p}_j$. One easily shows that the maximum likelihood estimate of $\widehat{p}_j$ is $\widehat{p}_j^{\mathrm{ML}} = n_j/n$ where $n = n_1 + ... + n_m$. Since the maximum likelihood does not depend on the parametrization, any vector $\boldsymbol{p}$ satisfying $\mathbf{P}\boldsymbol{p} = \widehat{\boldsymbol{p}}_j^{\mathrm{ML}}$ is a maximum likelihood estimate of $\boldsymbol{p}_t$. In general it will not be possible to invert $\mathbf{P}$. The next section defines constraints on observation strategies to ensure that $\mathbf{P}$ is somehow "well behaved".

The method however ignores that different observations have different uncertainties. The second option is to consider that the true distribution is parametrized by a vector of parameters $\boldsymbol{\alpha}$ and to look for its value giving the posterior likelihood of the distribution $\mathcal{L}$.

$$\mathcal{L}(\boldsymbol{\alpha}) = \prod_{k=1}^{M} \mathcal{L}_k(\boldsymbol{\alpha}) \tag{3.34}$$

$$\mathcal{L}_k(\boldsymbol{\alpha}) = \int p(\boldsymbol{y_k}|\boldsymbol{\theta_k})p(\boldsymbol{\theta_k}|\boldsymbol{\alpha})d\boldsymbol{\theta_k} \tag{3.35}$$

$$p(\boldsymbol{\theta_k}|\boldsymbol{\alpha}) = \frac{p_{\boldsymbol{\alpha}}(e)p(\tilde{\boldsymbol{\theta}}_{\boldsymbol{k}})}{p_0(e)} \tag{3.36}$$

where $\boldsymbol{\theta_k} = [e, \tilde{\boldsymbol{\theta}}_{\boldsymbol{k}}]$ are the set of parameters of observation k (these are equation 7 in Hogg et al. (2010)). Since the method relies on writing the Bayes formula several times (here twice) on different levels (population, individual system), this one is called hierarchical Bayes modelling. This has been applied more recently to the mass radius relationship of Kepler data (Wolfgang et al. 2016). The first equation holds assuming that the observations are independent from each others, and the third one subtends that the prior on eccentricity is independent from the prior of the other parameters. The latter assumption is debatable, since eccentricity is unlikely to be independent from the period of the orbit. Note that when computing the posterior of the distribution, credible intervals can be computed like in section 3.4.2. The maximum *a posteriori* distribution is here a point estimate, and is not the most interesting quantity from the scientific point of view.

The two methods mentioned above do not take into account the selection effect on the detection, and we saw in section 3.3.5 that the measured distribution depends non negligibly on the detection efficiency. The two frameworks presented above can be adapted easily by adding to the unknown distribution a parameter, the probability that there is no planet, and to include non detection as an outcome. Let us see how that would change the analysis based on equations (3.31) and (3.32).

If we include a possibility for non detection and that there is no planet, equation (3.32) becomes

$$p(\widehat{e}) = \int_0^1 p(\widehat{e}|e)p_t(e)de + \Pr\{\widehat{e}|\text{NP}\}\Pr\{\text{NP}\} \tag{3.37}$$

$$\Pr\{\text{ND}\} = \int_0^1 p(\text{ND}|e)p_t(e)de + \Pr\{\text{ND}|\text{NP}\}\Pr\{\text{NP}\} \tag{3.38}$$

where ND and NP stand for "No Detection" and "No Planet". These terms account for the situation where nothing is detected, while a planet might be present, and when no planetary signal is in the data but still we might have a false detection. We now approximate equations (3.37) and (3.38) by a discrete version. Then equation (3.33) becomes

$$\Pr\{\widehat{e} \in [e_j, e_{j+1}]\} = \sum_{k=0}^{m} \Pr\{\widehat{e} \in [e_j, e_{j+1}]\,|e_t \in [e_k, e_{k+1}]\}\Pr\{e_t \in [e_k, e_{k+1}]\} + \Pr\{\widehat{e}|\text{NP}\}\Pr\{\text{NP}\} \tag{3.39}$$

Now if we denote by $\widehat{p}_j := \Pr\{\widehat{e} \in [e_j, e_{j+1}]\}$, $P_{jk} = \Pr\{\widehat{e} \in [e_j, e_{j+1}]\,|e_t \in [e_k, e_{k+1}]\}$ and $p_{t,k} := \Pr\{e_t \in [e_k, e_{k+1}]\}$, we have

$$\begin{bmatrix} \Pr\{\text{ND}\} \\ \widehat{p}_0 \\ \widehat{p}_1 \\ \\ \vdots \\ \widehat{p}_m \end{bmatrix} = \left[\begin{array}{c|ccc} \Pr\{\text{ND}|\text{NP}\} & \Pr\{\text{ND}|e_0\} & \dots & \Pr\{\text{ND}|e_m\} \\ \hline \Pr\{\widehat{e_0}|\text{NP}\} & & & \\ \vdots & & & \\ & & (P_{kl})_{k,l=0..m} & \\ \Pr\{\widehat{e}_m|\text{NP}\} & & & \end{array}\right] \begin{bmatrix} \Pr\{\text{NP}\} \\ p_{t,0} \\ p_{t,1} \\ \\ \vdots \\ p_{t,m} \end{bmatrix} \tag{3.40}$$

The same equation can be written more compactly in a matrix form,

$$\widehat{\boldsymbol{p}} = \mathbf{L}\boldsymbol{p}_t \tag{3.41}$$

Note that if we were to represent $\mathbf{L}$ in color code, it would look like the transpose of Fig. 3.13.b. In the limit of an infinite signal-to-noise ratio, the matrix tends to the identity. On the

contrary, when the signal to noise ratio tends to zero, the probability of non detection reaches its maximum and all the lines of rows of $\mathbf{P}$ tend to the distribution of eccentricity estimate when the input is pure noise.

3.4.4 Observation strategies

The previous section is concerned with estimating the true distribution from existing measurement. One can change the viewpoint and ask the following question: what are the requirements on the observations to ensure that the distribution of an orbital element is close to the true one? In that case, the unknown is the array of measurement times, or some of its characteristics (number of measurements, timespan...), and we allow the measurement strategy to depend on the star under study. We will now specialize to eccentricity distributions, but the reasoning can be applied to any parameter.

Our problem is to retrieve the true distribution of eccentricity, denoted by $p_t(e)$, of a certain type of planets (that have a certain range of mass, semi-major axis or type of the host star etc...). Let us consider a sequence of eccentricities $(e_i)_{i=1..m}$ and suppose that n_i eccentricities were measured in the bin $[e_i, e_{i+1}]$ and n_0 observations lead to no detections. If we denote by $\widehat{p}_i$ the probability $\Pr\{\widehat{e} \in [e_i, e_{i+1}]\}$, $\widehat{p}_0$ the probability of detecting nothing, then the vector $\boldsymbol{n} = (n_i)_{i=0..m}$ follows the multinomial law

$$\Pr\{n_1, ...n_m\} = \frac{n!}{n_0!..n_m!}\widehat{p_0}^{n_0}...\widehat{p_m}^{n_m}, \tag{3.42}$$

n being defined as $n = n_0 + ...n_m$. As a consequence, from equation (3.41), the likelihood function of the true eccentricity parameters is

$$L(\boldsymbol{n}|\boldsymbol{p}_t) = \Pr\{n_1, ...n_k\} = \frac{n!}{n_0!..n_m!}(\boldsymbol{L}_0\boldsymbol{p}_t)^{n_0}...(\boldsymbol{L}_m\boldsymbol{p}_t)^{n_m} \tag{3.43}$$

where $\boldsymbol{L}_i$ designates the i-th row of matrix $\mathbf{L}$. Assuming that $p(\widehat{e}|e_t)$ is a Rice distribution of parameters e_t and $1/S$, as in section C, then we can compute the matrix $\mathbf{L}$. Then, by performing a principal component analysis (PCA) (or equivalently a singular value decomposition) of $\mathbf{L}$ we find the components of the eccentricity distribution that can be resolved. Fig. 3.21 and 3.23 show the matrix constructed as (3.40) in color code for respectively $S = 6.6$ and $S = 66$. In the second case the matrix is very close to the identity, so that all the components of e_t can be resolved.

The PCA consists in writing $\mathbf{L} = \mathbf{USV}'$ where $\mathbf{U}$ and $\mathbf{V}$ are orthonormal matrices and $\mathbf{S}$ is a diagonal matrix. One can think of the columns of $\mathbf{V}$ as a basis of vectors whose linear combinations will be used to represent the true distribution. However, it is in general impossible to constrain all the components of the true distribution expressed in basis $\mathbf{V}$. As the signal to noise ratio decreases, the number of non zero elements of $\mathbf{S}$ decreases. If d elements of $\mathbf{S}$ are non-zero, ordered in decreasing order along the diagonal, it means that the components of the true distribution along the last $m - d$ columns of $\mathbf{V}$ cannot be resolved. Interestingly enough, the columns of $\mathbf{V}$ are behaving like frequencies. Fig. 3.22 and 3.24 represent the behaviour of the five first columns of $\mathbf{V}$ and the corresponding elements of $\mathbf{S}$. As the order of the principal component increases, the resolution on the true eccentricity becomes finer. In conclusion, by constructing the matrix $\mathbf{L}$ and performing the PCA, the signal to noise ratio and detection power and false alarms might be tuned to obtain a certain

resolution on the eccentricity distribution. The SNR is $S = \sqrt{2/(N-p+1)}\sigma/K$, where N is the number of observations, σ the RMS error, K is the semi-amplitude. One might either pick a minimum semi amplitude and design the whole strategy accordingly, or possibly stop if the planet is found to have a K large enough in the course of the observations.

The other parameter to be tuned is the number of systems to observe. This can be done approximately by first rewriting the likelihood as a function of the coefficient in the basis generated by the d first columns of $\mathbf{V}'$. Then one computes the likelihood Hessian, whose inverse is asymptotically equal to the covariance matrix of the maximum likelihood estimate. This Hessian can then be evaluated at different points.

3.5 Application to real cases

3.5.1 Outline

Section 3.3 was devoted to understanding the origin of the detection and estimation biases on the orbital eccentricity. We found that in some sense the geometry of the model is responsible for the estimation bias, the fact that some eccentricities go to one when fitted, and the difficulty of detecting very eccentric planets. We also found that modelling errors not only lead to underestimated uncertainties, but also increased bias in eccentricity. This bias on individual systems translates to population estimates, which are worsen by the bias against detection of high eccentricities. In section 3.4, we advocated the use of interval estimations rather than point estimate, since they are closer to the process of rejecting hypotheses. We showed that adding an error so that the reduced χ^2 is one is an efficient technique to correctly estimate uncertainties in both Bayesian and frequentist approaches. We also suggested methods to check the consistency of the model based on the analysis of residuals. In this section, we will use the tools developed to analyse data, focusing on single system analysis.

3.5.2 Example: analysis of CoRoT-9 RV data

Past analyses

The photometry data from the space mission CoRoT allowed to discover a transiting warm Jupiter of period 95.3 days, $a = 0.41$ AU and mass 0.84 ± 0.07 M_{jup} (Deeg et al. 2010). The data was re-analysed by Southworth (2011) who found similar values and error bars. In Bonomo et al. (2017b), the analysis of the system is done with two more transits measurements (18 June 2010, Spitzer, and 4 July 2011 simultaneously with CoRoT and Spitzer) and 28 HARPS measurements spanning from September 2008 to August 2013. One of the goals of the paper was to test whether a zero eccentricity is compatible with the data. If not, it means the planet might have undergone secular chaos or planet-planet scattering.

To do so, they analyse jointly the transit and radial velocity data sets. Orbital elements are sampled from their posterior distribution with a differential evolution Monte-Carlo Markov chain algorithm (DE-MCMC Eastman et al. 2013). Then the Bayes factor as defined equation (3.18) is computed and they find 124.8, which is a strong support for the eccentric scenario.

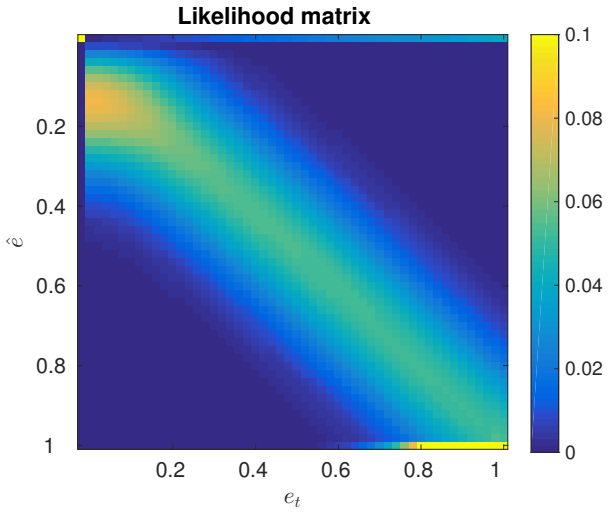

Figure 3.21: Representation of the discretized conditional density $p(\hat{e}|e_t) \sim \text{Rice}(e_t, 1/S)$ where S is the signal-to-noise ratio, $S = 6.6$.

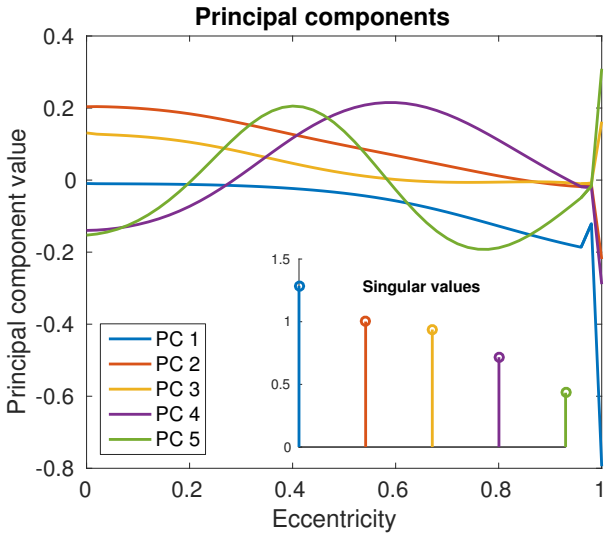

Figure 3.22: Principal components of the matrix $\mathbf{L}$, which are vectors — whose linear combination will represent the estimate of the density of eccentricity — and their associated singular values, $S = 6.6$.

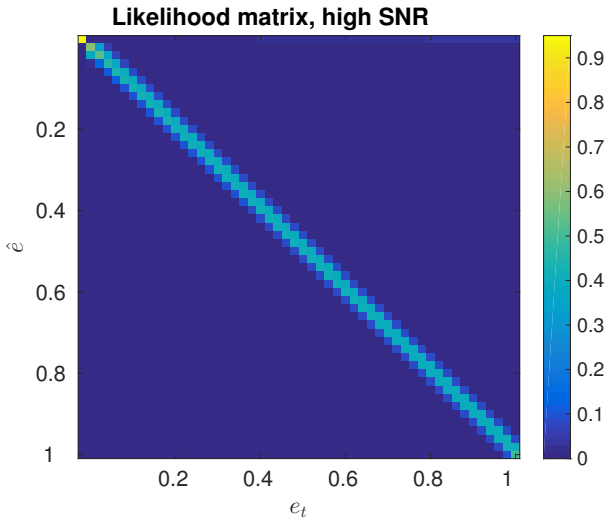

Figure 3.23: Same quantities as above but ten times higher signal-to-noise ratio, $S = 66$.

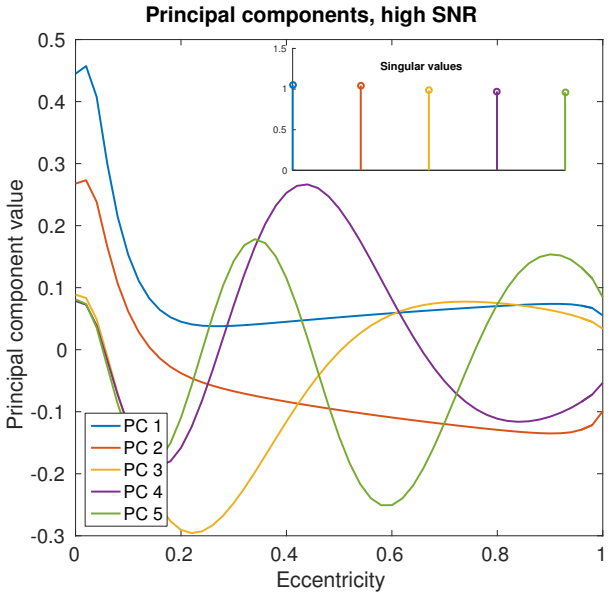

Figure 3.24: Same quantities as above but ten times higher signal-to-noise ratio, $S = 66$.

Re-analysing RV data

We first use the frequentist method 3.4.2 and compute for each eccentricity the likelihood ratio as defined in Eq. (3.23) and the associated probability. Let us recall that this consists in computing the sum of squared residuals of a model that has a fixed eccentricity e and the best fit model. One can then compute the probability to obtain a difference at least as great under the hypothesis that e is the true eccentricity. The lower is that probability, the more unlikely is the hypothesis $e_t = e$. We fit a Keplerian model plus a constant and obtain the FAP per eccentricity shown in Fig. 3.25 (left). If the FAP is lower than 0.05 then the corresponding eccentricity is rejected. This threshold means that the likelihood ratio would be as extreme as it is in 5% of the cases, assuming the correct eccentricity is the one under study. We find a minimum non rejected eccentricity of 0.08, which is consistent with the findings of Bonomo et al. (2017b).

To check this result, we plot the residuals of the fit $\boldsymbol{r}_W$ as defined in eq. (3.30) of section 3.4.2 as a function of the time interval between subsequent ones, that is $r_{W,i} - r_{W,j}$ as a function of $t_i - t_j$ for all $i > j$. We obtain Fig. 3.25 (right).

Now to finish, let us consider in that case what scenario could have lead to a spurious claim. Since the residuals of the best fit have a reduced χ^2 lower than one, what could have happened is that the model is in fact capable of fitting anything, and the errors are under estimated. As advocated section 3.4.2 we compute the average squared sum of residuals when fitting the model to noise only, and find 23.3, which is close to what we would expect from a linear fit, $N - p = 28 - 6 = 22$. The other possibility is that there was another source of noise η whose components are such that $y_t + \eta$ is still an acceptable model. Let us consider the possibility of an inner resonant planet in 2:1 mean motion resonance. When adjusting a two circular planet + constant model, considering nominal uncertainties, the χ^2 of the residuals is 16.05 versus 19.88 for a Keplerian model. It basically means that the two scenarios cannot be distinguished with radial velocity alone and without physical considerations. Some of the photometric data used in Bonomo et al. (2017a) for the joint analysis of photometry and radial velocity were not publicly available, we did not carry out the same analysis as Bonomo et al. (2017a) for the resonant model.

3.6 Discussion

Our initial goal was to give a consistent picture of the behaviour of eccentricity estimation. We give a justification for the fact that eccentricity estimates follow a Rice distribution, which was noted by Shen & Turner (2008) and the analytical approximation in the case $e_t = 0$ obtained by Lucy & Sweeney (1971). The formula was adapted to be more realistic and take into account the number of parameters fitted. We also show that the validity of the formula depends on the condition number of the Fisher information matrix. The notion of condition number also explains the behaviour noted by Zakamska et al. (2011): bias increases with the period and as the phase coverage becomes poorer. We see that the analytic Lucy-Sweeney test might be misleading when the Fisher information is poorly conditioned. This enjoins to use simulated systems to compute e-values, that are $\Pr\{e \geqslant \widehat{e} | e_t = 0\}$, where $\widehat{e}$ is the best fit eccentricity.

We show that in the case $e \gtrsim 0.8$ region, eccentricity does not parametrize the models uniformly, which explains the downward bias on high eccentricity, along with the facts that

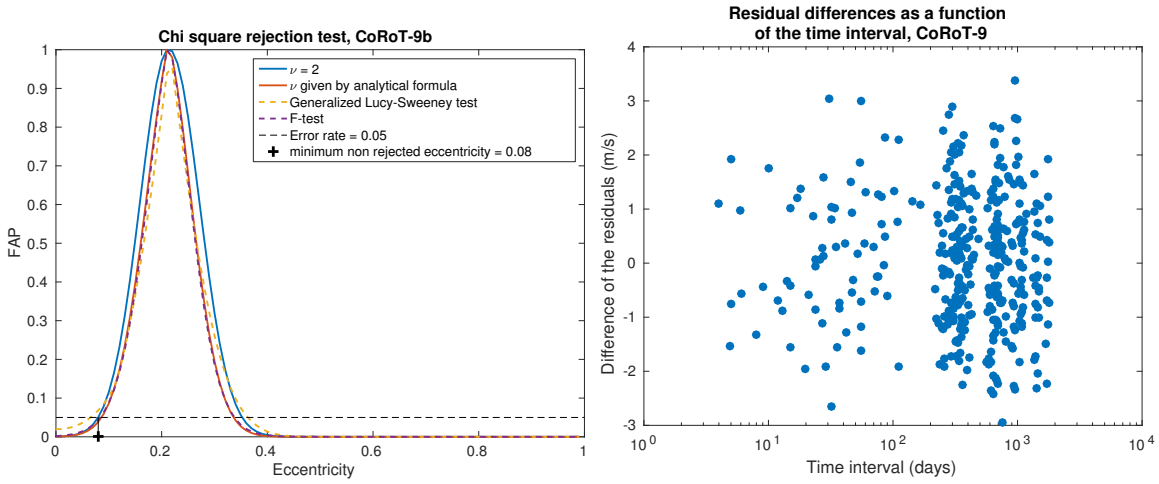

Figure 3.25: False alarm probabilities of eccentricity based on CoRoT-9 HARPS data (left). Difference between residuals as a function of the time interval between them.

some estimates are stuck at one because of the model boundary. The complicated shape of the model in this high eccentricity region accounts for the numerous local minima and the fact that high eccentricity models might spuriously be compatible with the data. The geometry of this region also explains the detection bias on high eccentricities. Since high eccentricity models are "far" from circular ones, they might be unseen on periodograms and due to the complicated shape, there might be local minima that are not accessible when initializing a Keplerian fit at $e = 0$.

This analysis leads us to formulate two propositions, verified by simulations. First, when the model is incorrect, the bias on eccentricity is higher on average. When the model is too simplistic, since there are "more" high eccentric models, the chance that the best fit appears on a high eccentricity increases. When the model is overly complicated, the condition number of the Fisher information matrix increases, so does the uncertainty, which is proportional to the bias at low eccentricity. Secondly, we show that one should not correct the distributions by setting eccentricities to zero whenever it is an option compatible with the data set, since it biases the estimate of the distribution and prevents from spotting unexpected effects.

The fact that eccentricity is biased upwards is not a problem in itself, since point estimates do not really have a scientific meaning. We are rather searching for all the models that are compatible with the data, that are confidence or credible region estimates. The fact that eccentricities are overestimated might become a problem when using such estimates if error bars are underestimated, since the eccentricities will over-confidently be locked at high values.

We then consider two options: Bayes ratio or credible intervals based on posterior distributions and confidence intervals based on a maximum likelihood ratio. In the Bayesian case, the inference depends on the prior, therefore we advise to test several of them whenever possible. We also propose a formula to re-calculate the posterior distribution of eccentricity when the prior is enhanced on a certain region. The distribution of the likelihood ratio is approximated analytically and we generalize the Lucy-Sweeney test. This is shown to behave as expected on numerical simulations.

Since our conclusion depends on the validity of the model, it is desirable to have a manner

to check that a given model is reasonable. We show that in many cases, forcing the model to have a reduced χ^2 equal to one is surprisingly robust. We discuss in which case this correction is not sufficient. To push the analysis further one might check that the residuals are compatible with the assumed model. Some solutions are considered, but we stress that there are many possibilities of goodness-of-fit tests yet to be explored.

We wish to emphasize this latter point. Regardless of the sophistication of the analysis, it is important to check that the residuals are likely. For instance, in Lovis et al. (2011b), after fitting six planets, a final reduced χ^2 of 1.7 is found. This might not seem too far from one, but in fact for 190 measurement points and 31 parameters fitted, the probability of obtaining a reduced χ^2 of 1.7 or above is less than 10^{-16}.

Chapter 4

Conclusion

Do what you can, with what you have, where you are.
Theodore Roosevelt

4.1 Condition number

We first tackled the problem of spurious peaks of the periodogram, then underwent to understand the eccentricity bias and what an honest account of uncertainties would be. The reader might have noticed that the notion of condition number, and more generally of the eigenvalues of the Fisher information matrix, kept appearing (see definition 3, page 69), be it for the stability of the basis pursuit algorithm (theorem 4), the amplitude of the bias on eccentricity in section 3.3.3 or the resolution on the eccentricity distribution for a given observation strategy (section 3.4.4). Bringing out this notion was not our initial purpose but we found — like many other data analysts (e.g. Demoment et al. 2001) — that it is an extremely powerful notion to quantify the stability of an inverse problem. The eigenvalues of the Fisher information matrix indicate which directions in the parameter space are well constrained and which ones are ill-constrained. This ubiquity motivated us in particular to try to use the restricted isometry constant (definition 6, page 72) as a generalization of the spectral window section 2.3.

Interestingly enough, in the radial velocity literature, analytical results on the false alarm probability of a detection have also been obtained via computations of the metric averaged in some complicated sense (Baluev 2015b). The metric describes the behaviour of the model in the vicinity of certain points. We believe that exploiting the "shape" of the model is an interesting trail for future theoretical work.

4.2 Method for data analysis

In this thesis we have tackled the problems of detection of exoplanets, the estimation of their orbital parameters, and of attaching uncertainties to those. These problems amount to select a family of hypotheses that are compatible with the data. Based on the theoretical considerations of section 1.3 and on practical experience with the data, we wish to propose a general method for data analysis that could be termed "closed-loop inference". This term is

chosen in reference to automated systems that update the control as a function of the error between the command and the measured state of the system. Closed-loop control is opposed to open-loop, where the command is sent to the system only once and no further adjustments are made, an instance of which are ballistic missiles. We do not claim originality here, many aspects of the following analysis are already implemented in existing techniques. Our goal is simply to clarify the guidelines for future work.

Our problem is to make a claim based on a certain amount of information. Here information is to be understood in a wide sense. For the discovery of an exoplanet, the type of the telescope, the sequence of decisions that took place to observe a certain star and the time series of observations are considered as part of the information we possess. Some of the aspects are hard to model statistically, but these should not be ignored.

The data itself (here the time series of observation) is supposed to have been generated by a certain random variable (or plausibility distribution, in a more general sense) and we wish to determine which one. We first define a family of hypotheses, where a hypothesis is either one or a collection of random variables. For instance, a hypothesis can be "3 planets with orbital parameters $\boldsymbol{\theta}$ and a Gaussian independent noise with standard deviation σ_k at measurement time t_k". The hypothesis "there are three planets and a Gaussian noise" is the set of all configurations with three planets plus any Gaussian noise, correlated or not.

We now have a family of hypothesis. To claim that a subset of them is a correct account of the data we must somehow prove existence and uniqueness. The uniqueness means that the probability of the subset of hypotheses is much higher than other alternatives. Existence means that the family of hypothesis can account for the data, there is a reasonable chance that the data has been generated from this family. It might mean that there is at least one hypothesis that correctly reproduces the data, or that there are so many hypotheses that even though none of them makes the data likely, the combined chance of occurrence of the data from this family is not "too low" in a sense to be made precise. When the alternatives are globally unsatisfying, it means that the set of alternatives has to be extended.

The point we would like to stress is that most of the ways to check existence of the solution are under exploited. When using a maximum likelihood estimator, posterior or not, posterior distributions or normalized confidence intervals, we are choosing among a set of alternatives that are potentially very bad. Running the analysis "open-loop", such situations are unseen. We stress that whenever possible, the models should not only be compared to each other. It must be checked that they are compatible with the data in an absolute sense.

A convenient way to do so is to check the residuals of the fit. Indeed, assuming that the true data has been generated by one of the hypothesis, one can compute the distribution followed by the residuals of the maximum likelihood. For instance, when the model $\boldsymbol{y} = \boldsymbol{f}(\boldsymbol{\theta}) + \boldsymbol{\epsilon}$ is linear and the noise ϵ is Gaussian, $\boldsymbol{y} = \mathbf{A}\boldsymbol{\theta} + \boldsymbol{\epsilon}$, appropriately scaled, the residuals of the least square fit must be distributed independently with Gaussian profiles. In the radial velocity literature, the residuals are often checked visually, but since there are many sources of correlated signal with unknown structure, studying the residuals (distribution of amplitudes, correlations...) in a systematic way could prove useful. We suggested some steps in that direction chapter 3, but many options are left to explore.

Ultimately, the usefulness of the methods presented in this thesis can only be evaluated in practical situations. We look forward to applying our methods to practical cases.

Bibliography

Affer, L., Micela, G., Damasso, M., et al. 2016, A&A, 593, A117

Aigrain, S., Gibson, N., Roberts, S., et al. 2011, in AAS Division for Extreme Solar Systems Abstracts, Vol. 2, AAS Division for Extreme Solar Systems Abstracts, 11.05

Aigrain, S., Pont, F., & Zucker, S. 2012, MNRAS, 419, 3147

Akaike, H. 1974, IEEE Transactions on Automatic Control, 19, 716

Amelunxen, D., Lotz, M., McCoy, M. B., & Tropp, J. A. 2013, ArXiv e-prints

Anderson, T. W. & Darling, D. A. 1954, Journal of the American Statistical Association, 49, 765

Anglada-Escudé, G., Amado, P. J., Barnes, J., et al. 2016, Nature, 536, 437

Anglada-Escudé, G. & Butler, R. P. 2012, ApJS, 200, 15

Anglada-Escude, G., Leger, A., & Malbet, F. 2014, Nearby Earth Astrometric Telescope, Tech. rep., Queen Mary Unversit of London, Institut d'Astrophysique Spatiale, Institut d'astrophyique de Grenoble

Anglada-Escudé, G., López-Morales, M., & Chambers, J. E. 2010, ApJ, 709, 168

Anglada-Escudé, G. & Tuomi, M. 2012, A&A, 548, A58

Artigau, É., Astudillo-Defru, N., Delfosse, X., et al. 2014, in Observatory Operations: Strategies, Processes, and Systems V, Vol. 9149, 914905

Bailey, J., Simpson, A., & Crisp, D. 2007, ApSP, 119, 228

Baliunas, S. L., Donahue, R. A., Soon, W., & Henry, G. W. 1998, in Astronomical Society of the Pacific Conference Series, Vol. 154, Cool Stars, Stellar Systems, and the Sun, ed. R. A. Donahue & J. A. Bookbinder, 153

Baluev, R. V. 2008, MNRAS, 385, 1279

Baluev, R. V. 2009, MNRAS, 393, 969

Baluev, R. V. 2011, Celestial Mechanics and Dynamical Astronomy, 111, 235

Baluev, R. V. 2013, Monthly Notices of the Royal Astronomical Society, 436, 807

Baluev, R. V. 2013a, Astronomy and Computing, 3, 50

Baluev, R. V. 2013b, MNRAS, 431, 1167

Baluev, R. V. 2013c, MNRAS, 429, 2052

Baluev, R. V. 2015a, MNRAS, 446, 1493

Baluev, R. V. 2015b, MNRAS, 446, 1478

Baranne, A., Queloz, D., Mayor, M., et al. 1996, A&A S, 119, 373

Bates, D. M. & Watts, D. G. 1980, Journal of the Royal Statistical Society. Series B (Methodological), 42, 1

Bauer, F. F., Zechmeister, M., & Reiners, A. 2015, A&A, 581, A117

Becker, S., Bobin, J., & Candès, E. J. 2011, SIAM Journal on Imaging Sciences, 4, 1

Benjamini, Y. & Hochberg, Y. 1995, Journal of the Royal Statistical Society. Series B (Methodological), 57, 289

Bobin, J., Starck, J.-L., & Ottensamer, R. 2008, IEEE Journal of Selected Topics in Signal Processing, 2, 718

Boisse, I., Bonfils, X., & Santos, N. C. 2012, A&A, 545, A109

Bonomo, A. S., Desidera, S., Benatti, S., et al. 2017a, ArXiv e-prints

Bonomo, A. S., Hébrard, G., Raymond, S. N., et al. 2017b, ArXiv e-prints

Bonomo, A. S., Santerne, A., Alonso, R., et al. 2010, A&A, 520, A65

Bouchy, F., Díaz, R. F., Hébrard, G., et al. 2013, A&A, 549, A49

Bouchy, F., Pepe, F., & Queloz, D. 2001, A&A, 374, 733

Bray, R. J., Loughhead, R. E., & Shore, B. W. 1965, Physics Today, 18, 68

Brewer, B. J. 2014, ArXiv e-prints

Brown, R. A. 2017, ArXiv e-prints

Bruzek, A. 1977, in Astrophysics and Space Science Library, Vol. 69, Illustrated Glossary for Solar and Solar-Terrestrial Physics, ed. A. Bruzek & C. J. Durrant, 71

Butler, R. P., Marcy, G. W., Williams, E., Hauser, H., & Shirts, P. 1997, ApJL, 474, L115

Campbell, B., Walker, G. A. H., & Yang, S. 1988, ApJ, 331, 902

Candès, E., Romberg, J., & Tao, T. 2006a, Information Theory, IEEE Transactions on, 52, 489

Candès, E. J. & Fernandez-Granda, C. 2013, Journal of Fourier Analysis and Applications, 19, 1229

Candes, E. J. & Tao, T. 2005, IEEE Transactions on Information Theory, 51, 4203

Candès, E., J., Romberg, J., K., & Tao, T. 2006b, Communications on Pure and Applied Mathematics, 59, 1207

Capitaine, N., Wallace, P. T., & Chapront, J. 2003, A&A, 412, 567

Casella, G. & Berger, R. 2001, Statistical Inference (Duxbury Resource Center)

Cavallini, F., Ceppatelli, G., & Righini, A. 1985, A&A, 143, 116

Cegla, H. M., Watson, C. A., Marsh, T. R., et al. 2012, MNRAS, 421, L54

Chabrier, G. & Küker, M. 2006, A&A, 446, 1027

Chandrasekaran, V., Recht, B., Parrilo, P., & Willsky, A. 2012, Foundations of Computational Mathematics, 12, 805

Chapman, G. A., Cookson, A. M., Dobias, J. J., & Walton, S. R. 2001, ApJ, 555, 462

Charbonneau, P. 2010, Living Reviews in Solar Physics, 7, 3

Chen, S. S., Donoho, D. L., & Saunders, M. A. 1995, Atomic Decomposition by Basis Pursuit

Chen, S. S., Donoho, D. L., & Saunders, M. A. 1998, SIAM JOURNAL ON SCIENTIFIC COMPUTING, 20, 33

Chen, Y. & Chi, Y. 2014, IEEE Transactions on Information Theory, 60, 6576

Christensen-Dalsgaard, J. 2014

Cohen, A., Dahmen, W., & Devore, R. 2009, J. Amer. Math. Soc, 211

Connes, P. 1985, ApSS, 110, 211

Cook, R. D. & Witmer, J. A. 1985, Journal of the American Statistical Association, 80, 872

Correia, A. C. M., Couetdic, J., Laskar, J., et al. 2010, A&A, 511, A21

Cosentino, R., Lovis, C., Pepe, F., et al. 2012, in Ground-based and Airborne Instrumentation for Astronomy IV, Vol. 8446, 84461V

Cotton, D. V., Bailey, J., & Kedziora-Chudczer, L. 2014, MNRAS, 439, 387

Cox, R. T. 1946, American Journal of Physics, 14, 1

Cox, R. T. 1961, The Algebra Of Probable Inference (John Hopkins University Press)

Cumming, A. 2004, MNRAS, 354, 1165

Cumming, A., Marcy, G. W., & Butler, R. P. 1999, ApJ, 526, 890

Cunha, D., Figueira, P., Santos, N. C., Lovis, C., & Boué, G. 2013, A&A, 550, A75

Cunha, D., Santos, N. C., Figueira, P., et al. 2014, A&A, 568, A35

Davies, R. B. 1977, Biometrika, 64, 247

Davies, R. B. 1987, Biometrika, 74, 33

Davies, R. B. 2002, Biometrika, 89, 484

Davis, A. B., Cisewski, J., Dumusque, X., Fischer, D. A., & Ford, E. B. 2017, ArXiv e-prints

Deeg, H. J., Moutou, C., Erikson, A., et al. 2010, Nature, 464, 384

Del Moro, D. 2004, A&A, 428, 1007

Delfosse, X., Forveille, T., Mayor, M., et al. 1998, A&A, 338, L67

Demoment, G., Idier, J., Giovanelli, J.-F., & Mohammad-Djafari, A. 2001, Techniques de l'ingénieur Traitement du signal : bases théoriques, base documentaire : TIB295DUO., fre

Deubner, F.-L. & Fleck, B. 1989, A&A, 213, 423

Donoho, D. 2006, Information Theory, IEEE Transactions on, 52, 1289

Donoho, D. L. 2005, Neighborly Polytopes and Sparse Solutions of Underdetermined Linear Equations, Tech. rep.

Donoho, D. L. & Elad, M. 2002, in PROC. NATL ACAD. SCI. USA 100 2197–202

Donoho, D. L., Elad, M., & Temlyakov, V. N. 2006, IEEE TRANS. INFORM. THEORY, 52, 6

Donoho, D. L. & Huo, X. 2001, IEEE Transactions on Information Theory, 47, 2845

Donoho, D. L. & Stark, P. B. 1989, SIAM J. Appl. Math., 49, 906

Duarte, M. F. & Baraniuk, R. G. 2013, Applied and Computational Harmonic Analysis, 35, 111

Dubois, D. & Prade, H. 2012, Possibility Theory (New York, NY: Springer New York), 2240–2252

Dumusque, X. 2016, A&A, 593, A5

Dumusque, X., Boisse, I., & Santos, N. C. 2014, ApJ, 796, 132

Dumusque, X., Borsa, F., Damasso, M., et al. 2016, ArXiv e-prints

Dumusque, X., Pepe, F., Lovis, C., et al. 2012, Nature, 491, 207

Dumusque, X., Santos, N. C., Udry, S., Lovis, C., & Bonfils, X. 2011a, A&A, 527, A82

Dumusque, X., Udry, S., Lovis, C., Santos, N. C., & Monteiro, M. J. P. F. G. 2011b, A&A, 525, A140

Eastman, J., Gaudi, B. S., & Agol, E. 2013, ApSP, 125, 83

Efron, B. 1979, Ann. Statist., 7, 1

Efroymson, M. A. 1960, Multiple Regression Analysis (John wiley)

Eggenberger, A. & Udry, S. 2010, in EAS Publications Series, Vol. 41, EAS Publications Series, ed. T. Montmerle, D. Ehrenreich, & A.-M. Lagrange, 27–75

Eldar, Y. C. 2009, IEEE Transactions on Signal Processing, 57, 2986

Endl, M., Robertson, P., Cochran, W. D., et al. 2012, ApJ, 759, 19

Faria, J. P., Haywood, R. D., Brewer, B. J., et al. 2016, A&A, 588, A31

Feng, F., Tuomi, M., & Jones, H. R. A. 2017, MNRAS, 470, 4794

Ferraz-Mello, S. 1981, AJ, 86, 619

Fienga, A., Manche, H., Laskar, J., & Gastineau, M. 2008, A&A, 477, 315

Figueira, P., Santos, N. C., Pepe, F., Lovis, C., & Nardetto, N. 2013, A&A, 557, A93

Firth, D. 1993, Biometrika, 80, 27

Fischer, D. A., Marcy, G. W., Butler, R. P., et al. 2008, ApJ, 675, 790

Ford+. 2017, in prep.

Ford, E. B. 2005, AJ, 129, 1706

Ford, E. B. 2006, ApJ, 642, 505

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, ApSP, 125, 306

Foster, G. 1995, AJ, 109, 1889

Foucart, S. 2010, Applied and Computational Harmonic Analysis, 29, 97

Foucart, S. & Lai, M.-J. 2009, Applied and Computational Harmonic Analysis, 26, 395

Foucart, S., Pajor, A., Rauhut, H., & Ullrich, T. 2010, Journal of Complexity, 26, 629

Foucart, S. & Rauhut, H. 2013, A Mathematical Introduction to Compressive Sensing

Gastineau, M. & Laskar, J. 2011, ACM Commun. Comput. Algebra, 44, 194

Ge, D., Jiang, X., & Ye, Y. 2011, Math. Program., 129, 285

Gelman, A. & Rubin, D. B. 1992, Statist. Sci., 7, 457

Goodman, J. & Weare, J. 2010, Communications in applied mathematics and computational science, 5, 65

Grant, M. & Boyd, S. 2008, in Recent Advances in Learning and Control, ed. V. Blondel, S. Boyd, & H. Kimura, Lecture Notes in Control and Information Sciences (Springer-Verlag Limited), 95–110, http://stanford.edu/~boyd/graph_dcp.html

Gregory, P. 2005, Bayesian Logical Data Analysis for the Physical Sciences (New York, NY, USA: Cambridge University Press)

Gregory, P. C. 2005, ApJ, 631, 1198

Gregory, P. C. 2007a, MNRAS, 374, 1321

Gregory, P. C. 2007b, MNRAS, 381, 1607

Grinsted. 2015, Ensemble MCMC sampler, https://fr.mathworks.com/matlabcentral/fileexchange/49820-ensemble-mcmc-sampler

Haldar, J. P., Hernando, D., & Liang, Z. P. 2011, IEEE Transactions on Medical Imaging, 30, 893

Hale, G. E. 1908, ApJ, 28, 315

Halpern, J. 1999a, Journal of AI research, 10, 67

Halpern, J. Y. 1999b, Journal of Artificial Intelligence Research, 11, 429

Hansen, C. J. & Kawaler, S. D. 1994, Stellar Interiors. Physical Principles, Structure, and Evolution., 84

Hara, N., Boué, G., & Laskar, J. 2018, in prep.

Hara, N. C., Boué, G., Laskar, J., & Correia, A. C. M. 2017, MNRAS, 464, 1220

Hartkopf, W. I., McAlister, H. A., & Franz, O. G. 1989, AJ, 98, 1014

Hartley, H. O. 1964, Biometrika, 51, 347

Harvey, J. 1985, in ESA Special Publication, Vol. 235, Future Missions in Solar, Heliospheric & Space Plasma Physics, ed. E. Rolfe & B. Battrick

Hastings, W. K. 1970, Biometrika, 57, 97

Hathaway, D. H. 2010, Living Reviews in Solar Physics, 7, 1

Haywood, R. D., Collier Cameron, A., Queloz, D., et al. 2014, MNRAS, 443, 2517

Haywood, R. D., Collier Cameron, A., Unruh, Y. C., et al. 2016, MNRAS, 457, 3637

Hébrard, G., Arnold, L., Forveille, T., et al. 2016, A&A, 588, A145

Hoeting, J. A., Madigan, D., Raftery, A. E., & Volinsky, C. T. 1999, Statist. Sci., 14, 382

Hogg, D. W., Myers, A. D., & Bovy, J. 2010, ApJ, 725, 2166

Horn, K. S. V. 2003, International Journal of Approximate Reasoning, 34, 3

Horne, J. H. & Baliunas, S. L. 1986, ApJ, 302, 757

Huber, P. J. 1985, Ann. Statist., 13, 435

Hurwicz, L. 1950, in Statistical inference in dynamic economic models, ed. T. Koopmans (New York: Wiley), 365–383

Husnoo, N., Pont, F., Mazeh, T., et al. 2012, MNRAS, 422, 3151

Husnoo, N., Pont, F., Mazeh, T., et al. 2011, in IAU Symposium, Vol. 276, The Astrophysics of Planetary Systems: Formation, Structure, and Dynamical Evolution, ed. A. Sozzetti, M. G. Lattanzi, & A. P. Boss, 243–247

Hussain, G. A. J. 2002, Astronomische Nachrichten, 323, 349

Ioannidis, J. P. A. 2005, PLOS Medicine, 2

Jacob, W. S. 1855, MNRAS, 15, 228

Jaynes, E. & Bretthorst, G. 2003, Probability Theory: The Logic of Science (Cambridge University Press Cambridge:)

Jeffreys, H. 1946, Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences, 186, 453

Jensen, J. L. W. V. 1906, Acta Math., 30, 175

Jurić, M. & Tremaine, S. 2008, ApJ, 686, 603

Kant, E. 1755, Allgemeine Naturgeschichte und Theorie des Himmels (Scottish Academic Press)

Kass, R. E. & Raftery, A. E. 1995, Journal of the American Statistical Association, 90, 773

Keynes, J. M. 1921, A Treatise on Probability (Macmillan and co.)

Kipping, D. M. 2014, MNRAS, 444, 2263

Kolmogorov, A. N. 1933, Giornale dell'Istituto Italiano degli Attuari, 4, 83

Kueveler, G. 1983, 88, 13

Kutyniok, G. 2012, ArXiv e-prints

Lagrange, A.-M., Desort, M., & Meunier, N. 2010, A&A, 512, A38

Laplace, P.-S. 1796, Exposition du système du monde (Imprimerie du Cercle-Social)

Laskar, J. 1986, A&A, 157, 59

Laskar, J. 1988, A&A, 198, 341

Laskar, J. 1990, Icarus, 88, 266

Laskar, J. 1993, Celestial Mechanics and Dynamical Astronomy, 56, 191

Laskar, J. 2003, ArXiv Mathematics e-prints

Laskar, J. 2008, Icarus, 196, 1

Laskar, J., Froeschlé, C., & Celletti, A. 1992, Physica D Nonlinear Phenomena, 56, 253

Latham, D. W., Stefanik, R. P., Mazeh, T., Mayor, M., & Burki, G. 1989, Nature, 339, 38

Lehmann, E. & Casella, G. 1998, Theory of Point Estimation (Springer Verlag)

Lehmann, E. L. & Romano, J. P. 2005, Testing statistical hypotheses, 3rd edn., Springer Texts in Statistics (New York: Springer), xiv+784

Levenberg, K. 1944, Quart. Applied Math., 2, 164

Llama, J., Jardine, M., Mackay, D. H., & Fares, R. 2012, MNRAS, 422, L72

Lockwood, G. W., Skiff, B. A., Henry, G. W., et al. 2007, ApJS, 171, 260

Lomb, N. R. 1976, ApSS, 39, 447

Lovis, C., Dumusque, X., Santos, N. C., et al. 2011a, ArXiv e-prints

Lovis, C. & Fischer, D. 2010, Radial Velocity Techniques for Exoplanets, 27–53

Lovis, C., Mayor, M., Pepe, F., et al. 2006, Nature, 441, 305

Lovis, C., Ségransan, D., Mayor, M., et al. 2011b, A&A, 528, A112

Lucy, L. B. 2013, A&A, 551, A47

Lucy, L. B. 2014, A&A, 565, A37

Lucy, L. B. & Sweeney, M. A. 1971, AJ, 76, 544

MacKay, D. J. C. 1999, Neural Computation, 11, 1035

MacKay, D. J. C. 2003, Information Theory, Inference, and Learning Algorithms (Cambridge University Press), available from http://www.inference.phy.cam.ac.uk/mackay/itila/

Mallat, S. & Zhang, Z. 1993, Trans. Sig. Proc., 41, 3397

Marcy, G. W., Butler, R. P., Fischer, D. A., et al. 2002, ApJ, 581, 1375

Marcy, G. W., Butler, R. P., Vogt, S. S., Fischer, D., & Lissauer, J. J. 1998, ApJL, 505, L147

Marcy, G. W., Butler, R. P., Williams, E., et al. 1997, ApJ, 481, 926

Marquardt, D. W. 1963, SIAM Journal on Applied Mathematics, 11, 431

Mayor, M., Pepe, F., Queloz, D., et al. 2003, The Messenger, 114, 20

Mayor, M. & Queloz, D. 1995, Nature, 378, 355

McArthur, B. E., Endl, M., Cochran, W. D., et al. 2004, ApJL, 614, L81

Melo. 2001, 2001

Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. 1953, The Journal of Chemical Physics, 21, 1087

Mishali, M., Eldar, Y. C., & Tropp, J. A. 2008, in Electrical and Electronics Engineers in Israel, 2008. IEEEI 2008. IEEE 25th Convention of, 290–294

Montes, D., Fernandez-Figueroa, M. J., de Castro, E., & Cornide, M. 1995, A&A, 294, 165

Mortier, A. & Collier Cameron, A. 2017, A&A, 601, A110

Mortier, A., Faria, J. P., Correia, C. M., Santerne, A., & Santos, N. C. 2015, A&A, 573, A101

Nardetto, N., Mourard, D., Kervella, P., et al. 2006, A&A, 453, 309

Needell, D. & Tropp, J. 2009, Applied and Computational Harmonic Analysis, 26, 301

Nelson, B. E., Ford, E. B., Wright, J. T., et al. 2014, MNRAS, 441, 442

Newton, I. 1687, Philosophiae naturalis principia mathematica (Jussu Societatis Regiæ ac Typis Josephi Streater)

O'Hara, R. B. & Sillanpää, M. J. 2009, Bayesian Anal., 4, 85

O'Toole, S. J., Tinney, C. G., Jones, H. R. A., et al. 2009, MNRAS, 392, 641

Paris, J. B. 1994, The Uncertain Reasoner's Companion: A Mathematical Perspective, Vol. 4 (Cambridge University Press)

Pasquini, L. & Pallavicini, R. 1991, A&A, 251, 199

Pati, Y. C., Rezaiifar, R., & Krishnaprasad, P. S. 1993, in Signals, Systems and Computers, 1993. 1993 Conference Record of The Twenty-Seventh Asilomar Conference on, 40–44 vol.1

Pelat, D. 2013, Bases et méthodes pour le traitement de données, Observatoire de Paris

Perruchot, S., Bouchy, F., Chazelas, B., et al. 2011, in Techniques and Instrumentation for Detection of Exoplanets V, Vol. 8151, 815115

Perruchot, S., Kohler, D., Bouchy, F., et al. 2008, in Ground-based and Airborne Instrumentation for Astronomy II, Vol. 7014, 70140J

Perryman, M. 2011, The Exoplanet Handbook (Cambridge University Press)

Petrovay, K. & van Driel-Gesztelyi, L. 1997, 176, 249

Petrovich, C., Tremaine, S., & Rafikov, R. 2014, ApJ, 786, 101

Phillips, A. C. 1999, The Physics of Stars, 2nd Edition, 262

Pont, F., Husnoo, N., Mazeh, T., & Fabrycky, D. 2011, MNRAS, 414, 1278

Potter, L. C., Ertin, E., Parker, J. T., & Cetin, M. 2010, Proceedings of the IEEE, 98, 1006

Prialnik, D. 2000, An Introduction to the Theory of Stellar Structure and Evolution

Prony, R. 1795, Journal de l'École Polytechnique

Queloz, D., Henry, G. W., Sivan, J. P., et al. 2001, A&A, 379, 279

Quenouille, M. H. 1949, Ann. Math. Statist., 20, 355

Rajpaul, V., Aigrain, S., Osborne, M. A., Reece, S., & Roberts, S. J. 2015, ArXiv e-prints

Rasmussen, C. E. & Williams, C. K. I. 2005, Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning) (The MIT Press)

Reegen, P. 2007, A&A, 467, 1353

Reeves, G. & Gastpar, M. 2009, in 2009 Conference Record of the Forty-Third Asilomar Conference on Signals, Systems and Computers, 1576–1580

Reiners, A. 2009, A&A, 498, 853

Rice, S. O. 1944, Bell System Technical Journal, 23, 282

Rickman, H. 2001, Transactions of the International Astronomical Union, Series B, 24

Rissanen, J. 1978, Automatica, 14, 465

Rivera, E. J., Laughlin, G., Butler, R. P., et al. 2010, The Astrophysical Journal, 719, 890

Rivera, E. J., Lissauer, J. J., Butler, R. P., et al. 2005, ApJ, 634, 625

Roberts, D. H., Lehar, J., & Dreher, J. W. 1987, AJ, 93, 968

Roudier, T., Malherbe, J. M., Vigneau, J., & Pfeiffer, B. 1998, A&A, 330, 1136

Rényi, A. 2007, Probability theory (Mineola, N.Y.: Dover Publications)

Sahnoun, S., Djermoune, E.-H., & Brie, D. 2012, in 20th European Signal Processing Conference, EUSIPCO 2012, Bucarest, Romania, 450–454

Sanchis-Ojeda, R. & Winn, J. N. 2011, ApJ, 743, 61

Sanchis-Ojeda, R., Winn, J. N., Holman, M. J., et al. 2011, ApJ, 733, 127

Santerne, A., Díaz, R. F., Almenara, J.-M., et al. 2015, MNRAS, 451, 2337

Satterthwaite, F. E. 1946, Biometrics Bulletin, 2, 110

Scargle, J. D. 1982, ApJ, 263, 835

Schmidt, R. 1986, IEEE Transactions on Antennas and Propagation, 34, 276

Schrijver, C. J. 2002, Astronomische Nachrichten, 323, 157

Schuster, A. 1898, Terrestrial Magnetism, 3, 13

Schwarz, G. 1978, Ann. Statist., 6, 461

Seifahrt, A., Käufl, H. U., Zängl, G., et al. 2010, A&A, 524, A11

Shafer, G. 1976, A Mathematical Theory of Evidence (Princeton: Princeton University Press)

Shannon, C. 1948, Bell system technical journal, 27

Shapiro, S. S. & Wilk, M. B. 1965, Biometrika, 52, 591

Shen, Y. & Turner, E. L. 2008, ApJ, 685, 553

Solomonoff, R. 1964a, Information and Control, 7, 1

Solomonoff, R. 1964b, Information and Control, 7, 224

Soric, B. 1989, Journal of the American Statistical Association, 84, 608

Southworth, J. 2011, MNRAS, 417, 2166

Spiegel, E. A. & Zahn, J.-P. 1992, A&A, 265, 106

Standish, Jr., E. M. 1990, A&A, 233, 252

Stoica, P. & Babu, P. 2012, Signal Processing, IEEE Transactions on, 60, 962

Strassmeier, K. G. 2009, 17, 251

Strassmeier, K. G. & Hall, D. S. 1994, in Astronomical Society of the Pacific Conference Series, Vol. 64, Cool Stars, Stellar Systems, and the Sun, ed. J.-P. Caillault, 483

Swedenborg, E. 1734, Opera Philosophica et Mineralia (Imprimerie du Cercle-Social)

Tang, G., Bhaskar, B., & Recht, B. 2013a, in Signals, Systems and Computers, 2013 Asilomar Conference on, 1043–1047

Tang, G., Bhaskar, B., Shah, P., & Recht, B. 2013b, Information Theory, IEEE Transactions on, 59, 7465

Tibshirani, R. 1994, Journal of the Royal Statistical Society, Series B, 58, 267

Tropp, J. A. 2004, IEEE Transactions on Information Theory, 50, 2231

Tropp, J. A., Laska, J. N., Duarte, M. F., Romberg, J. K., & Baraniuk, R. G. 2010, IEEE Transactions on Information Theory, 56, 520

Tuomi, M. 2011, A&A, 528, L5

Tuomi, M., Jones, H. R. A., Jenkins, J. S., et al. 2013, A&A, 551, A79

Tuomi, M. & Kotiranta, S. 2009, A&A, 496, L13

van de Kamp, P. 1963, AJ, 68, 515

van de Kamp, P. 1969, AJ, 74, 757

van den Berg, E. & Friedlander, M. P. 2008, SIAM Journal on Scientific Computing, 31, 890

Vaughan, A. H., Preston, G. W., & Wilson, O. C. 1978, ApSP, 90, 267

Walder, R., Folini, D., & Meynet, G. 2012, 166, 145

Welch, B. L. 1947, Biometrika, 34, 28

Wolfgang, A., Rogers, L. A., & Ford, E. B. 2016, ApJ, 825, 19

Wolszczan, A. & Frail, D. A. 1992, Nature, 355, 145

Zakamska, N. L., Pan, M., & Ford, E. B. 2011, MNRAS, 410, 1895

Zechmeister, M., Anglada-Escudé, G., & Reiners, A. 2014, A&A, 561, A59

Zechmeister, M. & Kürster, M. 2009a, A&A, 496, 577

Zechmeister, M. & Kürster, M. 2009b, A&A, 496, 577

Appendix A

A note on optimum weight

The goal of Bouchy et al. (2001) is to compute the fundamental noise limit on the radial velocity estimated from a spectrum. It follows a computation of (Connes 1985) that proceeds this way.

The Doppler shift δV is given in a non relativistic setting as

$$\frac{\delta V}{c} = \frac{\delta \lambda}{\lambda} \tag{A.1}$$

where $\delta\lambda$ is the shift in frequency. One considers then an ideal noiseless spectrum A_0, normalized, that corresponds to a reference velocity. Now let us consider a measured spectrum A also normalized, noisy, as an array whose i-th value is the power measured by a CCD pixel centred on λ_i. Then the shift in wavelength and velocity $\delta\lambda(i)$ respectively $\delta V(i)$ estimated from the pixel i verifies

$$A(i) - A_0(\lambda_i) = A(\lambda_i - \delta\lambda(i)) - A_0(\lambda_i) \approx -\frac{\partial A_0}{\partial \lambda}(\lambda_i)\delta\lambda(i) = -\frac{\partial A_0}{\partial \lambda}(\lambda_i)\frac{\lambda_i \delta V(i)}{c}. \tag{A.2}$$

We use notations $A(i)$ -and $A_0(\lambda_i)$ since A is measured on the CCD and A_0 is a theoretical spectrum. Then

$$\frac{\delta V(i)}{c} = -\frac{A(i) - A_0(\lambda_i)}{\lambda(i)(\partial A_0(\lambda_i))/\partial\lambda}. \tag{A.3}$$

The estimate of the Doppler velocity is now supposed to be a weighted sum of individual estimates as given in the equation above,

$$\frac{\delta V}{c} = \frac{\sum \frac{\delta V(i)}{c} W(i)}{\sum W(i)}.$$

It is then claimed that the optimal weighting is given by

$$W(i) = \frac{1}{\left(\frac{\delta V_{\mathrm{RMS}(i)}}{c}\right)^2}.$$

In Connes (1985), at this point, it is mentioned that the full calculation is available in Connes (1984), but which is unpublished and unavailable. Though not precisely defined, it seems clear

that in this context optimality means that the RV estimate is as precise as possible, since one wants to find the lowest error possible.

We just wish to make a simple point: P. Connes probably meant that $\delta V(i)$ are considered as independent random variables of known variances, and he wished to minimize the variance of their sum $\frac{\delta V}{c}$. Considering a weighted sum means that the expectancy remains the same as the mean of the arithmetic mean of the $\delta V(i)$ and is therefore unbiased. At least with those two assumptions on what the claim is, it is correct. To see that, let us remark that, denoting by $\epsilon_i = \delta V(i)/c$ the random variables and σ_i^2 their variances. As long as the ϵ_i are independent,

$$\mathrm{Var}\left\{\sum_{i=1}^{n} w_i \epsilon_i\right\} = \sum_{i=1}^{n} w_i^2 \sigma_i^2 \tag{A.4}$$

We then wish to solve

$$\arg\min_{w_1 \dots w_n} \sum_{i=1}^{n} w_i^2 \sigma_i^2 \quad s.t. \quad \sum_{i=1}^{n} w_i = 1. \tag{A.5}$$

The Lagrangian function associated to that minimization problem is

$$L(w, \lambda) = \sum_{i=1}^{n} w_i^2 \sigma_i^2 + \lambda \left(\sum_{i=1}^{n} w_i - 1 \right). \tag{A.6}$$

Differentiating with respect to w and λ yields

$$\forall i, \quad 2 w_i \sigma_i^2 + \lambda = 0 \tag{A.7}$$

$$\sum_{i=1}^{n} w_i = 1 \tag{A.8}$$

multiplying the first equations by $1/\sigma_i^2$ and summing up yields $\lambda = -2/\sum \frac{1}{\sigma_i^2}$. We then obtain the weights w_i as desired

$$\forall i, \quad w_i = \frac{\frac{1}{\sigma_i^2}}{\sum\limits_{i=1}^{n} \frac{1}{\sigma_i^2}} = \frac{\frac{1}{\left(\frac{\delta V_{\mathrm{RMS}(i)}}{c}\right)^2}}{\sum\limits_{i=1}^{n} \frac{1}{\left(\frac{\delta V_{\mathrm{RMS}(i)}}{c}\right)^2}} = \frac{W(i)}{\sum\limits_{i=1}^{n} W(i)} \tag{A.9}$$

As a concluding remark, $A_0(\lambda)$ is defined as the mean number of photon per unit time of the theoretical spectrum. Denoting by $N(i) = n(i) + n_D(i)$ the number of photons counted at pixel i during the integration time Δt, where n is the number of photons from the slit and n_D the error due to the detector, we have $A(i) = N(i)/\Delta t$ and

$$\sigma_i^2 \mathrm{Var}\left\{\frac{\delta V}{c}(i)\right\} = \frac{\frac{\sigma_D^2}{(\Delta t)^2} + \frac{A_0(\lambda_i)}{\Delta t}}{\lambda_i^2 \frac{\partial A_0}{\partial \lambda}(\lambda_i)} \tag{A.10}$$

where $\sigma_D^2(i) = \mathrm{Var}\{n_D(i)\}$. Given that $\mathrm{Var}\{\sum_i w_i \epsilon_1\} = \sum_i \sigma_i^2/\sigma_i^4/(\sum_i 1/\sigma_i^2)^2 = 1/(\sum_i 1/\sigma_i^2)$ we obtain the formula (1.10). Note that we express the formula in terms of flux and integration time instead of counted electrons as Bouchy et al. (2001). The reason behind this choice is that formula (1.10) emerges from the computation of a variance. As a consequence the mean quantities have to appear. The number of electrons received is, from a statistical point of view, an estimator of the mean of the Poisson process (the flux F) times the integration time.

Appendix B

Proof of theorem 5

Let us consider the following minimization problem

$$\underset{\boldsymbol{x}\in\mathbb{R}^n}{\arg\min} \quad \|\boldsymbol{x}\|_{\ell_q} \quad \text{s.t.} \quad \mathbf{A}\boldsymbol{x} = \boldsymbol{y}. \tag{P_{ℓ_q}}$$

We wish to prove the following theorem.

Theorem (Foucart & Lai (2009), modified)**.** *Provided that*

$$\gamma_{S,1} < s^{1/q-1/2}4(\sqrt{2}-1)+1 \tag{B.1}$$

and

$$\beta_2^2 - \alpha_2^2 \leqslant \beta_{S,1}^2 - \alpha_{S,1}^2 \tag{B.2}$$

then all vector $\boldsymbol{x}_t$ with support in S is exactly recovered by solving (P_{ℓ_q}) with $\boldsymbol{y} = \mathbf{A}\boldsymbol{x}_t$.

where

Definition (Condition number bound relative to S)**.** *Let S be a subset of s indices of $\{1..n\}$. Let be $\gamma_{S,1} = \beta_{S,1}^2/\ \alpha_{S,1}^2$ where $\beta_{S,1}$ and $\alpha_{S,1}$ are the tightest constants such that for any vector $\boldsymbol{a}_k$, k^{th} column of $\mathbf{A}$ with $k \in \bar{S}$, the concatenation of $\mathbf{A}_S$ and $\boldsymbol{a}_k$, $[\mathbf{A}_S, \boldsymbol{a}_k]$ verifies*

$$\forall \boldsymbol{z} \in \mathbb{R}^n \text{ or } \mathbb{C}^n, \|\boldsymbol{z}\|_{\ell_0} \leqslant s+1, \qquad \alpha_{S,1}\|\boldsymbol{z}\|_{\ell_2} \leqslant \|[\mathbf{A}_S, \boldsymbol{a}_k]\boldsymbol{z}\|_{\ell_2} \leqslant \beta_{S,1}\|\boldsymbol{z}\|_{\ell_2} \tag{B.3}$$

The following demonstration is very similar to the demonstration of step 1 of Theorem 2.1 in Foucart & Lai (2009). The change consists in obtaining a recovery condition not for all signals with s non zero components, but only the signals that has non zero components in a set of indices S. This leads to a weaker result in some sense, but also to weaker conditions for exact recovery that can be hoped to be checked in practice. In the proof, with the paper notation, we take $t = 1$ instead of $t \geqslant s$. We only prove the noiseless case, which necessitates only to prove one part of the theorem. Finally we use our hypothesis equation (B.13).

Proof. Let us consider a vector $\boldsymbol{v} \in \ker\mathbf{A}$. We will prove that under the conditions of the theorem, $\|\boldsymbol{v}_S\|_{\ell_1} < \|\boldsymbol{v}_{\bar{S}}\|_{\ell_1}$. In other words $\mathbf{A}$ satisfies the ℓ_q null space property relative to

S, which ensures uniqueness of the solution of (P_{ℓ_q}) for vectors with components in S. Let us note $\boldsymbol{v}_k$, $k = 1..n - |S|$ the components of $\boldsymbol{v}$ in $\bar{S}$ ordered by decreasing value. We have

$$\|\boldsymbol{v}_S\|^2 + \|\boldsymbol{v}_1\|^2 = \|\boldsymbol{v}_S + \boldsymbol{v}_1\|^2 \leqslant \frac{1}{\alpha_{S,1}^2}\|\mathbf{A}(\boldsymbol{v}_S + \boldsymbol{v}_1)\|^2 \tag{B.4}$$

$$= \frac{1}{\alpha_{S,1}^2}\langle \mathbf{A}(\boldsymbol{v} - \boldsymbol{v}_2 - ...), \mathbf{A}(\boldsymbol{v}_S + \boldsymbol{v}_1)\rangle \tag{B.5}$$

$$= \frac{1}{\alpha_{S,1}^2}\langle \mathbf{A}\boldsymbol{v}, \mathbf{A}(\boldsymbol{v}_S + \boldsymbol{v}_1)\rangle + \frac{1}{\alpha_{S,1}^2}\sum_{k\geqslant 2}[\langle \mathbf{A}(-\boldsymbol{v}_k), \mathbf{A}\boldsymbol{v}_S\rangle + \langle \mathbf{A}(-\boldsymbol{v}_k), \mathbf{A}\boldsymbol{v}_1\rangle] \tag{B.6}$$

$$= \frac{1}{\alpha_{S,1}^2}\sum_{k\geqslant 2}[\langle \mathbf{A}(-\boldsymbol{v}_k), \mathbf{A}\boldsymbol{v}_S\rangle + \langle \mathbf{A}(-\boldsymbol{v}_k), \mathbf{A}\boldsymbol{v}_1\rangle] \tag{B.7}$$

The last equality coming from $\mathbf{A}\boldsymbol{v} = 0$. Let us denote by $\boldsymbol{u}_k = -\boldsymbol{v}_k/\|\boldsymbol{v}_k\|$ and $\boldsymbol{u}_S = \boldsymbol{v}_S/\|\boldsymbol{v}_S\|$. Then

$$\frac{\langle \mathbf{A}(-\boldsymbol{v}_k), \mathbf{A}\boldsymbol{v}_S\rangle}{\|\boldsymbol{v}_k\|\|\boldsymbol{v}_S\|} = \langle \mathbf{A}\boldsymbol{u}_k, \mathbf{A}\boldsymbol{u}_S\rangle = \frac{1}{4}\left(\|\mathbf{A}(\boldsymbol{u}_k + \boldsymbol{u}_S)\|^2 - \|\mathbf{A}(\boldsymbol{u}_k - \boldsymbol{u}_S)\|^2\right) \tag{B.8}$$

$$\leqslant \frac{1}{4}\left(\beta_{S,1}^2\|\boldsymbol{u}_k + \boldsymbol{u}_S\|^2 - \alpha_{S,1}^2\|\boldsymbol{u}_k - \boldsymbol{u}_S\|^2\right) = \frac{1}{2}(\beta_{S,1}^2 - \alpha_{S,1}^2). \tag{B.9}$$

Since $\boldsymbol{u}_k$ and $\boldsymbol{u}_S$ are orthogonal, therefore $\|\boldsymbol{u}_k + \boldsymbol{u}_S\|^2 = \|\boldsymbol{u}_k\|^2 + \|\boldsymbol{u}_S\|^2 = 2$. Similarly, since $\boldsymbol{v}_k$ and $\boldsymbol{v}_1$ have only one non zero component,

$$\frac{\langle \mathbf{A}(-\boldsymbol{v}_k), \mathbf{A}\boldsymbol{v}_1\rangle}{\|\boldsymbol{v}_k\|\|\boldsymbol{v}_1\|} = \langle \mathbf{A}\boldsymbol{u}_k, -\mathbf{A}\boldsymbol{u}_1\rangle = \frac{1}{4}\left(\|\mathbf{A}(\boldsymbol{u}_k - \boldsymbol{u}_1)\|^2 - \|\mathbf{A}(\boldsymbol{u}_k + \boldsymbol{u}_1)\|^2\right) \tag{B.10}$$

$$\leqslant \frac{1}{4}\left(\beta_2^2\|\boldsymbol{u}_k - \boldsymbol{u}_1\|^2 - \alpha_2^2\|\boldsymbol{u}_k + \boldsymbol{u}_1\|^2\right) = \frac{1}{2}(\beta_2^2 - \alpha_2^2). \tag{B.11}$$

Combining the two inequalities above yields

$$\langle \mathbf{A}(-\boldsymbol{v}_k), \mathbf{A}\boldsymbol{v}_S\rangle + \langle \mathbf{A}(-\boldsymbol{v}_k), \mathbf{A}\boldsymbol{v}_1\rangle \leqslant \frac{\|\boldsymbol{v}_k\|}{2}\left((\beta_{S,1}^2 - \alpha_{S,1}^2)\|\boldsymbol{v}_S\| + (\beta_2^2 - \alpha_2^2)\|\boldsymbol{v}_1\|\right) \tag{B.12}$$

Now using our hypothesis $\beta_{S,1}^2 - \alpha_{S,1}^2 \geqslant \beta_2^2 - \alpha_2^2$,

$$\langle \mathbf{A}(-\boldsymbol{v}_k), \mathbf{A}\boldsymbol{v}_S\rangle + \langle \mathbf{A}(-\boldsymbol{v}_k), \mathbf{A}\boldsymbol{v}_1\rangle \leqslant \frac{\beta_{S,1}^2 - \alpha_{S,1}^2}{2}\|\boldsymbol{v}_k\|(\|\boldsymbol{v}_S\| + \|\boldsymbol{v}_1\|) \tag{B.13}$$

We then have

$$\|\boldsymbol{v}_S\|^2 + \|\boldsymbol{v}_1\|^2 \leqslant \left(\frac{\gamma_{S,1} - 1}{2}\sum_{k\geqslant 2}\|\boldsymbol{v}_k\|\right)(\|\boldsymbol{v}_S\| + \|\boldsymbol{v}_1\|) \tag{B.14}$$

Let us denote by $d = \frac{\gamma_{S,1}-1}{2}$ et $\Sigma = \sum_{k\geqslant 2} |v_k|$. The above inequality can be re-written

$$\left(\|\boldsymbol{v}_S\| - \frac{d\Sigma}{2}\right)^2 + \left(\|\boldsymbol{v}_1\| - \frac{d\Sigma}{2}\right)^2 \leqslant \frac{d^2\Sigma^2}{2} \tag{B.15}$$

from which we deduce

$$\|\boldsymbol{v}_S\| \leqslant \frac{1+\sqrt{2}}{2} d\Sigma \tag{B.16}$$

Moreover, $\Sigma = \sum_{k\geqslant 2} \|\boldsymbol{v}_k\| \leqslant \sum_{k\geqslant 1} \|\boldsymbol{v}_k\| = \sum_{k\geqslant 1} |v_k| = \|\boldsymbol{v}_{\bar{S}}\|_{\ell_1}$. Using Hölder inequality,

$$\|\boldsymbol{v}_S\|_{\ell_1} \leqslant s^{1/2-1/q}\|\boldsymbol{v}_S\| \leqslant \frac{1+\sqrt{2}}{2} d\|\boldsymbol{v}_{\bar{S}}\|_{\ell_1} \leqslant \frac{1+\sqrt{2}}{2} d\|\boldsymbol{v}_{\bar{S}}\|_{\ell_q}. \tag{B.17}$$

As a consequence, $\|\boldsymbol{v}_S\|_{\ell_q} < \|\boldsymbol{v}_{\bar{S}}\|_{\ell_q}$ as soon as $\gamma_{S,1} < s^{1/q-1/2}4(\sqrt{2}-1)+1$. In case $q = 1$, the condition becomes $\gamma_{S,1} < \frac{1}{\sqrt{s}}4(\sqrt{2}-1)+1$. □

Appendix C

Formulas

C.1 First order approximation

In this section, we developp the Keplerian model to first order in eccentricity to obtain an analytical expression of the bias. Within this approximation, the distribution of the least square fit knowing that $e = 0$ is given in Lucy & Sweeney (1971). This section extends their formula to small e, and takes into account the number of fitted parameters. First, we develop (1.2) to order one in e, obtaining

$$y(\lambda, K, P, e, \omega) = K(\cos(\lambda) + e\cos(2\lambda - \omega)) \tag{C.1}$$

where $\lambda = nt + \omega = \lambda_0 + 2\pi t/P$ is the mean longitude, λ_0 being its value at $t = 0$. Denoting by $n = 2\pi/P$ the mean motion, the above expression can be re-written

$$y^{(1)}(t, A, B, C, D, n) = A\cos nt + B\sin nt + C\cos 2nt + D\sin 2nt \tag{C.2}$$

where $A = K\cos\lambda_0$, $B = -K\sin\lambda_0$, $C = Ke\cos(2\lambda_0 - \omega)$, $D = -Ke\sin(2\lambda_0 - \omega)$. When other parameters are fitted, the uncertainties on A, B, C, D increases as well. To quantify this effect, we consider the problem of fitting the period and a constant.

$$y^{(2)}(t, A, B, C, D, E, F) = A\cos nt + B\sin nt + C\cos 2nt + D\sin 2nt + E\frac{\partial y}{\partial n}(t) + F. \tag{C.3}$$

which in a matrix form gives

$$\boldsymbol{y}^{(2)}(\boldsymbol{t}, A, B, C, D, E, F) = \mathbf{M}(P)\boldsymbol{x}. \tag{C.4}$$

Let us assume that the observations are $\boldsymbol{y}(\boldsymbol{t}) = \mathbf{M}(P)\boldsymbol{x}_t + \boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon}$ is a Gaussian noise, independent and identically distributed with variance σ^2. The least square estimate of $\boldsymbol{x}$ is $\widehat{\boldsymbol{x}} = (\mathbf{M}^T\mathbf{M})^{-1}\mathbf{M}^T\boldsymbol{y}$, and the estimate of eccentricity is

$$\widehat{e} = \sqrt{\frac{\widehat{C}^2 + \widehat{D}^2}{\widehat{A}^2 + \widehat{B}^2}} = \frac{\sqrt{\widehat{C}^2 + \widehat{D}^2}}{K_t}\left(\frac{\sqrt{\widehat{A}^2 + \widehat{B}^2}}{K_t}\right)^{-1}, \tag{C.5}$$

where K_t is the true semi-amplitude. By change of random variable we can obtain the law followed by $\widehat{e}$. If we assume that N is large enough then the columns of $\mathbf{M}(P)$ are

approximately orthogonal, the components of $\widehat{x}$ are independent Gaussian variables. Since the modulus of a sum of independent Gaussian variables follows a Rice distribution,

$$U \equiv \frac{\sqrt{\widehat{C}^2 + \widehat{D}^2}}{K_t} \sim g(u) = S^2 u \, \mathrm{e}^{-\frac{S^2}{2}(u^2+e_t^2)} \, I_0(S^2 e u) \tag{C.6}$$

$$W \equiv \frac{\sqrt{\widehat{A}^2 + \widehat{B}^2}}{K_t} \sim h(w) = S^2 w \, \mathrm{e}^{-\frac{S^2}{2}(w^2+1)} \, I_0(S^2 w) \tag{C.7}$$

where I_0 is a modified Bessel function of first kind, $S = K_t/\sigma$ is the signal to noise ratio, where σ is the standard deviation of $\widehat{A}, \widehat{B}, \widehat{C}$ and $\widehat{D}$. If K is sufficiently large, W is close to 1 and $g(u)$ gives a good approximation of the law followed by the eccentricity fitted. Within this approximation, one can obtain analytical formula for the bias b of the eccentricity that only depends on the true eccentricity and the signal to noise ratio,

$$b(e_t, S, n) = \frac{1}{S}\sqrt{\frac{\pi}{2}} L_{1/2}\left(\frac{S^2 e_t^2}{2}\right) - e_t. \tag{C.8}$$

where $L_{1/2}$ is the Laguerre polynomial of order $1/2$. In case K_t is small, one must use the formula for the law followed by the quotient of two random variables:

$$\widehat{e} = \frac{U}{W} \sim f(e) = \int_{-\infty}^{+\infty} g(u) h(ue) |u| du \tag{C.9}$$

but no simple analytical expression was found.

When fitting model (C.4) to $\boldsymbol{y}(\boldsymbol{t})$, the estimate $\widehat{\boldsymbol{\theta}}$ will have a covariance matrix $\boldsymbol{\Sigma}^{-1}$ where $\boldsymbol{\Sigma} = \sigma^2(\mathbf{M}(P)^T\mathbf{M}(P))$ (this is a classical statistical result, see for example Pelat (2013)). The variances of the components of $\widehat{\boldsymbol{x}}$ are given by the diagonal elements of $\boldsymbol{\Sigma}^{-1}$. Their approximate calculation is the object of the next section. For a deeper analysis, one can study the behaviour of the error on average or for each n, which is done section C.2.

C.1.1 Average error

First we consider the estimation of the error on A, B, C, D when averaging over the mean motion n. At little cost, we will generalize our claim to the fitting of model (C.3) plus fitting other linearised Keplerian model. This approximately corresponds to fitting a multi-planetary system starting closely from the correct local minimum of χ^2. Again, the model can be written as a linear one, $\boldsymbol{y} = \mathbf{M}\boldsymbol{x}$ but where $\mathbf{M}$ has $p = 6+5k$ columns, k being the number of additional planets. To facilitate the discussion, we will normalize the columns of $\mathbf{M}$. To have an expression of the model of the form (C.4), we have multiply the k^{th} component of θ by the norm of the k^{th} column of M. The variances of these new model parameters are still given by the diagonal elements of $\sigma^2\boldsymbol{\Sigma}^{-1}$ where $\boldsymbol{\Sigma} = (\mathbf{M}^T\mathbf{M})$, but now $\boldsymbol{\Sigma}$ has only ones on its diagonal.

Calculating precisely the uncertainty on A, B, C, D averaged over n and the phase of the signal as a function of the instant of observations t is complex since it requires the inversion of σ which is a $6 + 5k \times 6 + 5k$ matrix. We will instead use an approximation that grasps the effect we want to estimate: how the uncertainty worsens as more parameters are added to the model. We will consider that the elements of $\mathbf{M}$ are drawn from independent Gaussian laws

that have a variance $1/N$. To avoid confusion with the true model, the so defined random matrix is denoted by $\tilde{\mathbf{M}}$ and its covariance matrix by $\tilde{\mathbf{\Sigma}}$.

This approximation seems to be rough at first but will turn out to be surprisingly accurate as a lower bound in practice. We will list a few arguments that make it a reasonable guess:

- The variances of the entries were chosen such that the expectancy of a squared norm of a column is one, which is the value of $\mathbf{\Sigma}$ diagonal elements.
- The columns are cosines and sines, which are approximately orthogonal, and in the Gaussian case decorrelation implies independence. Furthermore, the average of the spectral window is equal to the expected value of a correlation between two random Gaussian variables.
- The normed vectors $\cos \nu \boldsymbol{t}$ and $\sin \nu \boldsymbol{t}$ are approximately distributed uniformly on the sphere of $\mathbb{R}^N$ when ν is distributed uniformly between 0 and $2\pi/T_{\text{obs}}$.

The expected value of the variance of any parameter is the expected value of any diagonal element of $\mathbf{\Sigma}^{-1}$, since all the columns of $\tilde{\mathbf{M}}$ follow the same law. To tackle that problem, we rewrite $\tilde{\mathbf{\Sigma}}$ as

$$\tilde{\mathbf{\Sigma}} = \sigma^2 \begin{pmatrix} \mathbf{\Sigma}_{11} & \mathbf{\Sigma}_1^T \\ \mathbf{\Sigma}_1 & \mathbf{\Sigma}_c \end{pmatrix}$$

Where $\mathbf{\Sigma}_{11}$ is $\tilde{\mathbf{\Sigma}}$ element at first row and first column and $\mathbf{\Sigma}_1$ is a column vector with $N-1$ entries. We now have

$$\mathbb{E}\{\tilde{\mathbf{\Sigma}}_{11}^{-1}\} = \frac{1}{\sigma^2}\mathbb{E}\left\{\frac{1}{\mathbf{\Sigma}_{11} - \mathbf{\Sigma}_1^T \mathbf{\Sigma}_c^{-1} \mathbf{\Sigma}_1}\right\}$$

By Jensen inequality (Jensen 1906), since $x \to 1/x$ is convex,

$$\mathbb{E}\{\tilde{\mathbf{\Sigma}}_{11}^{-1}\} \leqslant \frac{1}{\sigma^2}\frac{1}{\mathbb{E}\left\{\mathbf{\Sigma}_{11} - \mathbf{\Sigma}_1^T \mathbf{\Sigma}_c^{-1} \mathbf{\Sigma}_1\right\}}.$$

Now since for two independent variables X and Y, $\mathbb{E}\{XY\} = \mathbb{E}\{X\}\mathbb{E}\{Y\}$,

$$\mathbf{\Sigma}_1^T \mathbf{\Sigma}_c^{-1} \mathbf{\Sigma}_1 = \sum_{k=2}^{p} \mathbb{E}\{\mathbf{\Sigma}_{1k}^2\}\mathbb{E}\{\mathbf{\Sigma}_{c,kk}^{-1}\} \leqslant \sum_{k=2}^{p} \mathbb{E}\{\mathbf{\Sigma}_{1k}^2\} = 1 - \frac{p-1}{N}$$

As by construction $\mathbb{E}\left\{\mathbf{\Sigma}_{11}\right\} = 1$, we finally obtain

$$\mathbb{E}\{\tilde{\mathbf{\Sigma}}_{11}^{-1}\} \leqslant \frac{1}{\sigma^2}\frac{1}{1 - \frac{p-1}{N}}$$

where the inequality follows again from Jensen's inequality applied to matrix inversion. Finally, the standard deviation on $I = A, B, C, D$ is

$$\sigma_I \geqslant \sigma \sqrt{\frac{1}{1 - \frac{p-1}{N}}} \tag{C.10}$$

With the approximation $\|\cos \nu \boldsymbol{t}\| \approx \|\sin \nu \boldsymbol{t}\| \approx \sqrt{N/2}$, the errors on $k = C/\sqrt{A^2+B^2}$ and $k = D/\sqrt{A^2+B^2}$ then verify

$$\sigma_k \gtrsim \frac{\sigma}{K_t}\sqrt{\frac{2}{N}}\sigma_I = \frac{\sigma}{K_t}\sqrt{\frac{2}{N-p+1}} =: 1/S. \tag{C.11}$$

As k and h approximately follow a Gaussian law, $e = \sqrt{k^2+h^2}$ follows a Rice distribution, whose mean is given by

$$\mathbb{E}\{\widehat{e}\} = \frac{1}{S}\sqrt{\frac{\pi}{2}}L_{1/2}\left(\frac{S^2 e_t^2}{2}\right). \tag{C.12}$$

$$\mathbb{E}\{\widehat{e}|e_t = 0\} = \frac{\sigma}{K_t}\sqrt{\frac{\pi}{N-p+1}} \tag{C.13}$$

$L_{1/2}$ being the Laguerre polynomial of degree $1/2$. The relevance of formula (C.12) is checked on numerical examples next section. As we shall see, the lower bound is tight when p does not exceeds $\approx N/2$.

Let us finally stress that formula (C.12) approximates the bias averaged on the mean motion, that is the frequency of the orbit, not the period. Averaging on the period would give more weight to the bias at low frequencies, which is high, and would therefore lead to a greater average value of the bias.

C.2 Error per period

The formula (C.12) shows what happens on average for all periods. However, in some cases the model $\boldsymbol{M}(\boldsymbol{t})$ can have a behaviour that is very different from the asymptotic expressions for large N. We made some attempts to approximate the behaviour of the error as a function of n, and we found unlikely that a formula would sum up relevant information. It seems best to directly compute the diagonal elements of $\boldsymbol{\Sigma}^{-1}$ to obtain the variances of A, B, C, D and to compute numerically the mean fitted eccentricity via formula (C.5), where A, B, C, D follow Gaussian laws.

For the sake of completeness, let us note that the diagonal of $\boldsymbol{\Sigma}(n)^{-1}$ will drastically differ from one if $\boldsymbol{\Sigma}(n)$ differs much from identity. Interestingly enough, it is possible to express $\boldsymbol{\Sigma}(n)$ as a function of the spectral window W and its derivative,

$$W(n) = \frac{1}{N}\sum_{k=1}^{N} \mathrm{e}^{\mathrm{i}nt_k} \tag{C.14}$$

$$W'(n) = \frac{\mathrm{i}}{N}\sum_{k=1}^{N} t_k\, \mathrm{e}^{\mathrm{i}nt_k}. \tag{C.15}$$

We choose the origin of time such that $y(t) = K\cos nt$ and make the approximation

$$\|\cos n\boldsymbol{t}\| = \|\sin n\boldsymbol{t}\| = \|\cos 2n\boldsymbol{t}\| = \|\sin 2n\boldsymbol{t}\| = \sqrt{\frac{N}{2}} \tag{C.16}$$

$$\|\boldsymbol{t}\cos n\boldsymbol{t}\| = \sqrt{\frac{\|\boldsymbol{t}\| - N\mathcal{R}(W''(2n))}{2}} \approx \frac{\|\boldsymbol{t}\|}{\sqrt{2}} \tag{C.17}$$

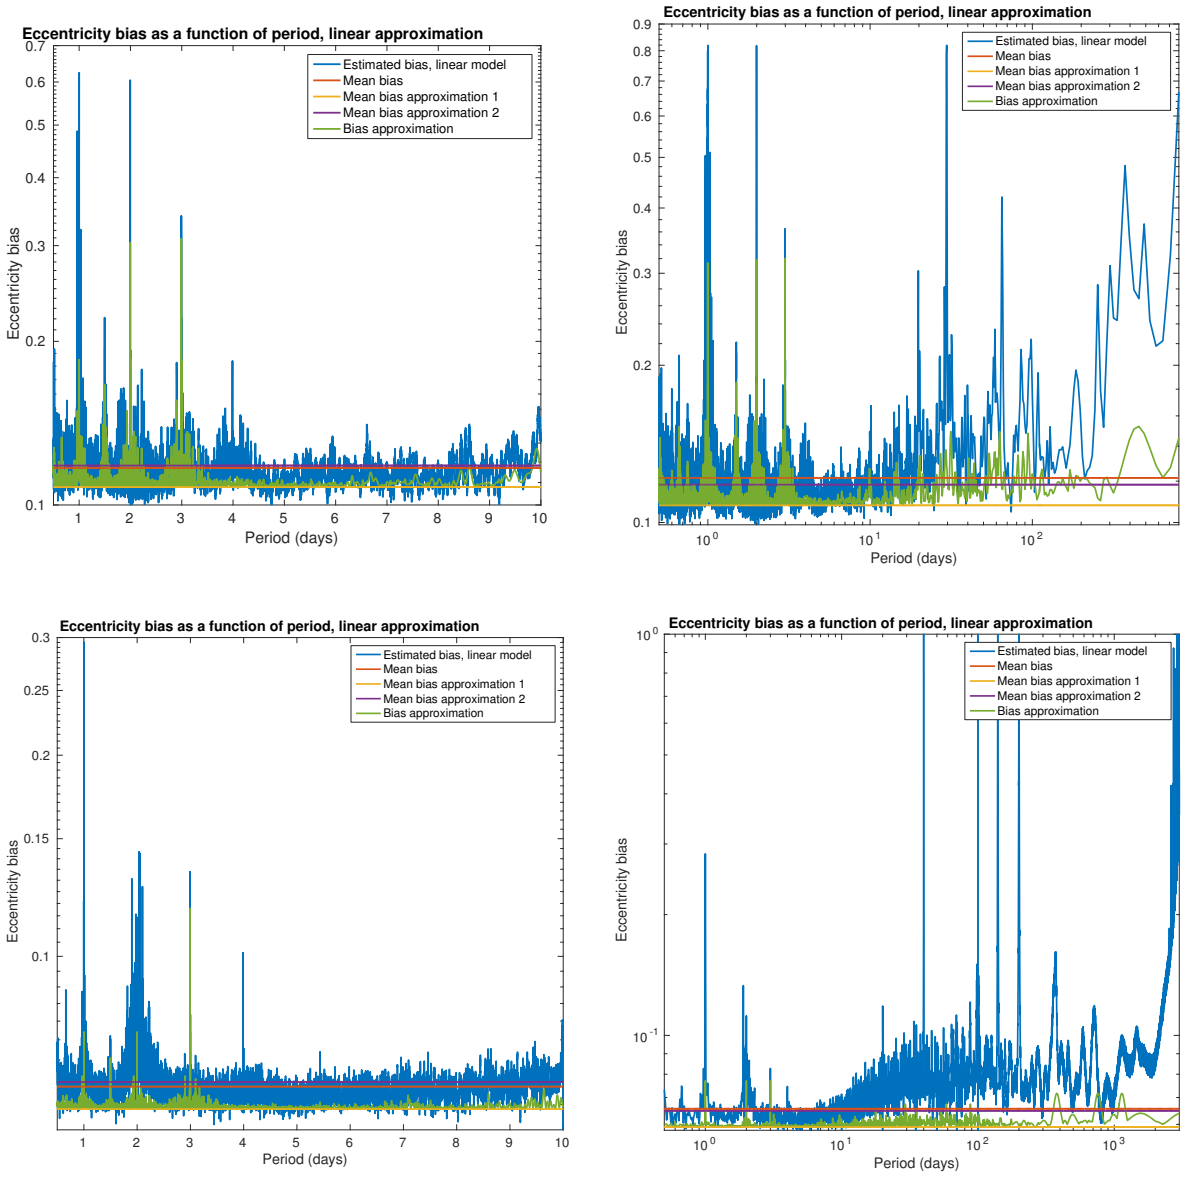

Figure C.1: Effect of the period of the signal on the bias when $e_t = 0$

where $\mathcal{R}(z)$ denotes the real part of complex number z, $\mathbf{\Sigma}$ reads

$$\mathbf{\Sigma}(t) = \begin{pmatrix} 1 & \mathcal{I}(W(2n)) & \mathcal{R}(W(n)+W(3n)) & \mathcal{I}(W(n)+W(3n)) & \frac{\|t\|_1+N\mathcal{I}(W'(2n))}{\sqrt{N}\|t\|} & 2\mathcal{R}(W(n)) \\ \mathcal{I}(W(2n)) & 1 & \mathcal{I}(-W(n)+W(3n)) & \mathcal{R}(W(n)-W(3n)) & -\frac{\sqrt{N}}{\|t\|}\mathcal{R}(W(2n)) & 2\mathcal{I}(W(n)) \\ \mathcal{R}(W(n)+W(3n)) & \mathcal{I}(-W(n)+W(3n)) & 1 & \mathcal{I}(W(4n)) & \frac{\sqrt{N}}{\|t\|}\mathcal{I}(W'(n)+W'(3n)) & 2\mathcal{R}(W(2n)) \\ \mathcal{I}(W(n)+W(3n)) & \mathcal{R}(W(n)-W(3n)) & \mathcal{I}(W(4n)) & 1 & -\frac{\sqrt{N}}{\|t\|}\mathcal{R}(W'(n)+W'(3n)) & 2\mathcal{I}(W(2n)) \\ \frac{\|t\|_1+N\mathcal{I}(W'(2n))}{\sqrt{N}\|t\|} & -\frac{\sqrt{N}}{\|t\|}\mathcal{R}(W(2n)) & \frac{\sqrt{N}}{\|t\|}\mathcal{I}(W'(n)+W'(3n)) & -\frac{\sqrt{N}}{\|t\|}\mathcal{R}(W'(n)+W'(3n)) & 1 & \frac{\sqrt{2}N}{\|t\|}\mathcal{I}(W'(n)) \\ 2\mathcal{R}(W(n)) & 2\mathcal{I}(W(n)) & 2\mathcal{R}(W(2n)) & 2\mathcal{I}(W(2n)) & 2\frac{\sqrt{2}N}{\|t\|}\mathcal{I}(W'(n)) & 1 \end{pmatrix}. \tag{C.18}$$

It means that specifically the value of W and W' at n, $2n$, $3n$, $4n$ are determinant to have a

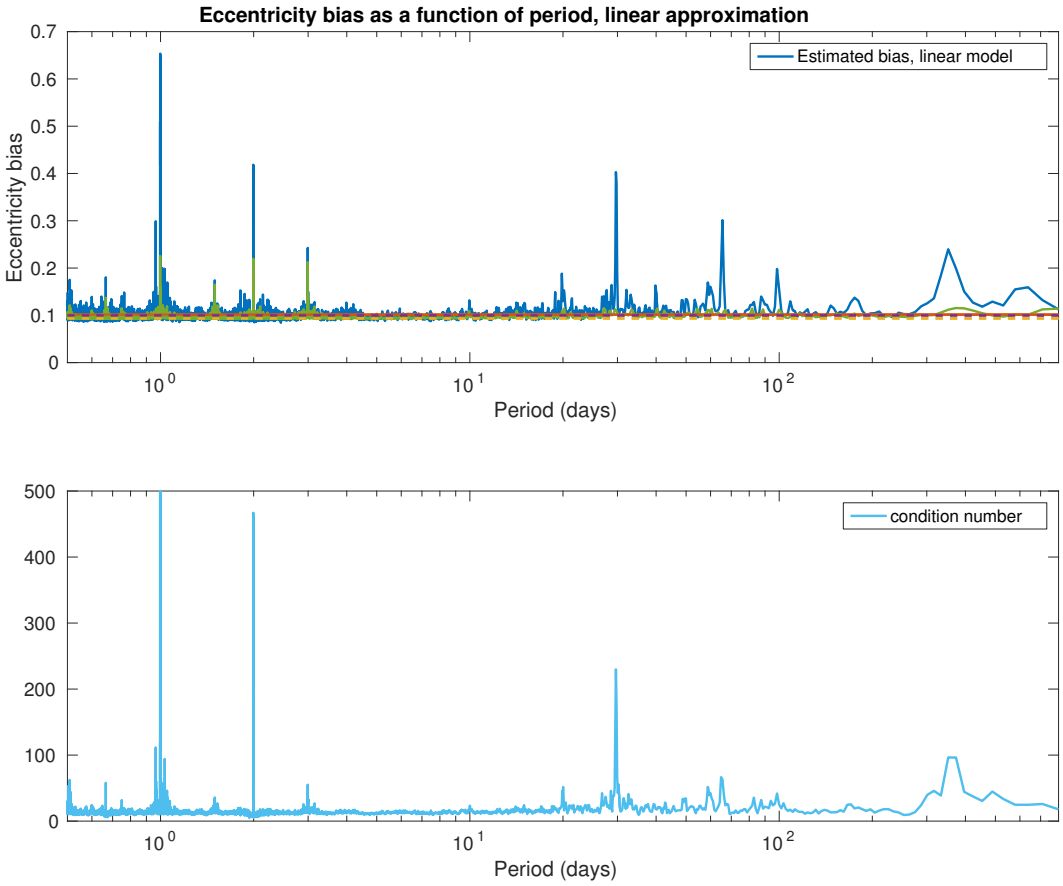

Figure C.2: Top: Eccentricity bias of the linear model as a function of the period. Bottom: condition number of the Fisher information matrix as a function of period. The measurement times are the first 40 measurements of GJ 876 (Correia et al. 2010).

good estimate of eccentricity.

To summarize our results, we plot the value of the bias per period in several situations. At each period P, the bias corresponds to the mean eccentricity obtained when the a planet injected in circular orbit at period p with $K = 3$ m.s^{-1} is fitted with the linearised model (the noise is set to 1 m.s^{-1}). In figures C.1 and C.2, the bias is computed on 1000 simulations with the linearised model. Figure C.1, left and right respectively correspond to a period range of 0.6 to 10 days and 0.6 to 1000 days. The two top figures represent the bias for the fitting of a single planet with the 30 first measurement times of HD10180. On the bottom figures, we take the 100 first measurements and fit a model with 2 planets, plus one planet at the period considered. The 2 planets account for the 4 peaks that reach $e = 1$ on Fig. C.1, bottom right. Which occur when $P = P_i$ or $P = P_i/2$, P_i being the periods of the two planets fitted.

The estimate of the bias amplitude provided by Lucy & Sweeney (1971) is represented in yellow. The mean bias computed from the simulation is represented in red, and our approximation (equation (C.13)) in purple. There is a good agreement between this latter formula and the mean bias computed with the linear model. In both cases, one sees that the bias increases at high period. We interpret this fact as the fact that the condition

number increases as the model $\cos 2\pi \boldsymbol{t}/P$, $\sin 2\pi \boldsymbol{t}/P$ becomes more and more correlated to the constant as P to infinity. To clarify the role of the condition number, on Fig. C.2, we plot the bias on eccentricity and the condition number of the linearised model as a function of period. There is a strong correlation between a high condition number and a high bias.

Appendix D

Optimal approximation by linear subspaces

In the following, $\boldsymbol{f}$ represents the function $\boldsymbol{u}(e,.)\colon \omega \in [0,2\pi] \mapsto \boldsymbol{u}(e,\omega) \in \mathbb{R}^n$, and $I = [0,2\pi]$.

Theorem 8. *Let $\boldsymbol{f}\colon t \in I \subset \mathbb{R} \to \mathbb{R}^n$ be a continuous function and $\mathbf{A}$ be the $n \times n$ positive semi-definite matrix defined as $\mathbf{A} = \int_I \boldsymbol{f}(t)\boldsymbol{f}(t)^T \,\mathrm{d}t$. Let us denote $P_S\colon \mathbb{R}^n \to \mathbb{R}^n$ the projection onto a vector subspace S, and by $\mathcal{S}_d$ the set of subspaces of $\mathbb{R}^n$ of dimension d. Then the subspace solution of*

$$S_d^\star = \arg\min_{S\in\mathcal{S}_d} \int_I \|\boldsymbol{f}(t) - P_S(\boldsymbol{f}(t))\|^2 \,\mathrm{d}t \tag{D.1}$$

is generated by the d eigenvectors of $\mathbf{A}$ with the highest eigenvalues.

Proof. The above problem is equivalent to finding an orthonormal base of $S_d^\star$,

$$\arg\min_{(\boldsymbol{e}_k)_{k=1..d}\in\mathbb{R}^n} \int_I \left\| \boldsymbol{f}(t) - \sum_{k=1}^{d} \langle \boldsymbol{f}(t), \boldsymbol{e}_k\rangle \boldsymbol{e}_k \right\|^2 \mathrm{d}t \tag{D.2}$$

$$\forall k,l \in \{1..d\}, k \neq l, \langle \boldsymbol{e}_k, \boldsymbol{e}_l\rangle = 0 \tag{D.3}$$

$$\forall k \in \{1..d\}, \|\boldsymbol{e}_k\| = 1 \tag{D.4}$$

where $\langle , \rangle$ denotes the usual Euclidian scalar product in $\mathbb{R}^n$. Hereafter, for brevity, conditions (3) and (4) are assumed to hold, but are not written anymore. Developping equation 2, due to the orthonormality of the $\boldsymbol{e}_k$, we obtain

$$\arg\min_{(\boldsymbol{e}_k)_{k=1..d}\in\mathbb{R}^n} \int_I \|\boldsymbol{f}(t)\|^2 \,\mathrm{d}t - \sum_{k=1}^{d} \int_I \langle \boldsymbol{f}(t), \boldsymbol{e}_k\rangle^2 \,\mathrm{d}t. \tag{D.5}$$

Since the norm of $\boldsymbol{f}(t)$ does not depend on the $\boldsymbol{e}_k$, our problem comes down to finding the solution of the maximisation problem

$$\arg\max_{(\boldsymbol{e}_k)_{k=1..d}\in\mathbb{R}^n} \sum_{k=1}^{d} \int_I \langle \boldsymbol{f}(t), \boldsymbol{e}_k\rangle^2 \,\mathrm{d}t. \tag{D.6}$$

Furthermore, for any two vectors $(\boldsymbol{u}, \boldsymbol{v}) \in \mathbb{R}^n$, we have $\langle \boldsymbol{u}, \boldsymbol{v}\rangle^2 = \langle \boldsymbol{v}, \boldsymbol{u}\rangle\langle \boldsymbol{u}, \boldsymbol{v}\rangle = \boldsymbol{v}^T\boldsymbol{u}\boldsymbol{u}^T\boldsymbol{v}$. Thus, moving the constant vectors $\boldsymbol{e}_k$ out of the integral, the problem can be written

$$\arg \max_{(\boldsymbol{e}_k)_{k=1..d}\in\mathbb{R}^n} \sum_{k=1}^{d} {\boldsymbol{e}_k}^T \left(\int_I \boldsymbol{f}(t)\boldsymbol{f}(t)^T \,\mathrm{d}t \right) \boldsymbol{e}_k \tag{D.7}$$

We recognize the matrix $\mathbf{A}$ of the theorem. By construction, this is a positive semi-definite matrix. We proceed by recurrence on d. For $d = 1$, a classical result of linear algebra is that the solution is the eigenvector associated with the maximum eigenvalue of $\mathbf{A}$. Let us suppose that for $d \geq 2$, the solution of $\mathcal{P}_{d-1}$ are the $d-1$ first normed eigenvectors of $\mathbf{A}$, $\boldsymbol{v}_1..\boldsymbol{v}_{d-1}$.

Let us consider a subspace $\mathcal{W}$ of dimension d. We have

$$\begin{aligned}\dim\left(\mathcal{W} \cap (\mathcal{S}_{d-1})^\perp\right) &= N - \dim\left(\mathcal{W} \cap (\mathcal{S}_{d-1})^\perp\right)^\perp \\ &= N - \dim\left(\mathcal{W}^\perp + \mathcal{S}_{d-1}\right)\end{aligned}$$

and

$$\dim\left(\mathcal{W}^\perp + \mathcal{S}_{d-1}\right) \leqslant \dim\left(\mathcal{W}^\perp\right) + \dim\left(\mathcal{S}_{d-1}\right) = N - d + d - 1 = N - 1$$

Therefore

$$\dim\left(\mathcal{W} \cap (\mathcal{S}_{d-1})^\perp\right) \geqslant 1$$

It means that there there exists a vector of $\mathcal{W}$ which is orthogonal to $\mathcal{S}_{d-1}$, which we can choose so that it has unit norm, and that we denote by $\boldsymbol{w}_d$. By the incomplete basis theorem, we there exists $d-1$ vectors $\boldsymbol{w}_k \in \mathcal{W}$ such that $(\boldsymbol{w}_k)_{k=1..d}$ forms a basis, which can be chosen to be orthonormal. Then

$$|\langle \mathbf{A}\boldsymbol{w}_d, \boldsymbol{w}_d\rangle| \leqslant \max_{\boldsymbol{x}\in(\mathcal{S}_{d-1})^\perp, \|\boldsymbol{x}\|=1} |\langle \mathbf{A}\boldsymbol{x}, \boldsymbol{x}\rangle| \leqslant \sigma_d^2 = \langle \mathbf{A}\boldsymbol{v}_d, \boldsymbol{v}_d\rangle \tag{D.8}$$

[1] where σ_d^2 is the d-th eigenvalue of $\mathbf{A}$ in decreasing order and $\boldsymbol{v}_d$ the associated eigenvector. Secondly, by inductive hypothesis we have

$$\sum_{k=1}^{d-1} |\langle \mathbf{A}\boldsymbol{w}_k, \boldsymbol{w}_k\rangle| \leqslant \sum_{k=1}^{d-1} |\langle \mathbf{A}\boldsymbol{v}_k, \boldsymbol{v}_k\rangle| \tag{D.9}$$

Summing respectively left hand sides and right hand sides of equations (D.8) and (D.9) we obtain

$$\sum_{k=1}^{d} |\langle \mathbf{A}\boldsymbol{w}_k, \boldsymbol{w}_k\rangle| \leqslant \sum_{k=1}^{d} |\langle \mathbf{A}\boldsymbol{v}_k, \boldsymbol{v}_k\rangle| \tag{D.10}$$

To conclude $\mathcal{W} = \mathcal{S}_d$ is an acceptable choice and equation (D.10) ensures that it maximizes equation (D.6). This achieves the induction step and the proof.

□

[1] Note that the absolute values can be dropped since $\mathbf{A}$ is positive.

Appendix E

Residual analysis

In this section we will compute the law followed by the residuals of a linear least-square fit. Let us suppose that we have a model

$$\boldsymbol{y} = \mathbf{A}\boldsymbol{x} + \boldsymbol{\epsilon}, \;\; \boldsymbol{\epsilon} \sim G(0, \mathbf{V})$$

where $\boldsymbol{y}$ is a vector of N observations, modelled as a linear combination of the column of the $N \times p$ matrix A, and $\boldsymbol{\epsilon}$ is a Gaussian noise of covariance matrix $\mathbf{V} =: \mathbf{W}^{-1}$. Assuming $\mathbf{V}$ and $\mathbf{A}$ are known, the least square estimate of $\boldsymbol{y}$ is $\widehat{\boldsymbol{y}} = \mathbf{A}(\mathbf{A}^T\mathbf{W}\mathbf{A})^{-1}\mathbf{A}^T\mathbf{W}\boldsymbol{y}$ Therefore

$$\begin{aligned}
\mathbf{W}^{1/2}(\boldsymbol{y} - \widehat{\boldsymbol{y}}) &= \mathbf{W}^{1/2}\left(\mathbf{A}\boldsymbol{x} + \boldsymbol{\epsilon} - \mathbf{A}(\mathbf{A}^T\mathbf{W}\mathbf{A})^{-1}\mathbf{A}^T\mathbf{W}(\mathbf{A}\boldsymbol{x} + \boldsymbol{\epsilon})\right) \\
&= \mathbf{W}^{1/2}(\mathbf{I}_N - \mathbf{B})\boldsymbol{\epsilon} \\
&=: \boldsymbol{r}_W
\end{aligned}$$

where $\mathbf{I}_N$ is the identity matrix of size N and $\mathbf{B} := \mathbf{A}(\mathbf{A}^T\mathbf{W}\mathbf{A})^{-1}\mathbf{A}^T\mathbf{W}$. The quantity $\boldsymbol{r}_W$, being a product of a matrix ($\mathbf{W}^{1/2}(\mathbf{I}_n - \mathbf{B})$) with a Gaussian random variable of covariance $\mathbf{V}$ has a covariance $\mathbf{U}$

$$\begin{aligned}
\mathbf{U} &= \mathbf{W}^{1/2}(\mathbf{I}_N - \mathbf{B}^T)\mathbf{V}(\mathbf{I}_N - \mathbf{B})\mathbf{W}^{1/2} \\
&= \mathbf{W}^{1/2}(\mathbf{V} - \mathbf{B}\mathbf{V} - \mathbf{V}\mathbf{B}^T + \mathbf{B}\mathbf{V}\mathbf{B}^T)\mathbf{W}^{1/2} \\
\text{since } \; \mathbf{W}^{1/2}\mathbf{V}\mathbf{B}^T &= \mathbf{W}^{-1/2}\mathbf{B}^T\mathbf{W}^{1/2} = \mathbf{W}^{1/2}\mathbf{B}\mathbf{V}\mathbf{B}^T\mathbf{W}^{1/2}, \\
\mathbf{U} &= \mathbf{I}_N - \mathbf{W}^{1/2}\mathbf{B}^T\mathbf{W}^{-1/2} \\
&= \mathbf{I}_N - \mathbf{C}(\mathbf{C}^T\mathbf{C})^{-1}\mathbf{C}^T
\end{aligned}$$

where $\mathbf{C} = \mathbf{W}^{1/2}\mathbf{A}$. This notation is convenient because it shows clearly that $\mathbf{P} = \mathbf{C}(\mathbf{C}^T\mathbf{C})^{-1}\mathbf{C}^T$ is a projection matrix on the space generated by the columns of $\mathbf{C}$. Finally

$$\mathbf{U} = \mathbf{I}_N - \mathbf{P} \tag{E.1}$$

Is a projection on the space orthogonal to the one generated by $\mathbf{C}$ columns. Therefore, there exists an orthonormal matrix $\mathbf{Q}$ such that

$$\mathbf{Q}^T\mathbf{U}\mathbf{Q} = \mathbf{J}_p$$

where $\mathbf{J}_p$ is a diagonal matrix whose first p elements are zero and the others are equal to one. Finally, let us remark that $\boldsymbol{r}_{QW} := \mathbf{Q}^T \boldsymbol{r}_W$ has a covariance matrix $\mathbf{Q}^T \mathbf{U} \mathbf{Q} = \mathbf{J}_p$, which shows the claim of section 3.4.2, $\mathbf{Q}^T \mathbf{W}^{1/2} (\boldsymbol{y} - \widehat{\boldsymbol{y}})$ has p 0 components and the others are Gaussian variables of mean 0 and variance 1.

Let us finally note that the covariance matrix $\mathbf{U} = \mathbf{I}_N - \mathbf{P}$ of $\boldsymbol{r}_W = \mathbf{W}^{1/2} (\boldsymbol{y} - \widehat{\boldsymbol{y}})$ will be close to identity if there are many more observations than parameters. This explains why the weighted residuals $\boldsymbol{r}_W$ almost behaves like independent Gaussian variables and, for instance, why plotting $\boldsymbol{r}_W(t_i) - \boldsymbol{r}_W(t_j)$ as a function of $t_i - t_j$ gives hints on the correlations.

Appendix F

Details on the frequentist methodology

In this section, we outline the calculation of the confidence intervals for eccentricity. Such an interval is constructed as a set of eccentricities that are not rejected by a hypothesis test. We choose the likelihood ratio test:

$$e \text{ is rejected if } \quad R := \frac{\max\limits_{\boldsymbol{\theta}\in\Theta_e} f(\boldsymbol{y}|\boldsymbol{\theta})}{\max\limits_{\boldsymbol{\theta}\in\Theta} f(\boldsymbol{y}|\boldsymbol{\theta})} \leqslant \beta$$

where $\boldsymbol{y}$ denotes the actual observations, $f(\boldsymbol{y}|\boldsymbol{\theta})$ denotes the likelihood, Θ_e is the set of parameters with eccentricity e, and β is a constant which will be made explicit later. Our aim is to compute the distribution of R under the assumption that the random variable giving the observations is $\boldsymbol{Y} = \boldsymbol{y}_t + \boldsymbol{\epsilon}$, $\boldsymbol{\epsilon}$ being a Gaussian noise. We further assume the noise is independent and identically distributed, the condition translates to

$$e \text{ is rejected if } \quad D := \|\boldsymbol{y} - \boldsymbol{y}(\boldsymbol{\theta}_e)\|^2 - \|\boldsymbol{y} - \boldsymbol{y}^\star\|^2 \geqslant -2\sigma^2 \ln\beta \tag{F.1}$$

where $\boldsymbol{\theta}_e = \arg\min\limits_{\boldsymbol{\theta}\in\Theta_e} \|\boldsymbol{y} - \boldsymbol{y}(\boldsymbol{\theta})\|^2$, σ^2 is the variance of the observations and $\boldsymbol{y}^\star$ is the global minimum. We will now compute the law followed by D, so that we can select a β that corresponds to a false alarm probability. Since D is defined implicitly, the calculation of its distribution is difficult. We will make two simplifying assumptions that will allow us to obtain an analytical expression. The expression will then be tested on real cases through numerical simulations.

Let us first consider the linear approximation $\boldsymbol{y} = \mathbf{M}\boldsymbol{x}_t + \boldsymbol{\epsilon}$ where $\mathbf{M}$ is defined as in (C.3) and (C.4). We further suppose that the columns of $\mathbf{M}$ are orthonormal. Since the columns are originally of the form $\cos n\boldsymbol{t}, \sin n\boldsymbol{t}, \cos 2n\boldsymbol{t}, \sin 2n\boldsymbol{t}$, they must be multiplied by $\sqrt{2/N}$ and the amplitude of the signal is no K_t but $K_t\sqrt{N/2}$. We look for the solution $\hat{\boldsymbol{\theta}}_e$ defined as

$$\hat{\boldsymbol{\theta}}_e = \arg\min_{\boldsymbol{x}\in\mathbb{R}^p} \|\boldsymbol{y} - \mathbf{M}\boldsymbol{x}\| \quad \text{subject to} \quad \sqrt{\frac{x_3^2 + x_4^2}{x_1^2 + x_2^2}} = e. \tag{F.2}$$

Thanks to the Lagrange multipliers theorem, $\hat{\boldsymbol{\theta}}_e$ satisfies the conditions

$$\frac{\partial L}{\partial \boldsymbol{x}} = \mathbf{0}, \quad \frac{\partial L}{\partial \lambda} = 0, \quad \text{where} \tag{F.3}$$

$$L(\boldsymbol{x}, \lambda) = \frac{1}{2}\|\boldsymbol{y} - \mathbf{M}\boldsymbol{x}\|^2 + \frac{\lambda}{2}\boldsymbol{x}^T\mathbf{E}\boldsymbol{x} \tag{F.4}$$

with

$$\mathbf{E} = \operatorname{diag}\left(-e^2, -e^2, 1, 1, 0, \ldots, 0\right). \tag{F.5}$$

The condition $\partial L/\partial \boldsymbol{x} = \mathbf{0}$ leads to

$$\left(\mathbf{M}^T\mathbf{M} + \lambda\mathbf{E}\right)\boldsymbol{x} = \mathbf{M}^T\boldsymbol{y}. \tag{F.6}$$

Since the columns of $\mathbf{M}$ are orthonormal, $\mathbf{M}^T\mathbf{M}$ is the identity, thus

$$x_1 = \frac{u_1}{1-\lambda e^2}, \quad x_2 = \frac{u_2}{1-\lambda e^2}, \quad x_3 = \frac{u_3}{1+\lambda}, \quad x_4 = \frac{u_4}{1+\lambda}, \tag{F.7}$$

and $x_j = u_j,\ \forall j \geq 5$, where we have defined $u_i = \boldsymbol{M}_i{}^T\boldsymbol{y}$, $\boldsymbol{M}_i$ being the i-th column of $\mathbf{M}$. The first four components of $\boldsymbol{x}$ are also constrained by $\partial L/\partial\lambda = 0$. Let $U = u_1^2 + u_2^2$ and $V = u_3^2 + u_4^2$. We get

$$\frac{-e^2}{(1-\lambda e^2)^2}U + \frac{1}{(1+\lambda)^2}V = 0, \tag{F.8}$$

or, equivalently,

$$e^2(e^2V - U)\lambda^2 - 2e^2(V+U)\lambda + V - e^2U = 0, \tag{F.9}$$

whose solutions are

$$\lambda_\pm = \frac{e^2(U+V) \pm e(1+e^2)\sqrt{UV}}{e^2(e^2V - U)}. \tag{F.10}$$

For the solution $\hat{\boldsymbol{\theta}}_e$ to actually be a minimum of L, all its eigenvalues must be positive, i.e., λ must verify $-1 < \lambda < 1/e^2$. Only λ_- fulfils this criterion, thus

$$\lambda = \frac{e^2(U+V) - e(1+e^2)\sqrt{UV}}{e^2(e^2V - U)}, \tag{F.11}$$

and

$$x_1 = \frac{1+e_0^2}{1+e^2}u_1, \quad x_2 = \frac{1+e_0^2}{1+e^2}u_2, \quad x_3 = \frac{e^2}{e_0^2}\frac{1+e_0^2}{1+e^2}u_3, \quad x_4 = \frac{e^2}{e_0^2}\frac{1+e_0^2}{1+e^2}u_4, \tag{F.12}$$

with $e_0^4 = e^2V/U$. After a few calculation, we show that

$$D = \sum_{k=1}^{4}(u_k - x_k)^2 = \frac{\left(e\sqrt{u_1^2+u_2^2} - \sqrt{u_3^2+u_4^2}\right)^2}{1+e^2}. \tag{F.13}$$

Let $x = e\sqrt{u_1^2+u_2^2}/K_t$ and $y = \sqrt{u_3^2+u_4^2}/K_t$. These two random variables follow Rice distributions with parameters

$$\nu_x = e\sqrt{\frac{N}{2}}, \quad \sigma_x = \frac{e\sigma}{K_t}, \quad \nu_y = e\sqrt{\frac{N}{2}}, \quad \sigma_y = \frac{\sigma}{K_t}. \tag{F.14}$$

An expansion of the product term shows that D behaves approximately as a weighted sum of variables following a χ^2 distribution. We can then use the Welch-Satterthwaite approximation (Satterthwaite 1946; Welch 1947): D approximately follows a χ^2 distribution whose

number of degrees of freedom ν is given by $\mathbb{E}\{D\}$. In the following, we denote by $S' = \frac{K_t}{\sigma}\sqrt{\frac{N}{2}}$ the signal to noise ratio. The expected value of D is

$$\begin{aligned}
\mathbb{E}\{D\} &= \frac{K_t^2}{1+e^2}\int_0^\infty\int_0^\infty (x-y)^2 f(x|\nu_x,\sigma_x) f(y|\nu_y,\sigma_y)\,\mathrm{d}x\mathrm{d}y, \\
&= \frac{K_t^2}{1+e^2}\left[2\sigma_x^2+\nu_x^2+2\sigma_y^2+\nu_y^2-\pi\sigma_x\sigma_y L_{\frac{1}{2}}\left(-\frac{\nu_x^2}{2\sigma_x^2}\right)L_{\frac{1}{2}}\left(-\frac{\nu_y^2}{2\sigma_y^2}\right)\right] \\
&= \frac{K_t^2}{1+e^2}\left[2\frac{\sigma^2}{K_t^2}(1+e^2)+Ne^2-e\pi\frac{\sigma^2}{K_t^2}L_{\frac{1}{2}}\left(-\frac{S'^2}{2}\right)L_{\frac{1}{2}}\left(-\frac{e^2S'^2}{2}\right)\right].
\end{aligned} \tag{F.15}$$

With $\nu = \mathbb{E}\{D\}/\sigma^2$, we get

$$\nu = 2 + 2S'^2\frac{e^2}{1+e^2} - \frac{\pi e}{1+e^2}L_{\frac{1}{2}}\left(-\frac{S'^2}{2}\right)L_{\frac{1}{2}}\left(-\frac{e^2S'^2}{2}\right). \tag{F.16}$$

To obtain a confidence level α, then we need to take $-2\ln\beta = F^{-1}_{\chi^2_\nu}(1-\alpha)$ where $F^{-1}_{\chi^2_\nu}$ is the inverse cumulative distribution function of a χ^2 distribution with ν degrees of freedom. Conversely, it is possible to convert a measured D to a probability simply by computing $\alpha_e = 1 - F_{\chi^2_\nu}(D)$. The hypothesis $e_t = e$ is rejected if α_e is below a certain threshold.

This formula was tested numerically. It is in very good agreements with the simulations as soon as S' is above ≈ 20. As it decreases, the average of estimated eccentricity increases (which is exactly saying that the bias increases) therefore the approximation of low eccentricities does not hold any more. The value of S' can be evaluated keeping in mind that when the linearised model at $e = 0$ is poorly conditioned, (matrix $\mathbf{M}$, as defined equations (C.3) and (C.4)), then the uncertainty on k and h is higher than given by the simple formula (C.11) and the S' analytical approximation is inoperative.

Appendix G

Hara, Boué, Laskar & Correia 2017

This section contains the last version of Hara et al. (2017). We refer to section "G.i" as the section i of the article for $i = 1..6$. Sections G.7 to G.10 refer to the appendices of the article (A to D).

Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **464,** 1220–1246 (2017)
Advance Access publication 2016 September 8

doi:10.1093/mnras/stw2261

Radial velocity data analysis with compressed sensing techniques

Nathan C. Hara,[1]★ G. Boué,[1] J. Laskar[1] and A. C. M. Correia[1,2]

[1]*ASD/IMCCE, CNRS-UMR8028, Observatoire de Paris, PSL, UPMC, 77 Avenue Denfert-Rochereau, F-75014 Paris, France*
[2]*CIDMA, Departamento de Física, Universidade de Aveiro, Campus de Santiago, P-3810-193 Aveiro, Portugal*

Accepted 2016 September 7. Received 2016 September 5; in original form 2016 April 2

ABSTRACT

We present a novel approach for analysing radial velocity data that combines two features: all the planets are searched at once and the algorithm is fast. This is achieved by utilizing compressed sensing techniques, which are modified to be compatible with the Gaussian process framework. The resulting tool can be used like a Lomb–Scargle periodogram and has the same aspect but with much fewer peaks due to aliasing. The method is applied to five systems with published radial velocity data sets: HD 69830, HD 10180, 55 Cnc, GJ 876 and a simulated very active star. The results are fully compatible with previous analysis, though obtained more straightforwardly. We further show that 55 Cnc e and f could have been respectively detected and suspected in early measurements from the Lick Observatory and Hobby–Eberly Telescope available in 2004, and that frequencies due to dynamical interactions in GJ 876 can be seen.

Key words: methods: data analysis – techniques: radial velocities – planets and satellites: detection.

1 INTRODUCTION

1.1 Overview

Determining the content of radial velocity data is a challenging task. There might be several companions to the star, unpredictable instrumental effects as well as astrophysical jitter. Fitting separately the different features of the model might distort the residual and prevent one from finding small planets, as pointed out for instance by Anglada-Escudé, López-Morales & Chambers (2010) and Tuomi (2012). There might even be the cases where, due to aliasing and noise, the tallest peak of the periodogram is a spurious one while being statistically significant. To overcome those issues, recent approaches privilege the fitting of the whole model at once. In those cases, the usual framework is the maximization of an a posteriori probability distribution. In order to avoid being trapped in a sub-optimal solution, random searches such as Monte Carlo Markov chain methods or genetic algorithm are used (e.g. Gregory 2011; Ségransan et al. 2011). The goal of this paper is to suggest an alternative method using convex optimization, therefore offering a unique minimum and faster algorithms.

To do so, we will not try to find directly the orbital parameters of the planets but to unveil the true spectrum of the underlying continuous signal, which is equivalent. The power spectrum is often estimated with a Lomb–Scargle periodogram (Lomb 1976; Scargle 1982) or generalizations (Ferraz-Mello 1981; Cumming, Marcy & Butler 1999; Zechmeister & Kürster 2009). However, as said above, the estimation of the power spectrum with one frequency at a time has severe drawbacks. To improve the estimate, we introduce a priori information: the representation of exoplanetary signal in the Fourier domain is sparse. In other words, the number of sine functions needed to represent the signal is small compared to the number of observations. The Keplerian models are not the only ones to verify this assumption, stable planetary systems are quasi-periodic as well (e.g. Laskar 1993). By doing so, the periodogram can be efficiently cleaned (see Figs 1–5).

The field of signal processing devoted to the study of sparse signals is often referred to as 'compressed sensing' or 'compressive sampling' (Candès, Romberg & Tao 2006b; Donoho 2006) – though it is sometimes restricted to sampling strategies based on sparsity of the signal. The related methods show very good performances and are backed up by solid theoretical results. For instance, compressed sensing techniques allow one to recover exactly a spectrum while sampling it at a much lower rate than the Nyquist frequency (Mishali, Eldar & Tropp 2008; Tropp et al. 2010). Its use was advocated to improve the scientific data transmission in space-based astronomy (Bobin, Starck & Ottensamer 2008). Sparse recovery techniques are also used in image processing (e.g. Starck, Elad & Donoho 2005).

It seems relevant to add to that list a few techniques developed by astronomers to retrieve harmonics in a signal. In the next section, we show that even though the term 'sparsity' is not explicitly used (except in Bourguignon, Carfantan & Böhm 2007), some of the existing techniques have an equivalent in the compressed sensing literature. After those remarks on our framework, the paper is organized as follows: in Section 2, the theoretical background and the associated algorithms are presented. Section 3 presents in detail the procedure we developed for analysing radial velocity data. This one is applied in Section 4 to simulated observations and four real radial velocity data sets: HD 69830, HD 10180, 55 Cnc and GJ 876 and to a simulated very active star. The performance of

★ E-mail: nathan.hara@obspm.fr

the method is discussed in Section 5, and conclusions are drawn in Section 6.

1.2 Previous work

The goal of this paper is to devise a method to efficiently analyse radial velocity data. As it builds upon the retrieval of harmonics, the discussion will focus on spectral synthesis of unevenly sampled data (see Kay & Marple 1981; Schwarzenberg-Czerny 1998; Babu & Stoica 2010, for surveys).

First, let us consider the methods that are efficient to spot one harmonic at a time. The first statistical analysis is given by Schuster (1898). However, the statistical properties of Schuster's periodogram only hold when the measurements are equispaced in time. When this is not the case, one can use Lomb–Scargle periodogram (Lomb 1976; Scargle 1982) or its generalization consisting in adding a constant to the model (Ferraz-Mello 1981; Cumming et al. 1999; Reegen 2007; Zechmeister & Kürster 2009). More recently, Mortier et al. (2015) derived a Bayesian periodogram associated with the maximum of an a posteriori distribution. Also, Cumming (2004) and O'Toole et al. (2009) define the Keplerian periodogram, which measures the χ^2 of residuals after the fit of a Keplerian curve. One can remark that 'Keplerian' vectors defined by P, e, ω and M_0 form a family of vectors in which the sparsity of exoplanetary signals is enhanced.

These methods can be applied iteratively to retrieve several harmonics. In the context of radial velocity data processing, one searches for the peak of maximum power, then the corresponding signal is subtracted and the search is performed again. This procedure is very close to CLEAN (Roberts, Lehar & Dreher 1987), which relies on the same principle of maximum correlation and subtraction. One of the first general algorithm exploiting sparsity of a signal in a given set of vectors (Matching Pursuit, Mallat & Zhang 1993) relies on the same iterative process. This method was formerly known as forward stepwise regression (e.g. Bellmann 1975). To limit the effects of error propagation in the residuals, one can use the orthogonal matching pursuit algorithm (Pati, Rezaiifar & Krishnaprasad 1993; Tropp & Gilbert 2007). In that case, when a harmonic is found to have maximum correlation with the residuals, it is not directly subtracted. The next residual is computed as the original signal minus the fit of all the frequencies found so far. The CLEANest algorithm (Foster 1995) and frequency map analysis (Laskar 1988; Laskar, Froeschlé & Celletti 1992; Laskar 1993; Laskar 2003), though developed earlier, are particular cases of this algorithm. To analyse radial velocity data, Baluev (2009) and Anglada-Escudé & Tuomi (2012) introduce what they call respectively the 'residual periodogram' and the 'recursive periodogram', which can be seen as pushing that logic one step further. The principle is to re-fit at each trial frequency the previous Keplerian signals plus a sine at the considered frequency.

Besides the matching pursuit procedures, there are two other popular algorithms in the compressed sensing literature: convex relaxations (e.g. Tibshirani 1994; Chen, Donoho & Saunders 1998; Starck et al. 2005) and iteratively re-weighted least squares (IRWLS; e.g. Gorodnitsky & Rao 1997; Candès, Romberg & Tao 2006a; Donoho 2006; Daubechies et al. 2010). In the context of astronomy, Bourguignon et al. (2007) implement a convex relaxation method using ℓ_1 norm weighting (see equation 2) to find periodicity in unevenly sampled signals, and Babu et al. (2010) present an IRWLS algorithm named IAA to analyse radial velocity.

The methods presented above are apparently very different, yet they can be viewed as a way to bypass the brute force minimization of

$$\arg\min_{K,\omega,\phi} \sum_{i=1}^{m} \left(\boldsymbol{y}(t_i) - \sum_{j=1}^{k} K_j \cos(\omega_j t_i + \phi_j) \right)^2 , \quad (1)$$

where $\boldsymbol{y}(\boldsymbol{t})$ is a vector made of m measurements, and $\boldsymbol{x}^\star = \arg\min f(\boldsymbol{x})$ denotes the element such that $f(\boldsymbol{x}^\star) = \min f(\boldsymbol{x})$ for a function f. This problem is very similar to 'best k-term approximation', and its link to compressed sensing has been studied in Cohen, Dahmen & Devore (2009) in the noise-free case. Solving that problem is suggested by Baluev (2013b) under the name of 'multifrequency periodograms'. However, finding that minimum by discretizing the values of $(K_j, \omega_j, \phi_j)_{j=1,\dots,k}$ depends exponentially on the number of parameters, and the multifrequency periodograms could hardly handle more than three or four sines with conventional methods. However, with parallel programming on GPUs one can handle up to $\approx$25 frequencies depending on the number of measurements (Baluev 2013a). Jenkins et al. (2014) explicitly mention the above problem and suggest a tree-like algorithm to explore the frequency space. They analyse GJ 876 with their procedure and find six significant harmonics, which we confirm in Section 4.5.2.

Let us mention that searching for a few sources of periodicity in a signal is not always done with the Fourier space. When the shape of the repeating signal or the noise structure is not well known, other tests might be more robust. A large part of those methods consists in computing the autocorrelation function or folding the data at a certain period and look for correlation. See Engelbrecht (2013) for a survey or Zucker (2015, 2016) in the context of radial velocity measurements. Finally, we point out that the use of the sparsity of the signal is not specific to compressed sensing. The number of planets in a model is often selected via likelihood ratio tests. A model with an additional planet must yield a significant improvement of the evidence. In general, the model with $k+1$ planets $\mathcal{M}_{k+1}$ is selected over a model with k planet if $\Pr\{\boldsymbol{y}(\boldsymbol{t})|\mathcal{M}_{k+1}\}/\Pr\{\boldsymbol{y}(\boldsymbol{t})|\mathcal{M}_k\}$ is greater than 150 (see Tuomi et al. 2014), $\boldsymbol{y}(\boldsymbol{t})$ being the observations. Indeed, adding more parameters to the model automatically decreases the χ^2 of the residuals. Putting a minimum improvement of the χ^2 acts against overly complicated models.

The discussion above points that searching planets one after another is already in the compressed sensing paradigm: this iterative procedure is close to the orthogonal matching pursuit algorithm. Donoho, Elad & Temlyakov (2006) show that for a wide range of signals, this algorithm is outperformed by ℓ_1 relaxation methods. Does this claim still apply to radial velocity signals ? In this paper, this question is not treated in full generality, but we show the interest of ℓ_1 relaxation on several examples. To address that question more directly, it is shown in Appendix C that in some cases, the tallest peak of the periodogram is spurious but ℓ_1 minimization prevents one from being misled.

2 METHODS

2.1 Minimization problem

Techniques based on sparsity are thought to enforce the 'Occam's razor' principle: the simplest explanation is the best. To apply that principle, we must have an idea of 'how' the signal is simple. In the compressed sensing framework (or compressive sampling), this is done by selecting a set of vectors $\mathcal{A} = (\boldsymbol{a}_j(\boldsymbol{t}))_{j\in I}$ such that the signal to be analysed $\boldsymbol{y}(\boldsymbol{t})$ is represented by a linear combination of

a few elements of $\mathcal{A}$. Such a set is often called the 'dictionary' and can be finite or not (the set of indices I can be finite or infinite). It is here made of vectors $\boldsymbol{a}_{-\omega}(\boldsymbol{t}) = \mathrm{e}^{-\mathrm{i}\omega t}$ and $\boldsymbol{a}_{\omega}(\boldsymbol{t}) = \mathrm{e}^{\mathrm{i}\omega t}$, where $\boldsymbol{t}$ is the array of measurement times.

Before going into the details, let us define some quantities.

(i) $\boldsymbol{y}(\boldsymbol{t})$ denotes the vector of observations at times $t = t_1, \ldots, t_m$, $\boldsymbol{y}(\boldsymbol{t}) \in \mathbb{R}^m$ for radial velocity data sets.

(ii) The ℓ_p norm of a complex or real vector x with n components is defined as

$$\|\boldsymbol{x}\|_{\ell_p} := \left(\sum_{k=1}^{n} |x_k|^p \right)^{1/p} \tag{2}$$

for $p > 0$. In particular, $\|\boldsymbol{x}\|_{\ell_1}$ is the sum of absolute values of the vector components and $\|\boldsymbol{x}\|_{\ell_2} = \sqrt{\sum_{k=1}^{n} |\boldsymbol{x}_k|^2}$ is the usual Euclidian norm. When $p = 0$, $\|\boldsymbol{x}\|_{\ell_0}$ is the number of non-zero components of $\boldsymbol{x}$.

(iii) For a function f defined on a set E, $\arg\min_{\boldsymbol{x}\in E} f(\boldsymbol{x})$ is the element for which the minimum is attained, that is $\boldsymbol{x}^\star$ of E such that $f(\boldsymbol{x}^\star) = \min_{\boldsymbol{x}\in E} f(\boldsymbol{x})$. We denote by the superscript $\star$ the solution of the minimization problem under consideration. In all the cases considered here except equations (1) and (3), the minimum is attained as we consider convex functions on convex sets.

Let us consider combinations of S elements of the dictionary $(\boldsymbol{a}_j(\boldsymbol{t}))_{j=1,\ldots,S}$ and their corresponding amplitudes x_j. To enhance the sparsity of the representation, one can think of solving

$$\underset{\substack{\boldsymbol{a}_{j(t)}\in\mathcal{A} \\ S\in\mathbb{C}}}{\arg\min} \quad S \quad \text{s.t.} \quad \left\| \sum_{j=1}^{S} x_j \boldsymbol{a}_j(\boldsymbol{t}) - \boldsymbol{y}(\boldsymbol{t}) \right\|_{\ell_2} \leqslant \epsilon \tag{3}$$

that is finding the smallest number of elements of $\mathcal{A}$ required to approximate $\boldsymbol{y}(\boldsymbol{t})$ with a certain tolerance ϵ. This one is a priori a combinatorial problem which seems unsolvable if $\mathcal{A}$ is infinite or of an exponential complexity if the dictionary is finite. In the latter case, $\mathcal{A}$ can be viewed as an $m \times n$ matrix $\mathbf{A}$. In that case, on can re-write equation (3) like

$$\boldsymbol{x}^\star = \arg\min_{\boldsymbol{x}\in\mathbb{C}^n} \|\boldsymbol{x}\|_{\ell_0} \quad \text{s.t.} \quad \|\mathbf{A}\boldsymbol{x} - \boldsymbol{y}(\boldsymbol{t})\|_{\ell_2} \leqslant \epsilon. \tag{4}$$

This problem is in general combinatorial (Ge, Jiang & Ye 2011), therefore computationally intractable. Fortunately, when replacing the ℓ_0 norm by the ℓ_1 norm,

$$\boldsymbol{x}^\star = \arg\min_{\boldsymbol{x}\in\mathbb{C}^n} \|\boldsymbol{x}\|_{\ell_1} \quad \text{s.t.} \quad \|\mathbf{A}\boldsymbol{x} - \boldsymbol{y}(\boldsymbol{t})\|_{\ell_2} \leqslant \epsilon, \tag{5}$$

the problem becomes convex and still enhances sparsity efficiently. In the signal processing literature, this problem is referred to as basis pursuit denoising (Chen et al. 1998), and is sometimes denoted by BP_ϵ. At this point, one can ask what is lost by considering equation (5) instead of equation (3). Let us cite a few results – among many: when $\boldsymbol{y}(\boldsymbol{t})$ is noise free, Donoho (2006) shows that under certain hypotheses the solution to equation (5) is equal to the solution of equation (3); more generally, denoting by $\boldsymbol{y}_t = \mathbf{A}\boldsymbol{x}_t$ the true signal, such that $\boldsymbol{y} = \boldsymbol{y}_t + \boldsymbol{e}$, $\boldsymbol{e}$ being the error, there is a theoretical bound on $\|\mathbf{A}\boldsymbol{x}^\star - \boldsymbol{y}_t\|_{\ell_2}$ (Candès et al. 2006b). One can also obtain constraints on $\|\boldsymbol{x}^\star - \boldsymbol{x}_t\|_{\ell_2}$ or conditions to have $\mathrm{supp}(\boldsymbol{x}^\star) \subset \mathrm{supp}(\boldsymbol{x}_t)$, where $\mathrm{supp}(\boldsymbol{x})$ is the set of indices with x being non-zero (e.g. Donoho et al. 2006). In summary, there are results guaranteeing the performance for denoising, compression and also for inverse problems, the search for planets being a particular case of the latter.

These results apply to a finite dictionary $\mathcal{A}$, but the periods of the planets could be anywhere: $\mathcal{A}$ is infinite for our purposes. We will eventually go back to solving a modified version of the discrete problem (5) and smooth its solution with a moving average. Beforehand, we will present in the next section what seems to be the most relevant theoretical background for our studies, 'atomic norm minimization', in particular used in 'super-resolution theory' (Candès & Fernandez-Granda 2014). This one will give guidelines to improve our procedure.

2.2 Atomic norm minimization

If $\mathcal{A}$ is infinite, the ℓ_1 norm cannot be used straightforwardly. Chandrasekaran et al. (2010) suggest to use an 'atomic norm' that extends equation (5) to infinite dictionaries. Practical methods to solve the new minimization problem are designed in Candès & Fernandez-Granda (2013) and Tang et al. (2013b). The atomic norm $\|\boldsymbol{y}\|_{\mathcal{A}}$, of $\boldsymbol{y} \in \mathbb{R}^m$ or $\mathbb{C}^m$ defined for a dictionary $\mathcal{A}$, is the smallest ℓ_1 norm of a combination of vectors of the dictionary reproducing y:

$$\|\boldsymbol{y}\|_{\mathcal{A}} = \inf \left\{ \sum_j |x_j|, \boldsymbol{y} = \sum_j x_j \boldsymbol{a}_j(\boldsymbol{t}) \right\}. \tag{6}$$

If the observations were not noisy, computing the atomic norm of $\boldsymbol{y}$ would be sufficient. As this is obviously not the case, the following problem is considered,

$$\boldsymbol{u}^\star = \underset{\boldsymbol{u}\in\mathbb{C}^m}{\arg\min} \quad \|\boldsymbol{u} - \boldsymbol{y}(\boldsymbol{t})\|_{\ell_2}^2 + \lambda \|\boldsymbol{u}\|_{\mathcal{A}}, \tag{7}$$

where λ is a positive real number fixed according to the noise. This problem is often referred to as atomic norm denoising. The coefficient λ can be interpreted as a Lagrange multiplier, and this problem can be seen as maximizing a posterior likelihood with a prior on $\boldsymbol{u}$. The quantities we are interested in are the dictionary elements $\boldsymbol{a}_j^\star$ and the coefficients $\boldsymbol{x}^\star$ selected by the minimization, where $\boldsymbol{u}^\star = \sum_{j=1}^{S^\star} x_j^\star \boldsymbol{a}_j^\star(\boldsymbol{t})$.

2.3 More complex noise models

If exoplanetary signals are arguably a sum of sines plus noise, the noise variance is not constant. Even more, the noise might not be independent nor Gaussian. Recent papers such as Tuomi et al. (2013) and Rajpaul et al. (2015) stress that the detection efficiency and robustness improve as the noise model becomes more realistic. Aigrain et al. (2011) suggest to consider the RV time series as Gaussian processes: the noise $\boldsymbol{n}(\boldsymbol{t})$ is then characterized by its covariance matrix $\mathbf{V}$ which is such that $V_{kl} = \mathbb{E}\{\boldsymbol{n}(t_k)\boldsymbol{n}(t_l)\}$, $\mathbb{E}$ being the mathematical expectancy. When the noise is stationary, by definition there exists a covariance function R such that $V_{kl} = R(|t_l - t_k|)$; therefore, choosing $\mathbf{V}$ is equivalent to choosing R. This approach is similar to Sulis, Mary & Bigot (2016), which normalizes the periodogram by the power spectrum of the stationary part of the stellar noise. The similarity comes from the fact that the power spectrum of the noise is $P(\omega) = |\mathcal{F}(R)|^2$, where $\mathcal{F}$ denotes the Fourier transform.

Here, the noise is assumed to be Gaussian of covariance matrix $\mathbf{V}$. In that case, the logarithm of the likelihood is (e.g. Baluev 2011, equation 21; Pelat 2013)

$$\ln(L) = -\frac{m}{2}\ln(2\pi) - \frac{1}{2}\det(\mathbf{V}) - \frac{1}{2}(\boldsymbol{y} - \mathbf{A}\boldsymbol{x})^T \mathbf{V}^{-1} (\boldsymbol{y} - \mathbf{A}\boldsymbol{x}), \tag{8}$$

where the subscript T denotes the matrix transpose. Assuming that the matrix $\mathbf{V}$ is fixed, we wish to minimize $(\boldsymbol{y}-\mathbf{A}\boldsymbol{x})^T\mathbf{V}^{-1}(\boldsymbol{y}-\mathbf{A}\boldsymbol{x})$. If $\mathbf{V}^{-1}$ admits a square root, then $\mathbf{W}$ is chosen such that $\mathbf{W}^2=\mathbf{V}^{-1}$. This is the case when $\mathbf{V}$ is symmetric positive definite, which is the case for covariance matrices of stationary processes. Consequently, $\|\mathbf{W}(\mathbf{A}\boldsymbol{x}-\boldsymbol{y})\|_{\ell_2}^2=(\boldsymbol{y}-\mathbf{A}\boldsymbol{x})^T\mathbf{V}^{-1}(\boldsymbol{y}-\mathbf{A}\boldsymbol{x})$ is always ensured for Gaussian noises. We then obtain the minimization

$$\arg\min_{\boldsymbol{u}\in\mathbb{C}^m} \quad \|\mathbf{W}(\boldsymbol{u}-\boldsymbol{y}(\boldsymbol{t}))\|_{\ell_2}^2+\lambda\|\boldsymbol{u}\|_{\mathcal{A}}. \tag{9}$$

Handling problem (5) with correlated measurements and noise has been investigated by Arildsen & Larsen (2014). However, to the best of our knowledge, the formulation above is not mentioned in the literature; thus, we will briefly discuss its features.

The ability of problem (5) to unveil the true non-zero coefficients of x improves as the so-called mutual coherence of matrix $\mathbf{A}$ diminishes (Donoho et al. 2006). This one is defined as the maximum correlation between two column vectors of $\mathbf{A}$. We here consider a weight matrix, but we can go back to the previous problem by noting that $\mathbf{W}(\mathbf{A}\boldsymbol{x}-\boldsymbol{y})$ can be re-written as $\mathbf{A}'\boldsymbol{x}-\boldsymbol{y}'$, where $\mathbf{A}'=\mathbf{W}\mathbf{A}$ and $\boldsymbol{y}'=\mathbf{W}\boldsymbol{y}$. If we now consider two column vectors of $\mathbf{A}'$, $\boldsymbol{a}_1'=\mathbf{W}\boldsymbol{a}_1$ and $\boldsymbol{a}_2'=\mathbf{W}\boldsymbol{a}_2$, their correlation is $\boldsymbol{a}_1'^T\boldsymbol{a}_2'=\boldsymbol{a}_1W^TW\boldsymbol{a}_2=\boldsymbol{a}_1V^{-1}\boldsymbol{a}_2$. In other words, introducing a matrix $\mathbf{W}$ only comes down to changing the scalar product. This should not be surprising. The matched filter technique (Kay 1993) proposes to detect a model $\boldsymbol{x}$ in a signal $\boldsymbol{s}=\boldsymbol{x}+\boldsymbol{n}$, where $\boldsymbol{n}$ is a noise of covariance matrix $\mathbf{V}$ if $\boldsymbol{x}\mathbf{V}^{-1}\boldsymbol{s}\leq\gamma$, where γ is a threshold. This means if the correlation is sufficient for a non-trivial scalar product.

In the case of an independent Gaussian noise, its covariance matrix $\mathbf{V}$ is diagonal and its elements are σ_k^2, where σ_k is the measurement error at time t_k. $\mathbf{W}$ is defined as $\mathbf{V}^{-1/2}$ so is a diagonal matrix of elements $w_{kk}=1/\sigma_k$. Therefore, $\boldsymbol{a}_1'^T\boldsymbol{a}_2'=\boldsymbol{a}_1\mathbf{W}^T\mathbf{W}\boldsymbol{a}_2=\sum_{k=1}^n\frac{\boldsymbol{a}_1(t_k)\boldsymbol{a}_2(t_k)}{\sigma_k^2}$. This is compatible with the behaviour we intuitively expect: the less precise is the measurement, the lesser the correlation between the signals matter through the weighting by σ_k.

Unfortunately, having a non-identically independent distributed (i.i.d.) Gaussian noise model biases the estimates of the true signals as it acts as a frequency filter. Whether this bias prevents one from having the benefits of a correct noise model is discussed in Appendix B. We show that choosing an appropriate weight matrix $\mathbf{W}$ indeed allows us to see signals that would be buried in the red noise otherwise.

3 IMPLEMENTATION

3.1 Overview

As said above, stable planetary systems are quasi-periodic. This means in particular that radial velocity measurements are well approximated by a linear combination of a few vectors $\mathrm{e}^{-\mathrm{i}\omega t}$ and $\mathrm{e}^{\mathrm{i}\omega t}$. The minimization problem (7) seems therefore well suited for searching for exoplanets. This section is concerned with the numerical resolution, and the numerous issues it raises: the numerical scheme to be used, the choice of the algorithm parameters and the evaluation of the confidence in a detection.

Solving equation (7) is done either by reformulating it as a quadratic program (Candès & Fernandez-Granda 2013; Chen & Chi 2014; Tang et al. 2013b) or by discretizing the dictionary (Tang, Bhaskar & Recht 2013a). The first one necessitates to see the sampling as a regularly spaced one with missing samples. As the measurement times are far from being equispaced in the considered applications, the required time discretization results in large matrices. Therefore, the second approach is used.

Let us pick a set of frequencies equispaced with interval $\Delta\omega$, $\Omega=\{\omega_k=k\Delta\omega,\ k=0,\ldots,n\}$ and an $m\times 2n$ matrix $\mathbf{A}$ whose columns are $\mathrm{e}^{-\mathrm{i}\omega_k t}$ and $\mathrm{e}^{\mathrm{i}\omega_k t}$. In that case, equation (9) reduces to

$$\arg\min_{\boldsymbol{x}\in\mathbb{C}^{2n}} \quad \|\mathbf{W}(\mathbf{A}\boldsymbol{x}-\boldsymbol{y})\|_{\ell_2}^2+\lambda\|\boldsymbol{x}\|_{\ell_1}, \tag{10}$$

which is often referred to as the Least Absolute Shrinkage and Selection Operator (LASSO) problem when W is the identity matrix. As the parameter λ is not so easy to tune, an equivalent formulation of discretized equation (9) is chosen,

$$\boldsymbol{x}^\star=\arg\min_{\boldsymbol{x}\in\mathbb{C}^{2n}} \quad \|\boldsymbol{x}\|_{\ell_1} \quad \text{s. t.}$$
$$\|\mathbf{W}(\mathbf{A}\boldsymbol{x}-\boldsymbol{y})\|_{\ell_2}\leqslant\epsilon, \tag{11, BP$_{\epsilon,\mathbf{W}}$}$$

where ϵ is a positive number. By 'equivalent', we mean there exists a λ_ϵ such that the solution of equation (10) is equal to the solution of equation (11,BP$_{\epsilon,W}$) (Rockafellar 1970). As this problem will often be referred to, we add to the equation number BP$_{\epsilon,W}$ in the rest of the text, BP standing for basis pursuit. There are several codes written to solve equation (5). The existing codes we have tested for analysing radial velocity data sets are ℓ_1-magic (Candès et al. 2006a), SparseLab (Donoho 2006), NESTA (Becker, Bobin & Candès 2011), CVX (Grant & Boyd 2008), spectral compressive sampling (Duarte & Baraniuk 2013) and SPGL1 (van den Berg & Friedlander 2008). The latter gave the best results in general for exoplanetary data and consequently is the one we selected (the code can be downloaded from this link[1]).

The solution of equation (11,BP$_{\epsilon,W}$) offers an estimate for the periods, but the efficiency of the method can be improved by using a moving average on $\boldsymbol{x}^\star$, to approximate better equation (9). Indeed, if a sine of frequency ω_0 and amplitude K is in the signal, corollary 1 (Tang et al. 2013a) shows that the solution of equation (5) $\boldsymbol{x}^\star$ verifies

$$K \quad\approx\quad \sum_{|\omega_k|\in[\omega_0-\eta,\omega_0+\eta]} x^\star(\omega_k) \tag{12}$$

rather than $|x(\omega_0)|\approx K$. The coefficients $x^\star(\omega_k)$ are added up for ω_k lying in a certain interval of length 2η (see Section 3.6).

Finally, the confidence in the detection must be estimated. Problem (11,BP$_{\epsilon,W}$) selects significant frequencies in the data, but the estimates of their amplitude are biased due to the ℓ_1 norm minimization. To obtain unbiased amplitudes, we first check that the peaks are not aliases of each other. Then the most significant peaks are fitted until non-significant residuals are obtained (see Section 3.7.4).

In summary, the method follows a seven-step process.

(i) Pre-process the data: remove the mean in radial velocity data or an estimate of the stellar noise.

(ii) Choose the discrete grid Ω, tolerance ϵ, weighting matrix $\mathbf{W}$ and the width η of the interval over which the result of equation (11,BP$_{\epsilon,W}$) is averaged.

(iii) Define the dictionary $\mathcal{A}$ and normalize the columns of $\mathbf{WA}$.

(iv) Run the program solving the convex optimization (11,BP$_{\epsilon,W}$) to obtain $\boldsymbol{x}^\star$.

(v) Denoting $\Omega=[\omega_{\min},\omega_{\max}]$ for each frequency $\omega\in\{\omega_{\min}+\eta,\ldots,\omega_{\max}-\eta\}$, sum up the amplitudes of $x^\star(\omega')$ from $\omega'\in[-\omega-\eta,-\omega+\eta]\cup[\omega-\eta,\omega+\eta]$ to obtain a smoothed figure $\boldsymbol{x}^\sharp$.

[1] https://www.math.ucdavis.edu/~mpf/spgl1/supplement.html

(vi) Plot $\boldsymbol{x}^\sharp$ as a function of the frequencies or the periods.
(vii) Evaluate the significance of the main peaks (Fig. 6).

Each of these steps is detailed in the following sections.

3.2 Optimization routine

Many solvers can handle problem (5); however, their precision and speed vary. Among the solvers tested, SPGL1 (van den Berg & Friedlander 2008) gives the best results in general. This one has several user-defined parameters such as a stopping criterion that must be tuned. For a given tolerance, this one is $\frac{|\|\mathbf{A}\boldsymbol{x}-\boldsymbol{y}\|_{\ell_2}-\epsilon|}{\max\left(1,\|\mathbf{A}\boldsymbol{x}-\boldsymbol{y}\|_{\ell_2}\right)} < \text{tol}$. The default parameters seem acceptable, in particular $\text{tol} = 10^{-4}$.

3.3 Dictionary A

To estimate the spectrum, a natural choice for the columns of matrix $\mathbf{A}$ is $(\mathrm{e}^{-\mathrm{i}\omega t}, \mathrm{e}^{\mathrm{i}\omega t})$. However, the data might not contain only planetary signals. In the case of a binary star, a linear trend t and a quadratic term t^2 are added. If the star is active, the ancillary measurements are also added.

The method described in Section 3 is applicable to a wider range of dictionary. As the timespan of the observations is in general a few years, the signal might be more sparsely represented either by Poisson terms $[(a_0 + a_1 t + a_2 t^2 + \cdots)\cos(\omega t + \phi)]$ or Keplerian motions. In the latter case, column vectors would be of the form $\frac{r}{a}\mathrm{e}^{\mathrm{i}\boldsymbol{\nu}(t)}$, where $\boldsymbol{\nu}(t)$ is a vector of true anomalies depending on the period P, eccentricity e and initial mean anomaly M_0 (or any combination of three variables that cover all possible orbits). Unfortunately, the size of A increases exponentially with the number of parameters describing the dictionary elements (here P, e, M_0).

3.4 Pre-processing

Theoretical results in Tang et al. (2013a) guarantee that the solution to equation (5) will be close to equation (7) as the discretization gets finer, provided the dictionary is continuous. As linear trends or stellar activity-related signals are not sine, removing these from the data before solving equation (11,$\mathrm{BP}_{\epsilon,W}$) is crucial. The mean, a linear trend and estimates of the stellar noise can be fitted and removed. We reckon this is contrary to the philosophy of fitting the whole model at once. However, the vectors fitted are included again in the dictionary which allows us to mitigate the distortions induced by their removal.

Secondly, to make the precision of the SPGL1 solver independent from the value of $\mathbf{W}\boldsymbol{y}$, the weighted observations $\mathbf{W}\boldsymbol{y}$ are normed by $\|\mathbf{W}\boldsymbol{y}\|_{\ell_2}$, and the columns of the matrix $\mathbf{WA}$ are also normed. Denoting by $\boldsymbol{y}' = \frac{1}{\epsilon}\mathbf{W}\boldsymbol{y}/\|\mathbf{W}\boldsymbol{y}\|_{\ell_2}$ and $\mathbf{A}' = \frac{1}{\epsilon}(\mathbf{W}A_k/\|\mathbf{W}A_k\|_{\ell_2})_{k=1,\ldots,n}$, we set in input of the solver

$$\arg\min_{\boldsymbol{x}\in\mathbb{C}^n} \|\boldsymbol{x}\|_{\ell_1} \quad \text{s.t.} \quad \left\|\mathbf{A}'\boldsymbol{x} - \boldsymbol{y}'\right\|_{\ell_2} \leqslant 1, \tag{13}$$

to always be in the same kind of use of the solver and ensure that the accuracy of the result does not depend on its units. Going back to the correct units in the post-processing step is described in Section 3.6.

3.5 Tuning

Choice of $\mathbf{W}$. We have seen in Section 2.3 that the weight matrix $\mathbf{W}$ is characterized by the covariance function R via $\mathbf{W}_{kl} = R(|t_k - t_l|)$. Several forms for the covariance functions were suggested (e.g. Rajpaul et al. 2015). Here we only consider exponential covariances, which are

$$\begin{aligned} R(\Delta t) &= \sigma_R^2 \mathrm{e}^{-\frac{|\Delta t|}{\tau}}, \quad \Delta t \neq 0 \\ R(0) &= \sigma_W^2 + \sigma_R^2, \end{aligned} \tag{14}$$

where the subscripts W and R stand respectively for white and red. As the red and white noises are here supposed independent, the covariance function of their sum is the sum of their covariance functions. Therefore, the matrix $\mathbf{W}$ is such that its diagonal terms are $V_{kk} = \sigma_k^2 + \sigma_W^2 + \sigma_R^2$ and $V_{kl} = \sigma_R^2 \mathrm{e}^{-\frac{|t_k - t_l|}{\tau}}$ for $k \neq l$.

Choice of Ω. We have two parameters to choose: the grid span and the grid spacing. For the first one, we take 1.5 cycles d^{-1} as a default value but it is also advisable to re-do the analysis for 0.95 cycles d^{-1}, as discussed in the examples in Section 2. We ensure that if the signal is made of sinusoids (a.k.a. it is quasi-periodic), there exists at least one vector x verifying $\|\mathbf{W}(\mathbf{A}\boldsymbol{x} - \boldsymbol{y})\|_{\ell_2} < \epsilon$ that has the correct ℓ_0 norm. Let us consider a signal made of p pure sinusoids sampled at times $\boldsymbol{t} = (t_k)_{k=1,\ldots,m}$, $\boldsymbol{y}(\boldsymbol{t}) = \sum_{j=1}^{p} c_j \mathrm{e}^{\mathrm{i}\omega_j \boldsymbol{t}}$. Assuming that the frequencies on the grid are regularly spaced with step $\Delta\omega$, this leads to the condition (see Appendix A for calculation details)

$$\Delta\omega \leqslant \frac{4}{T} \arcsin \frac{\epsilon}{2\sqrt{\sum_{j=1}^{p} |c_j|^2}\sqrt{\sum_{k=1}^{m} \frac{1}{\sigma_k^2}}}. \tag{15}$$

Let us note that the values of c_j are a priori unknown, so the term $\sqrt{\sum_{j=1}^{p} |c_j|^2}$ has to be approximated. Supposing the signal is made of sinusoids plus small noise, $\sqrt{\sum_{j=1}^{p} |c_j|^2} \approx \|\boldsymbol{y}\|_{\ell_2}/\sqrt{m}$. Furthermore, it must be ensured that all possible significant frequencies are in the signal.

The choice of the grid spacing can be based on other criteria: Stoica & Babu (2012) suggest to choose a spacing such that the 'practical rank of matrix $M_{kl} = \mathrm{e}^{\mathrm{i}\Delta\omega(t_k - t_l)}$ is equal to one. This term designates the number of singular values above a certain threshold. Here the condition states that only one singular value is non-negligible. Let us also mention that one can perform the reconstruction with different grids and average out the results. However, this approach does not practically generate better results than using a finer grid.

Choice of ϵ. The error is due to two sources: grid discretization which gives an error ϵ_{grid} and noise, which yields ϵ_{noise}. Supposing the noise is Gaussian, denoting by $\boldsymbol{y}_t$ the underlying non-noisy observations, $\|\mathbf{W}(\boldsymbol{y}_t - \boldsymbol{y})\|_{\ell_2}^2$ as a function of random variable $\boldsymbol{y} = \boldsymbol{y}_t + \boldsymbol{n}$ follows a χ^2 distribution with m degrees of freedom. Denoting its cumulative distribution function (CDF) by $F_{\chi_m^2}$, the probability $1 - \alpha$ that the true signal $\boldsymbol{y}_t$ is in the set $\{\boldsymbol{y}', \|\mathbf{W}(\boldsymbol{y}' - \boldsymbol{y})\|_{\ell_2}^2 \leq \epsilon_{\text{noise}}\}$ is

$$F_{\chi_m^2}(\epsilon_{\text{noise}}^2) = 1 - \alpha. \tag{16}$$

The bound ϵ_{noise} is determined according to the equation above for a small α. Once ϵ_{noise} is chosen, rearranging equation (15) gives a minimal value of ϵ_{grid} that ensures that a signal with a correct ℓ_0 norm exists,

$$\epsilon_{\text{grid}} = 2\sqrt{\sum_{j=1}^{p} |c_j|^2}\sqrt{\sum_{k=1}^{m} \frac{1}{\sigma_k^2}} \sin \frac{\Delta\omega T_{\text{obs}}}{4}. \tag{17}$$

An alternative is to set ϵ to zero and let the algorithm find a representation for the noise, which will not be sparse. In that case,

one must obviously not perform the re-normalization of the columns of $\mathbf{WA}$ by ϵ of Section 3.4. Below a certain amplitude, a 'forest' of peaks would be seen on the ℓ_1-periodogram. This has the advantage to give an estimation of the noise structure. However, this method is more sensitive to the solver inner uncertainties and requires more time; it was not retained for this work.

Choice of η. See the next section.

3.6 Post-processing

Once the solution to equation (11,$\mathrm{BP}_{\epsilon,W}$) is computed, the spectrum $\boldsymbol{x}^\star$ is filtered with a moving average. We expect from discretization (9) that the frequencies might leak to close frequencies. Indeed, the amplitude of the solution to equation (11,$\mathrm{BP}_{\epsilon,W}$) might be untrustworthy. When the signal is made of several frequencies, the solution might overestimate the one with the greatest amplitude, and underestimate the others; this problem arises especially when less than a hundred observations are available. To mitigate this effect, one can sum up the contribution of subsequent frequencies and estimate the amplitude of the resulting signal. If $\boldsymbol{x}^\star$ is the solution to equation (11,$\mathrm{BP}_{\epsilon,W}$), denoting by $x^\star(\omega)$ the coefficient corresponding to frequency ω, we compute

$$\hat{\boldsymbol{y}}_\omega(\boldsymbol{t}) = \|\mathbf{W}\boldsymbol{y}\|_{\ell_2} \sum_{\substack{\omega' \in \Omega \\ \omega - \eta \leqslant |\omega'| \leqslant \omega + \eta}} \frac{x^\star(\omega')\boldsymbol{a}_{\omega'}(\boldsymbol{t})}{\|\mathbf{W}\boldsymbol{a}_{\omega'}(\boldsymbol{t})\|_{\ell_2}}, \tag{18}$$

where $\boldsymbol{a}_{\omega'}(\boldsymbol{t})$ is the column of A corresponding to frequency ω'. The terms $\|\mathbf{W}\boldsymbol{y}\|_{\ell_2}$ and $1/\|\mathbf{W}\boldsymbol{a}_{\omega'}(\boldsymbol{t})\|_{\ell_2}$ appear because the columns of WA and the weighted observations Wy were normalized in Section 3.4. The vector $\hat{\boldsymbol{y}}_\omega(\boldsymbol{t})$, $t = t_1, \ldots, t_m$ is approximately a sine function; the new estimation of the signal power is

$$x^\sharp(\omega) = \max_{t_1..t_m} |\hat{y}_\omega(t_k)|. \tag{19}$$

Other estimates are possible, such as the power of a sine at frequency ω fitted on $\hat{y}(\omega)$. Though the choice of η is heuristic, corollary 1 of Tang et al. (2013a) is used as a guideline. It indeed states that the summed amplitudes of coefficients of $\boldsymbol{x}^\star$ within a certain distance η_0 from the actual peak in the signal tend to the appropriate value as the discretization step tends to zero. In the proof, they choose ϵ such that the balls of width η_0 centred around the true peaks have a null intersection. Thus, it seems reasonable to select η as the largest interval within which the probability to distinguish frequencies is low. Values such as $\approx 0.5\pi/T_{\mathrm{obs}}$ to π/T_{obs} are robust in practice.

3.7 Significance and uncertainties

3.7.1 Detection threshold

It is simple to associate a 'global' false alarm probability (FAP) to the ℓ_1-periodogram similar to the classical FAP of the Lomb–Scargle periodogram (Scargle 1982, equation 14). Let us consider the probability that '$x = 0$ is not a solution knowing the signal is pure independent Gaussian noise'. Denoting this probability $\tilde{\alpha}$, following notations of Section 2.1, $\epsilon^2 = {F_{\chi^2}}^{-1}(1-\tilde{\alpha})$. As $\epsilon \approx \epsilon_{\mathrm{noise}}$, the value of $\tilde{\alpha}$ is close to the user-defined parameter α. In the Lomb–Scargle case, the FAP obeys: 'if the maximum of the periodogram is z, then the FAP is $\beta(z)$', where β is an increasing function of z [often taken as $\beta(z) = 1 - (1 - \mathrm{e}^{-z})^M$, where M is a parameter fitted with numerical simulations; Scargle 1982; Horne & Baliunas 1986; Cumming 2004. Here the formulation is 'if the solution to equation (11,$\mathrm{BP}_{\epsilon,W}$) is not zero, then a signal has been detected with an FAP lower than or equal to α'.

3.7.2 Statistical significance of a peak

The discussion above points out similarities with the FAP defined for periodograms. This one and the global FAP share in particular that they only allow one to reject the hypothesis that the signal is pure Gaussian noise of covariance matrix $\mathbf{W}$. However, the problem is rather to determine if a given peak indicates a true underlying periodicity, and if this one is due to a planet.

In that scope, our goal is to test if the harmonics spotted by the ℓ_1-periodogram are statistically significant. Ultimately, one can use statistical hypothesis testing, which can be time consuming. To quickly assess the significance of the peaks, two methods seem to be efficient.

(i) Re-sampling: taking off randomly 10–20 per cent of the data and re-computing the ℓ_1-periodogram. The peaks that show great variability should not be trusted.

(ii) Using the formulae of the 'residual/recursive periodograms' (Cumming 2004; Baluev 2008, 2009, 2015a; Anglada-Escudé & Tuomi 2012).

The first case is easy to code and has the advantage to implement implicitly a time–frequency analysis. Indeed, we might expect from stellar variability some wavelet-like contributions: a signal with a certain frequency arises and then vanishes. The timespan of observation might be short enough so that feature is mistaken for a truly sinusoidal component. By taking off some of the measurements, we can see if the amplitude of a given frequency varies through time. However, this method requires to re-compute the ℓ_1-periodogram several times and might not be suited for systems with numerous measurements.

3.7.3 Model

As the re-sampling approach is straightforward to code, we will now focus on the recursive periodogram formulae. These ones should be useful for readers more interested in speed than comprehensiveness. In this section, the relevant signal models are defined. We consider that the signal is of the form

$$f_K\left(\boldsymbol{\theta}_0, \left(\boldsymbol{\theta}_{K\,j}\right)_{j=1..np}\right) = \text{non-planetary}\,(\boldsymbol{\theta}_0) + \sum_{j=1}^{n_p} \text{Keplerian}_j\,\left(\boldsymbol{\theta}_{K\,j}\right) \tag{20}$$

or

$$f_C\left(\boldsymbol{\theta}_0, \left(\boldsymbol{\theta}_{C\,j}\right)_{j=1..np}\right) = \text{non-planetary}\,(\boldsymbol{\theta}_0) + \sum_{j=1}^{n_p} \text{Circular}_j\,\left(\boldsymbol{\theta}_{C\,j}\right). \tag{21}$$

That is a sum of a model accounting for non-planetary effects, non-planetary($\boldsymbol{\theta}_0$), $\boldsymbol{\theta}_0$ being a real vector with n_θ components, and a sum of Keplerian or circular curves depending on five resp. three parameters, $\boldsymbol{\theta}_{K\,j} = (k_j, h_j, P_j, A_j, B_j)$ and $\boldsymbol{\theta}_{C\,j} = (P_j, A_j, B_j)$

$$\text{Keplerian}(\boldsymbol{\theta}_K) = A\dot{U}(k, h, P) + B\dot{V}(k, h, P) \tag{22}$$

$$\text{Circular}\,(\boldsymbol{\theta}_C) = A\cos\left(\frac{2\pi t}{P}\right) + B\sin\left(\frac{2\pi t}{P}\right), \tag{23}$$

where $k = e\cos\varpi$, $h = e\sin\varpi$, $\varpi = \omega + \Omega$ is the sum of the argument of periastron and right ascension at ascending node, U, V are the position on the orbital plane rotated by angle ϖ. These variables are chosen to avoid poor determination of the eccentricity and time at periastron for low eccentricities.

We compare subsequently the χ^2 of residuals of a model with n_p and $n_p + 1$ planets. In practice, one selects the tallest peak of the ℓ_1-periodogram, and uses this frequency to initialize a least-squares fit of a circular or Keplerian orbit. Then the two tallest peaks are selected and so on.

To clarify the meaning of the computed FAP, let us define the recursive periodogram, depending on a frequency ω. We denote the χ^2 of the residuals by

$$\chi^2_{K,C}\left(\boldsymbol{\theta}_0^{\text{fit}}, \boldsymbol{\theta}_{n_p}^{\text{fit}}, \omega\right) = \left[\boldsymbol{y} - f_{K,C}\left(\boldsymbol{\theta}_0^{\text{fit}}, \boldsymbol{\theta}_{n_p}^{\text{fit}}, \omega^{\text{fit}}\right)\right]^T \mathbf{V}^{-1} \left[\boldsymbol{y} - f_{K,C}\left(\boldsymbol{\theta}_0^{\text{fit}}, \boldsymbol{\theta}_{n_p}^{\text{fit}}, \omega^{\text{fit}}\right)\right] \tag{24}$$

$$\chi^2_{K,C}\left(\boldsymbol{\theta}_0^{\text{fit}}, \boldsymbol{\theta}_{n_p}^{\text{fit}}\right) = \left[y - f_{K,C}\left(\boldsymbol{\theta}_0^{\text{fit}}, \boldsymbol{\theta}_{n_p}^{\text{fit}}\right)\right]^T \mathbf{V}^{-1} \left[\boldsymbol{y} - f_{K,C}\left(\boldsymbol{\theta}_0^{\text{fit}}, \boldsymbol{\theta}_{n_p}^{\text{fit}}\right)\right]. \tag{25}$$

$f_{K,C}\left(\boldsymbol{\theta}_0^{\text{fit}}, \boldsymbol{\theta}_{n_p}^{\text{fit}}, \omega^{\text{fit}}\right)$ is the model fitted depending on the non-planetary effects $\boldsymbol{\theta}_0$, the (Keplerian or circular) $\boldsymbol{\theta}_{n_p} = (\boldsymbol{\theta}_{K,C\,j})_{j=1,\ldots,n_p}$ parameters of n_p planets plus a circular or Keplerian orbit initialized at frequency ω. $\mathbf{V}$ designates the covariance matrix of the noise model ($\mathbf{V}^{-1} = \mathbf{W}^2$ with the notations above). This one is often assumed to be diagonal but this is not necessary as all the properties of those periodograms come from the fact that they are likelihood ratios. The model fit can be done linearly (Baluev 2008) or non-linearly (Anglada-Escudé & Tuomi 2012). By linear, we mean that among the five or three parameters defined in equations (22) and (23), only $(A_j)_{j=1,\ldots,n_p+1}$ and $(B_j)_{j=1,\ldots,n_p+1}$ are fitted and the non-planetary effects are modelled linearly: there exists a matrix $\boldsymbol{\Phi}$ such that non-planetary$(\boldsymbol{\theta_0}) = \boldsymbol{\Phi\theta_0}$. In the second option, the orbital elements of previously selected planets, the non-planetary effects and the signal at the trial frequency are re-adjusted non-linearly for each trial frequency.

3.7.4 FAP formulae for recursive periodograms

Recursive periodogram is a term that refers to a general concept for comparing the residuals of a model with or without a signal at a given frequency. Here we specialize the formulae we use. Denoting by $P_C(\omega)$ and $P_K(\omega)$ in the circular resp. Keplerian case,

$$P_C(\omega) = N\frac{\chi^2_C(n_p, \omega) - \chi^2_C(n_p)}{\chi^2_C(n_p)} \tag{26}$$

$$P_K(\omega) = \frac{1}{2}\left(\chi^2_K(n_p) - \chi^2_K(n_p, \omega)\right), \tag{27}$$

where $N = m - 2n_p - n_\theta$. The circular case is expression 'z_1' in equation 2 of Baluev (2008), and the Keplerian one is expression 'z' in equation 4 of Baluev (2015a). In what follows, only the circular case will be used.

The quantity we are interested in is the probability that a selected peak is not a planet. We here use the FAP as a proxy for that quantity:

$$\text{FAP}(Z) = \Pr\left\{\max_{\omega\in[0,\omega_{\max}]} P(\omega) > Z \mid \text{non-planetary effects}, n_p\right\}, \tag{28}$$

where $\omega_{\max}$ is the maximum frequency of the periodogram that has been scanned. This FAP is the probability to obtain a peak at least as high as Z while there are only non-planetary effects and n_p planets. Baluev (2008) has computed tight bounds for that quantity in the case of a circular model and a linear fit (corresponding to subscript C), which we reproduce here:

$$\text{FAP}(z, \omega_{\max}) \approx W\gamma\left(\frac{2z}{N_{\mathcal{H}}}\right)^{\frac{1}{2}}\left(1 - \frac{2z}{N_{\mathcal{H}}}\right)^{\frac{N_{\mathcal{H}}+1}{2}}, \tag{29}$$

where $N_{\mathcal{H}}$ is the number of degrees of freedom of the model without the sine at frequency ω, $\gamma = \Gamma(N_{\mathcal{H}}/2)/\,\Gamma((N_{\mathcal{H}}+3/)2)$, Γ being the Euler Γ function, and $W = \omega_{\max}\sqrt{(\bar{\boldsymbol{t}}^2 - \bar{\boldsymbol{t}}^2)/\pi}$, $\boldsymbol{t}$ being the array of measurement times and $\bar{\boldsymbol{t}}$ is the mean value of $\boldsymbol{t}$. We have also tried the exact expression of the so-called Davies bound provided by equations 8, B5 and B7 of Baluev (2008), but the results were very similar to the simpler formula. In the case of Keplerian periodogram, we used equation 21 and 24 from Baluev (2015a).

Again, we emphasize that the interest of the present method is to select candidates for future observations or unveiling signals unseen on periodograms. The FAP formulae used here do not guarantee the planetary origin of a signal. For robust results, statistical hypothesis testing (e.g. Díaz et al. 2016) can be used.

4 RESULTS

4.1 Algorithm tuning

For all the systems analysed in the following sections, the figures called ℓ_1-periodogram represent $x^\sharp(\omega)$ as defined in equation (19) plotted versus periods. The name ℓ_1-periodogram was chosen to avoid the confusion with the generalized Lomb–Scargle (GLS) periodogram defined by Zechmeister & Kürster (2009). In each case, the algorithm is tuned in the following way.

(i) The problem (11,BP$_{\epsilon,W}$) is solved with SPGL1 (van den Berg & Friedlander 2008).
(ii) The solution of SPGL1 is averaged on an interval $\eta = 2\pi/(3T_{\text{obs}})$ according to Section 3.5.
(iii) The grid spacing is chosen according to equation (15).

The importance of the grid span and the tolerance ϵ will be discussed in the examples.

The FAPs are computed according to the procedure described in Section 3.7.4 and are represented in Fig. 6 with decreasing FAP. The ticks in abscissa correspond to the period of the signals and the flag to their semi-amplitude after a non-linear least-squares fit.

In the following, we will present our results for HD 69830, HD 10180, 55 Cnc, GJ 876 and a simulated very active star from the RV Challenge (Dumusque et al. 2016). For each system, the GLS periodogram is plotted along with the ℓ_1-periodogram.

4.2 HD 69830

In Lovis et al. (2006), three Neptune-mass planets are reported around HD 69830 based on 74 measurements of HARPS spanning over 800 d. The precision of the measurements given in the raw data set (from now on called nominal precision) is between 0.8 and

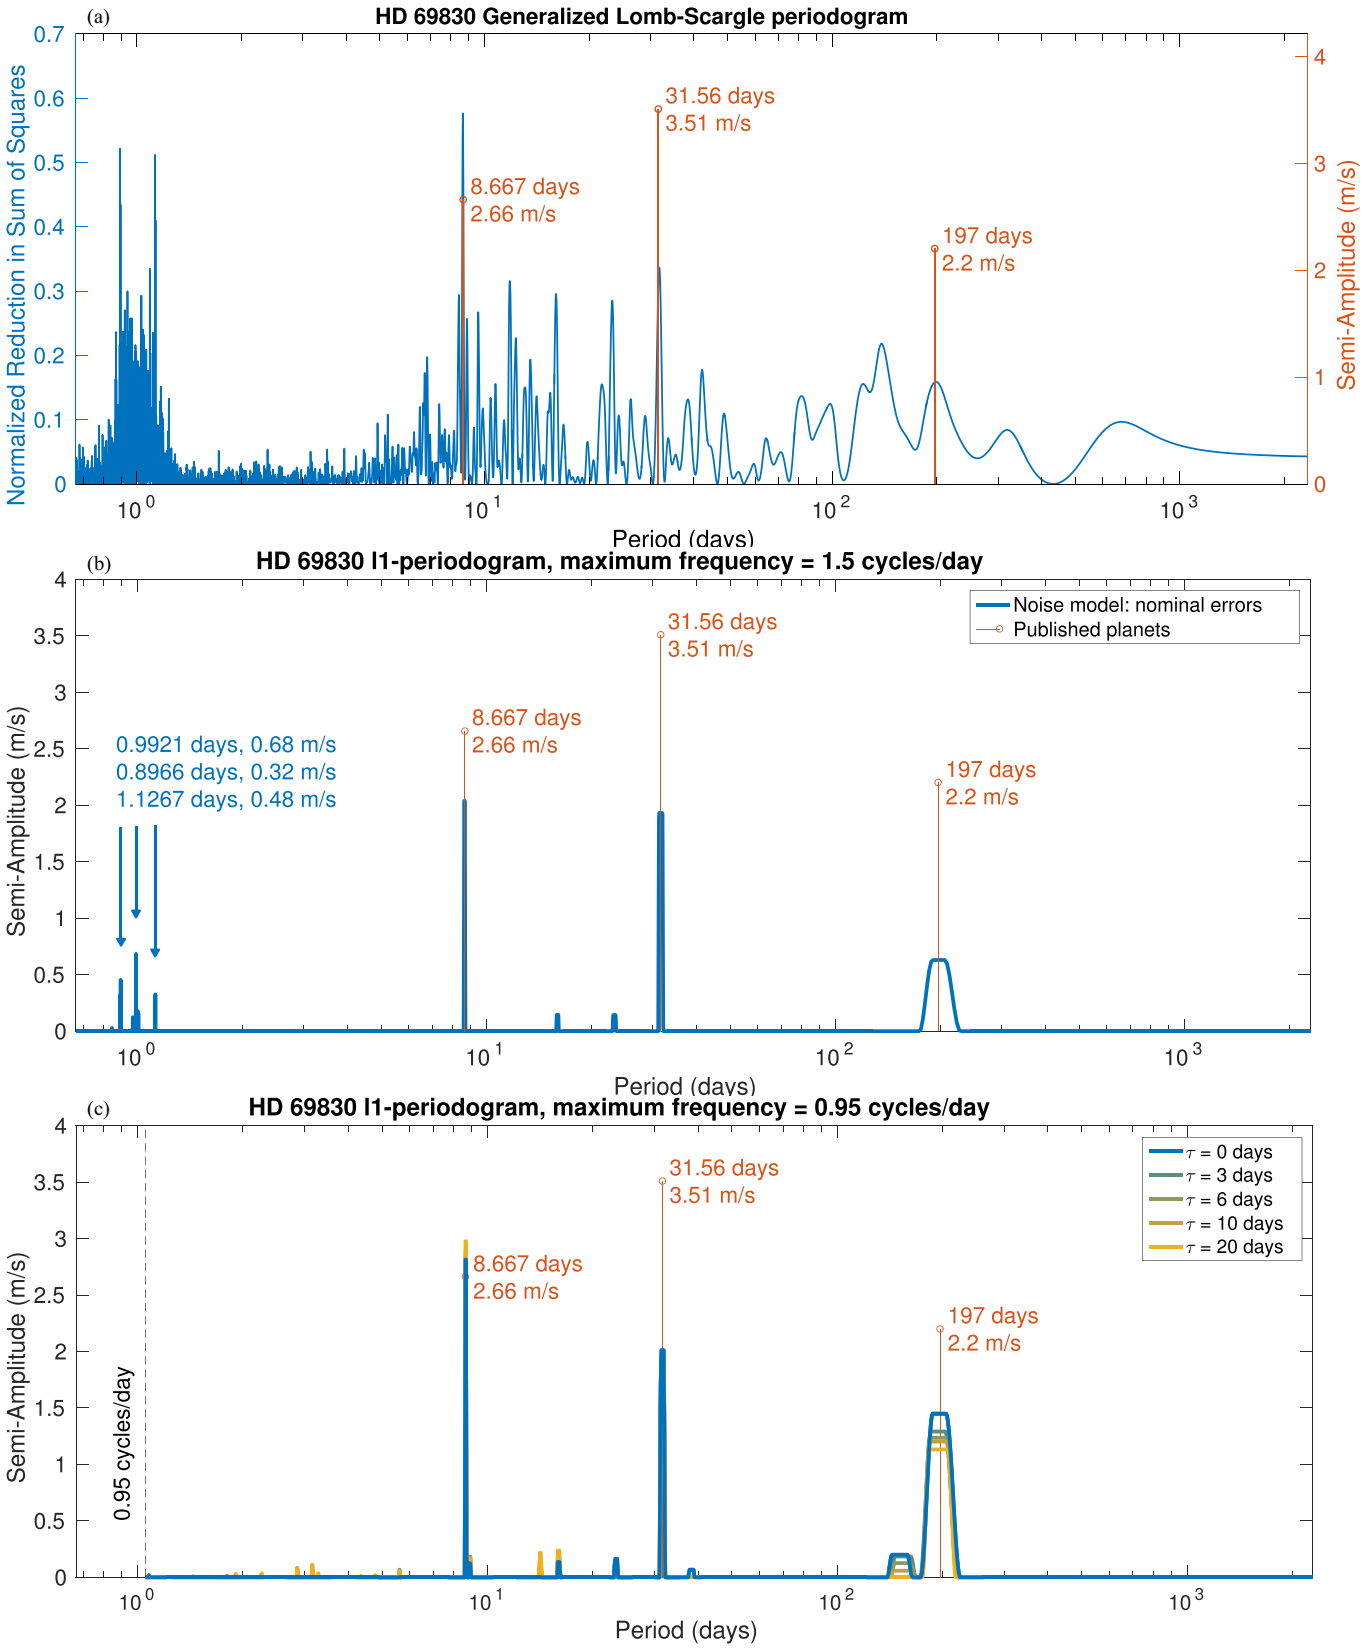

Figure 1. Generalized Lomb–Scargle periodogram and ℓ_1-periodogram of HD 69830 in blue, and published planets are represented by the red stems. The frequency spans used for panels (b) and (c) are respectively 1.5 and 0.95 cycles d^{-1}. The other signals mentioned in Section 1 are spotted by the blue arrows. For all the noise model considered for matrix $\mathbf{W}$, $\sigma_W = 0$, $\sigma_R = 1$ m s^{-1}.

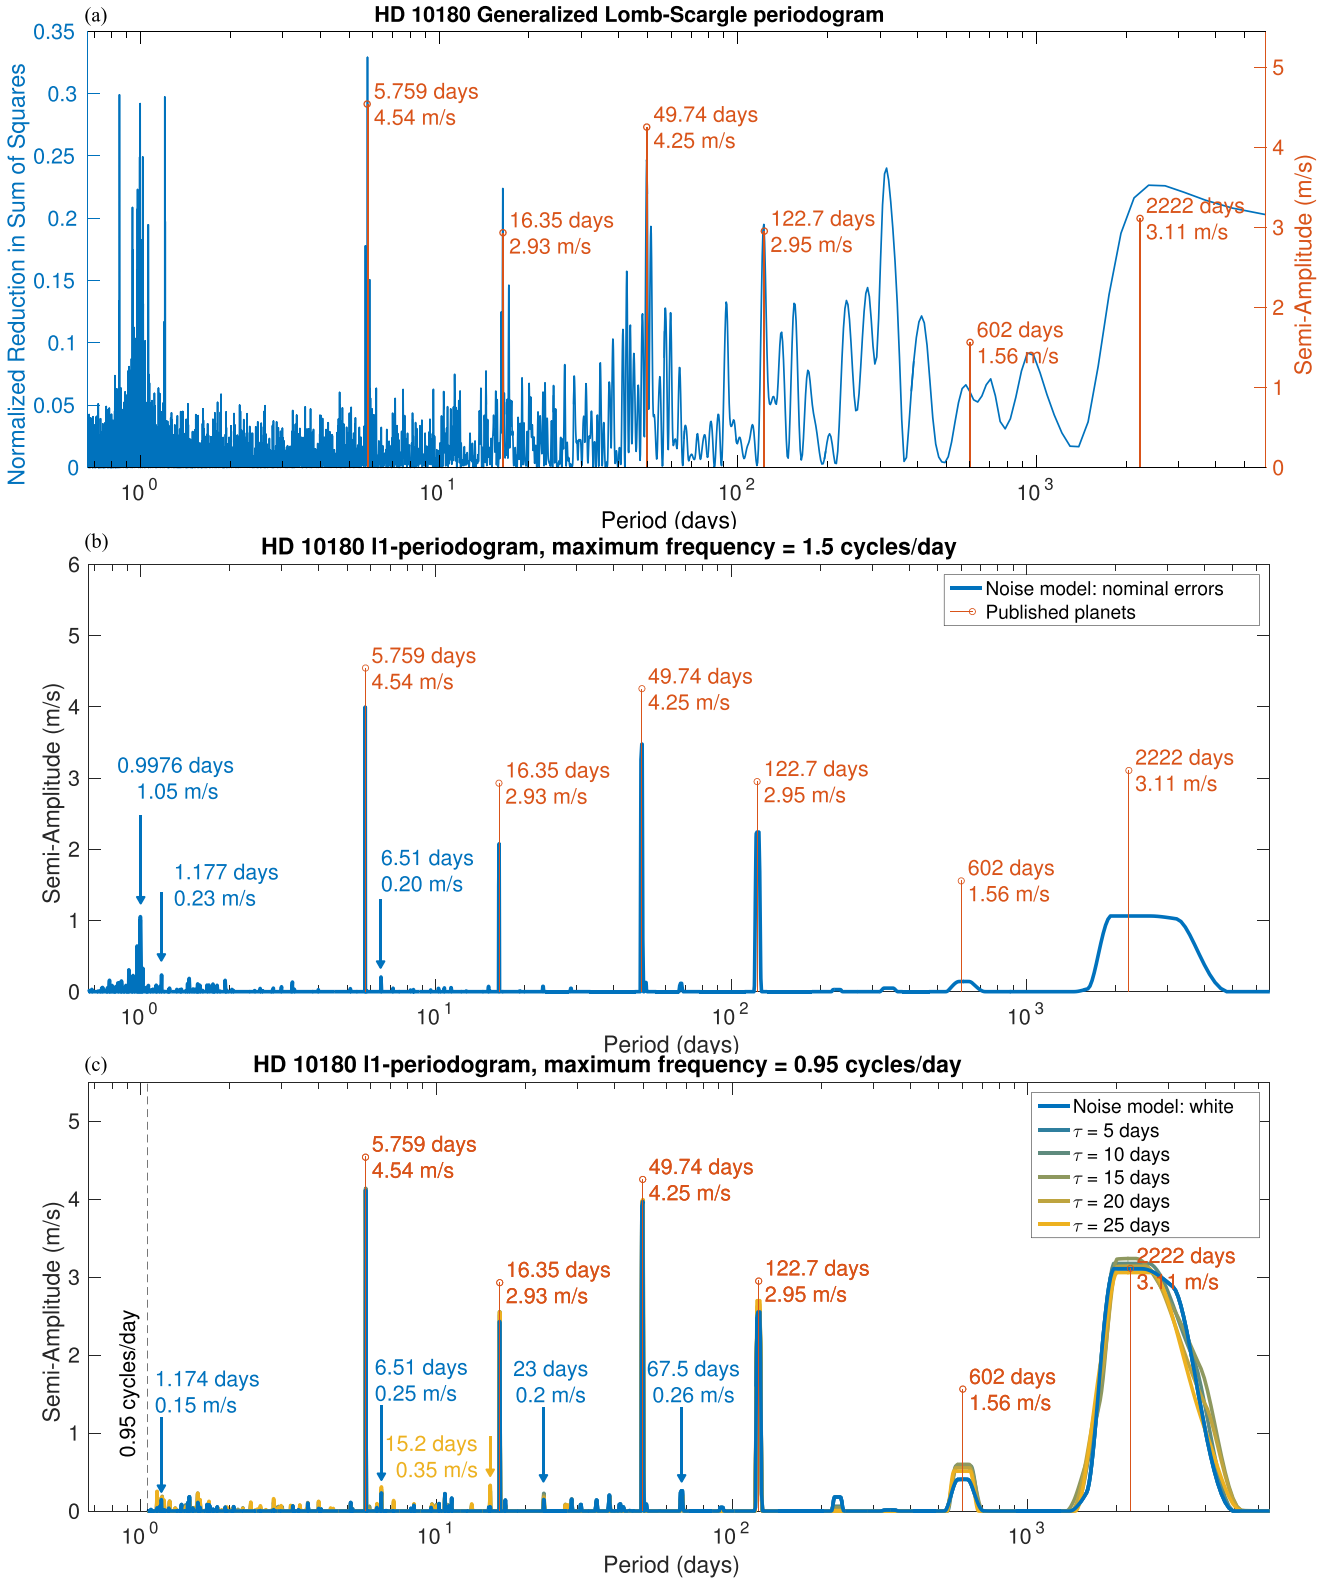

Figure 2. GLS and ℓ_1-periodograms of HD 10180 data set with mean subtracted. The red stems have the periods and amplitude of published planets. The other signals mentioned in Section 2 are spotted by the blue arrows. For all the noise model considered for matrix $\mathbf{W}$, $\sigma_W = 0$, $\sigma_R = 1$ m s^{-1}.

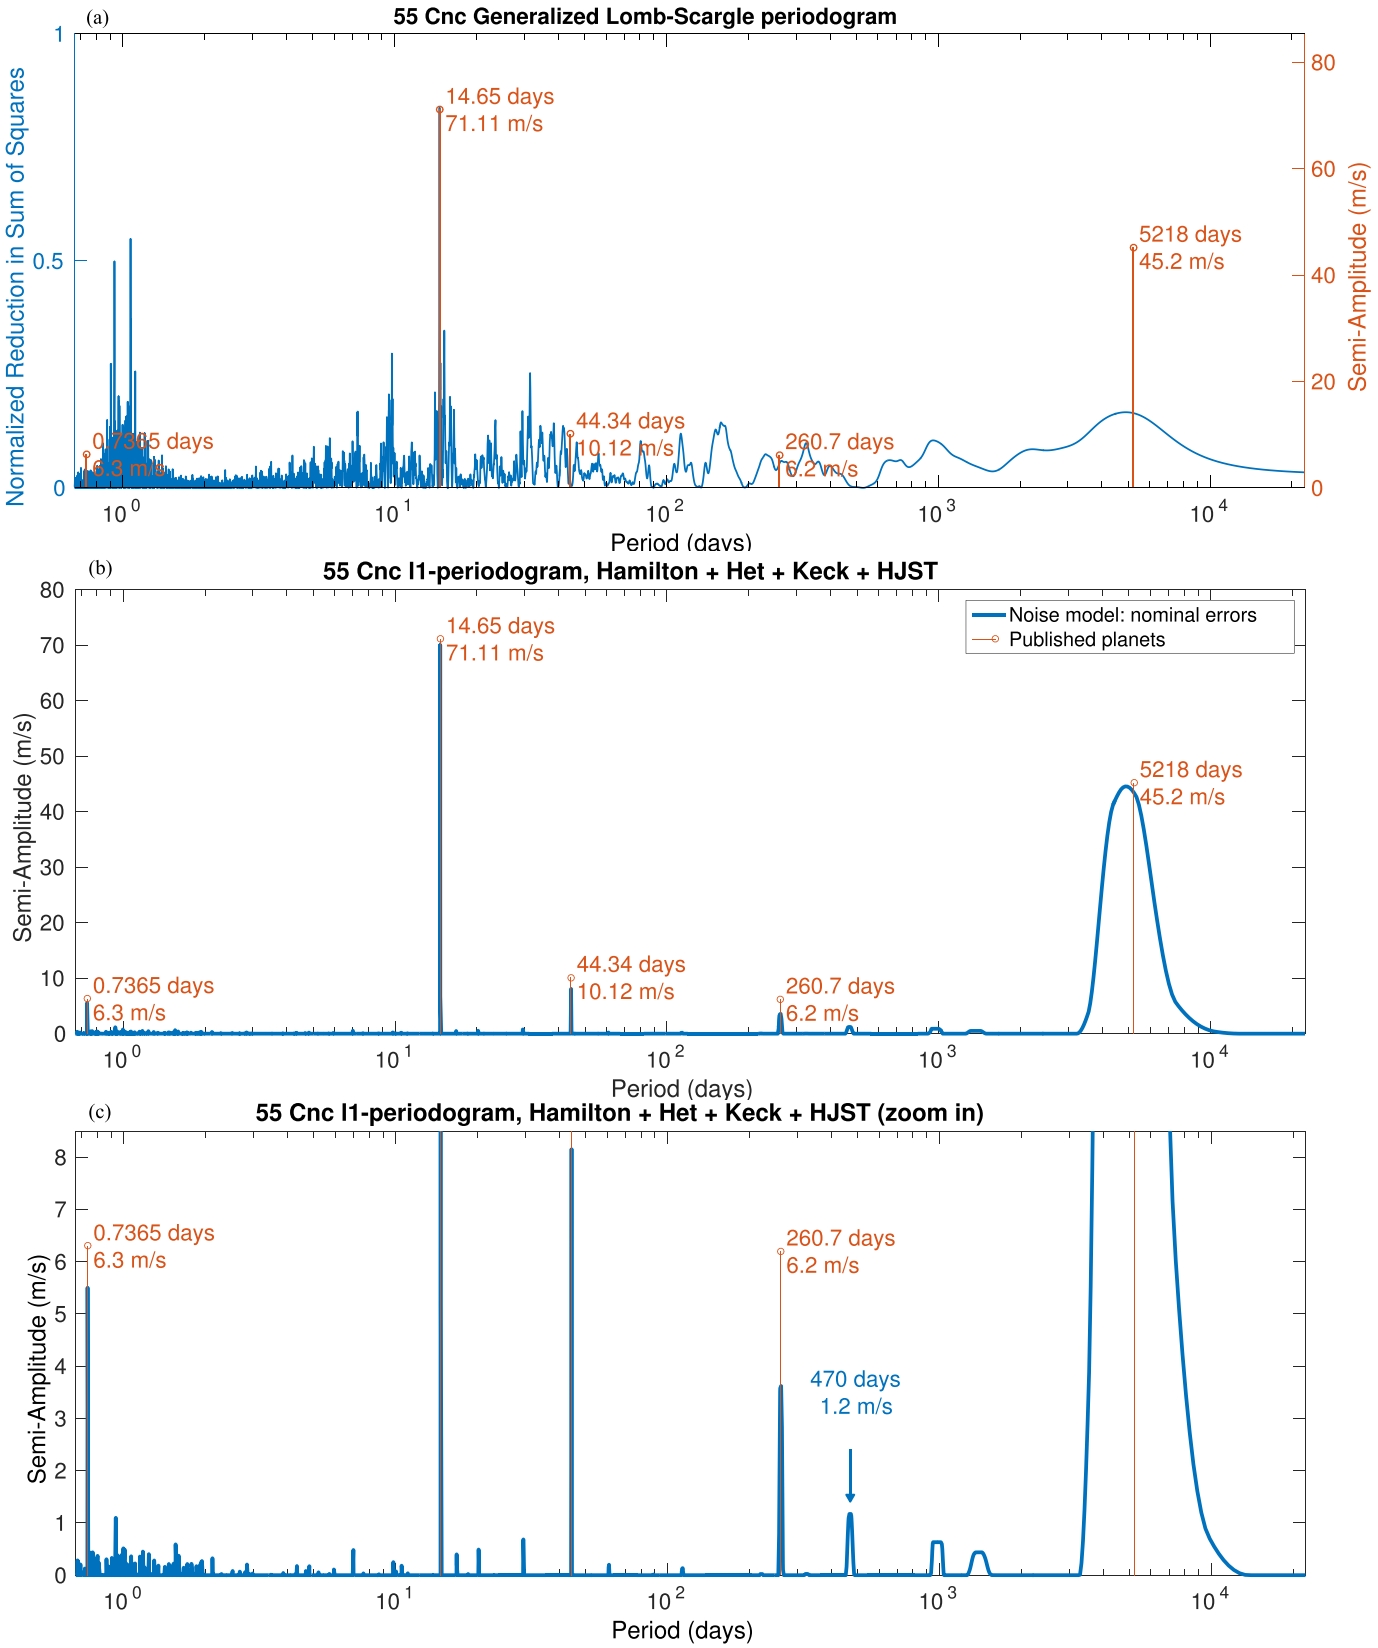

Figure 3. GLS and ℓ_1-periodograms of 55 Cnc data set with mean subtracted. The red stems have the periods and amplitude of published planets. The other signals mentioned in Section 4.4 are indicated by the blue arrows.

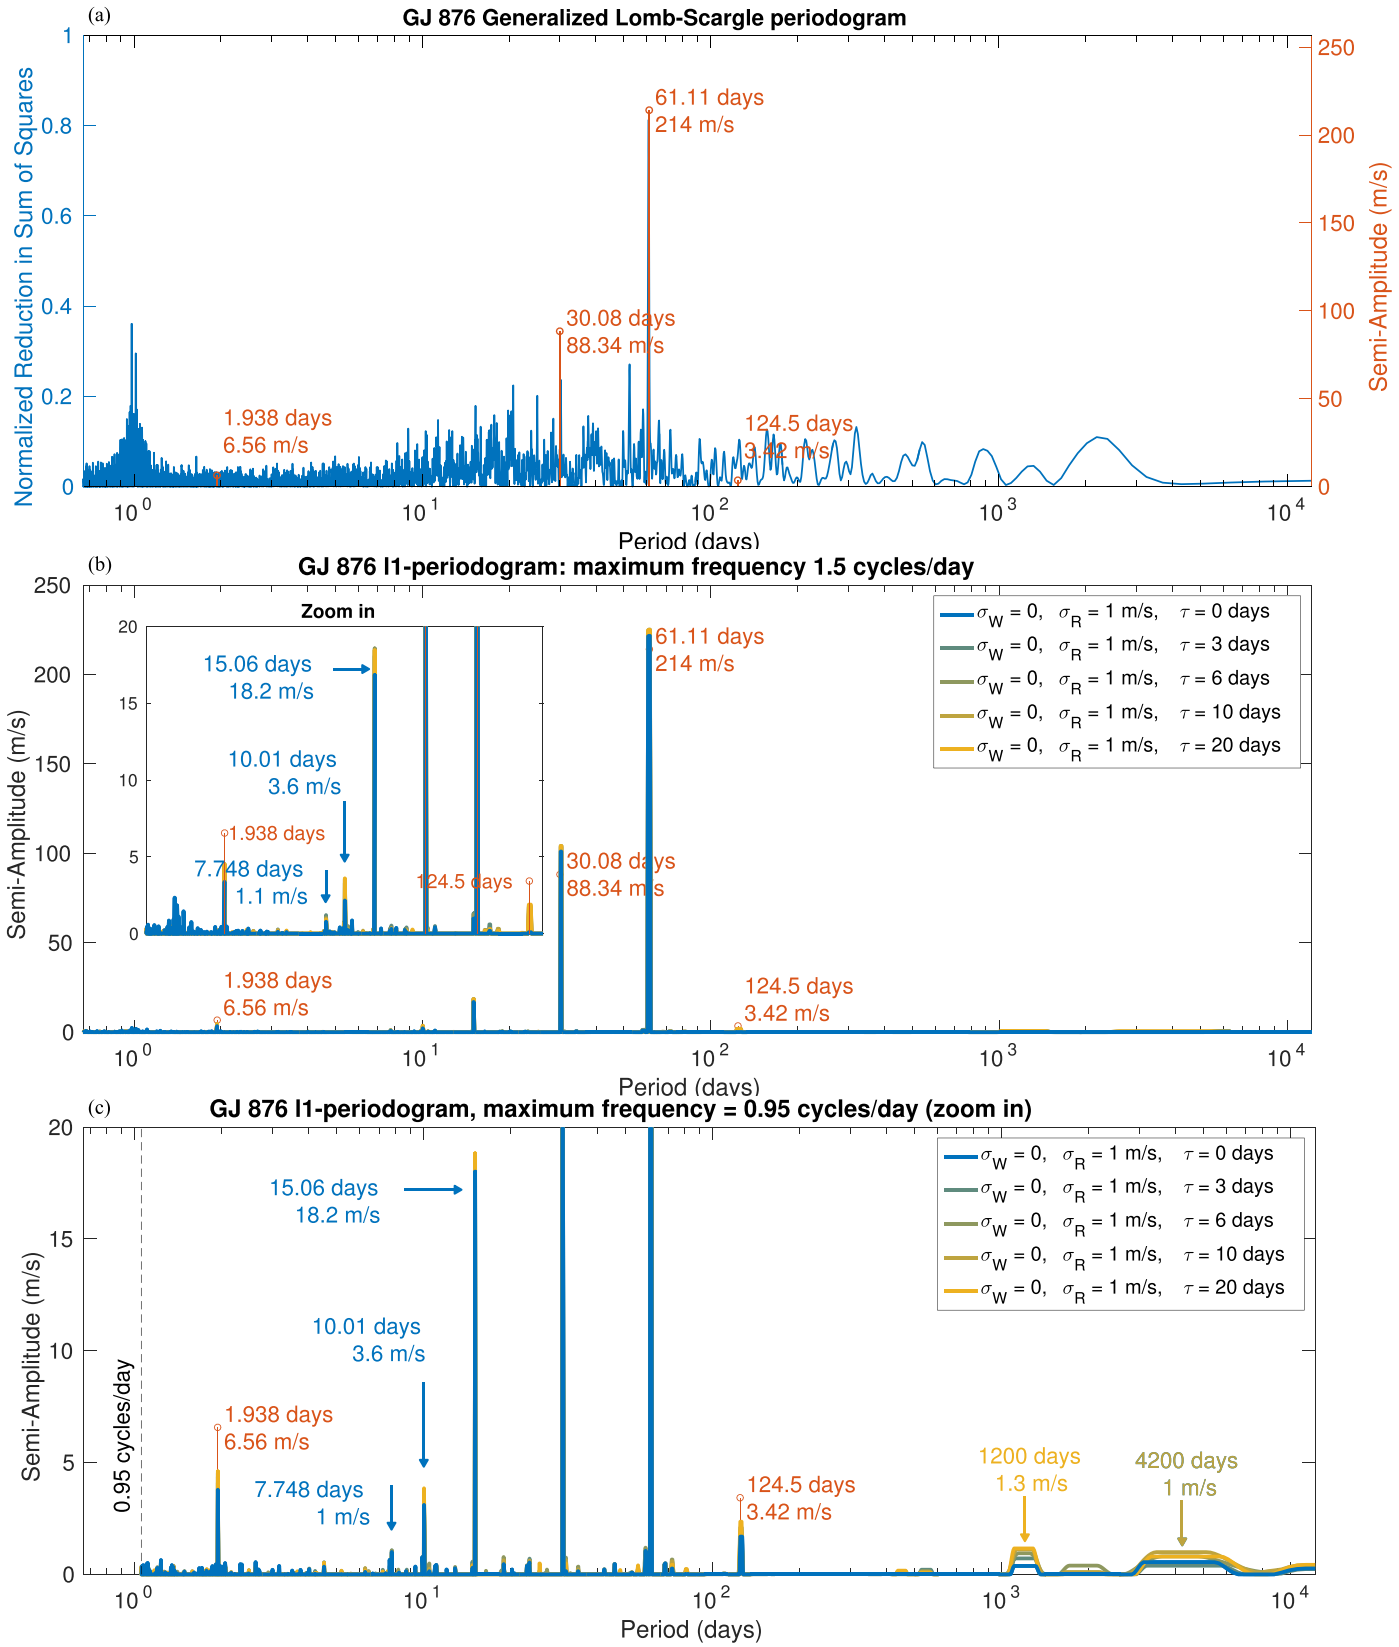

Figure 4. GLS and ℓ_1-periodograms of GJ 876 data set with means of Keck and HARPS measurement respectively subtracted. The red stems have the periods and amplitude of published planets. The other signals mentioned in Section 4.5 are indicated by the blue arrows.

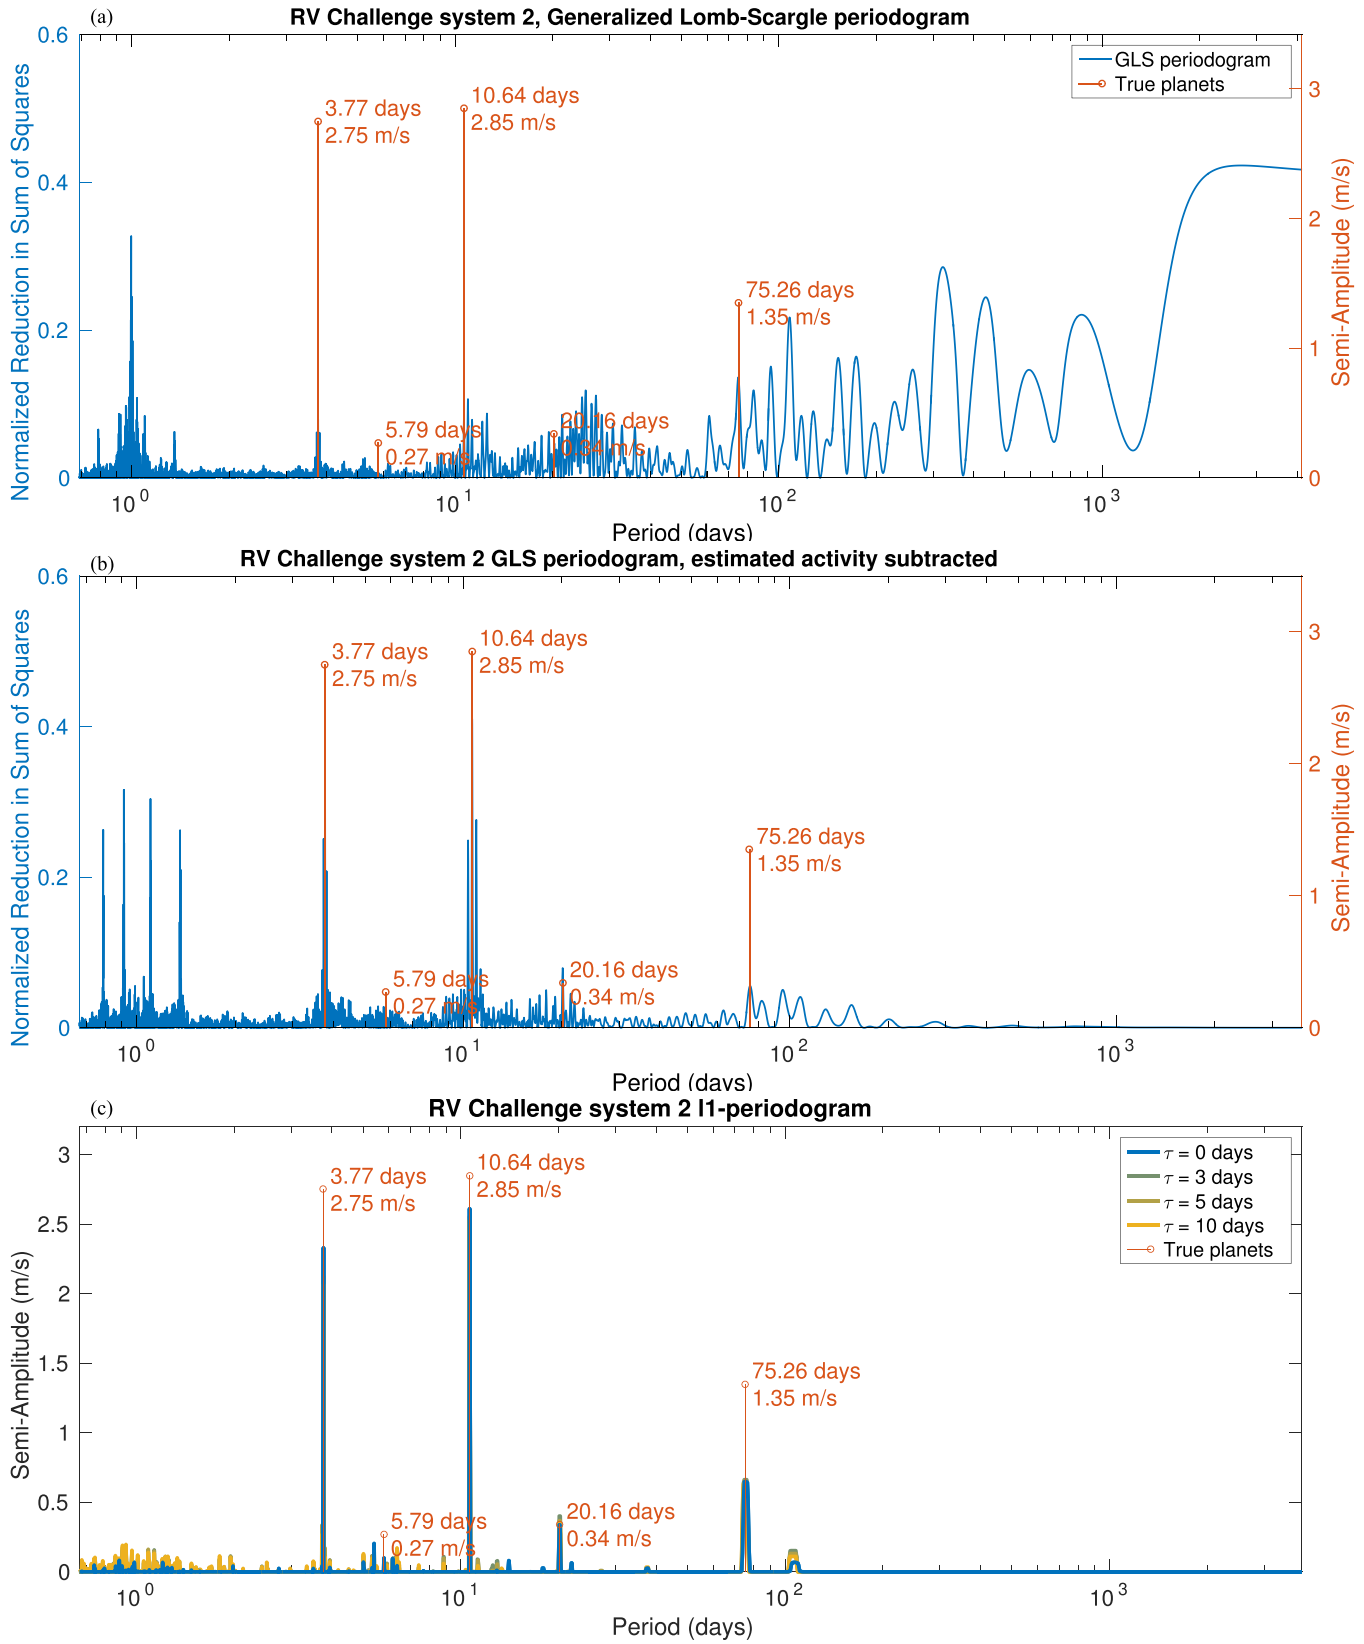

Figure 5. Top: GLS of the RV Challenge system 1 (simulated signal). Top: GLS of raw data. Middle: GLS after fitting ancillary measurements. Bottom: ℓ_1-periodogram after fitting ancillary measurements. True planets are represented by red lines.

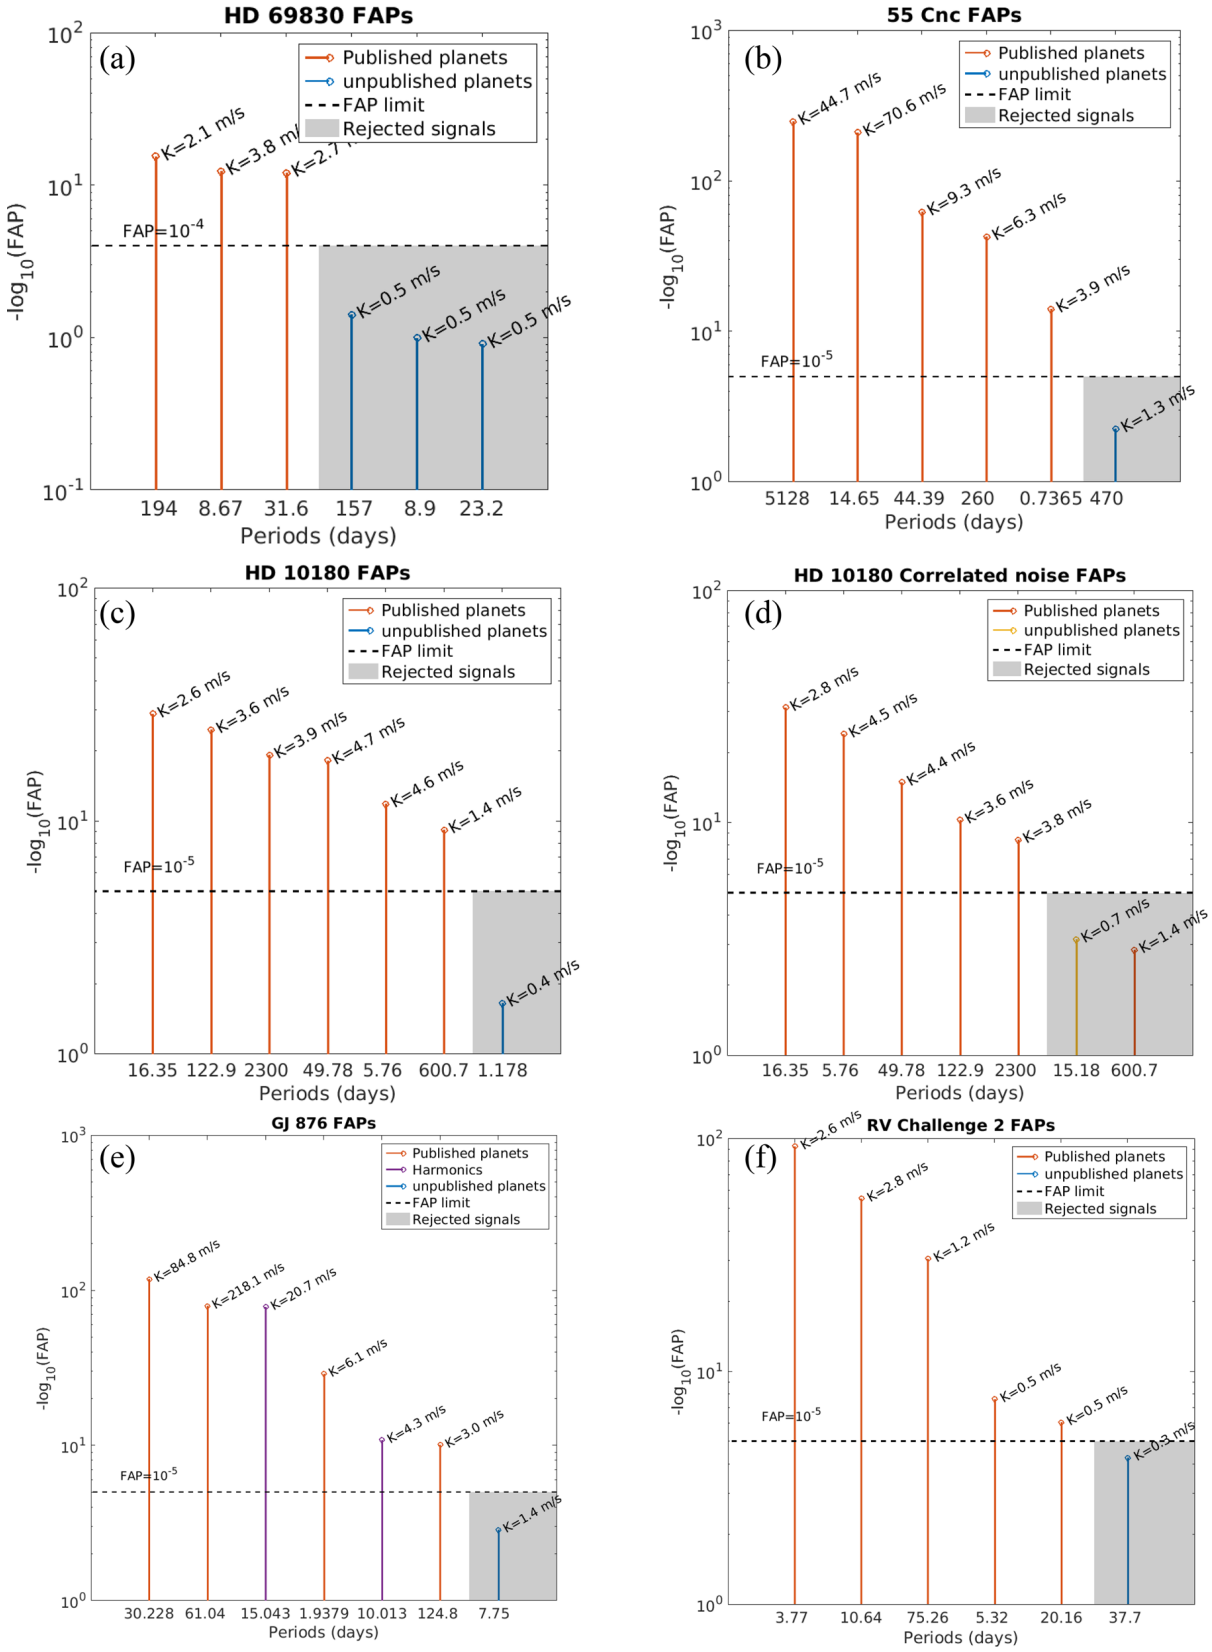

Figure 6. Peak amplitudes and associated FAPs for the four systems analysed.

1.6 m s^{-1}. The host star is a quiet K dwarf with a $\log R'_{HK} = -4.97$ and an estimated projected rotational velocity of $1.1^{+0.5}_{-1.1}$ m s^{-1}; therefore, the star jitter should not amount to more than 1 m s^{-1} (Lovis et al. 2006).

Our method consists in solving the minimization problem (11,BP$_{\epsilon,W}$) and averaging the solution as explained in Section 3.6. The resulting array $x^\sharp(\omega)$ (see equation 19) is plotted versus frequency, here giving Figs 1(b) and (c). The tallest peaks are then fed to a Levenberg–Marquardt algorithm, and the FAPs of models with an increasing number of planets are computed. We represent the FAPs of the signals when fitted from the tallest peaks to the lowest – disregarding aliases – Fig. 6(a). The FAP corresponding to an FAP of 10^{-4} is represented by a dotted line.

The values of most of the algorithm parameters defined in Section 3.5 are fixed in the previous section. In this example, the method is performed for two grid spans: 0–1.5 and 0–0.95 cycles d$_{-1}$ (Figs 1b resp. c).

We first apply the method on a grid spanning between 0 and 1.5 cycles d^{-1}. The weight matrix is diagonal, $W_{kk} = 1/\sigma_k$ (not $1/\sigma_k^2$), where σ_k is the error on measurement k. In Fig. 1(b), the peaks of published planets appear, as opposed to the GLS periodogram (Fig. 1a). However, there are still peaks around 1 d. The three main peaks in that region have periods of 0.9921, 0.8966 and 1.1267. The maximum of the spectral window occurs at $\omega_M = 6.300\,84$ rad d^{-1}. Calculating $2\pi/(\omega - \omega_M)$ yields 194.06, 8.8877 and −8.6759, respectively for $\omega = 2\pi/0.9921$, $2\pi/0.8966$ and $2\pi/1.1267$, suggesting that the short-period peaks are aliases of the true periods.

We now apply the method described in Section 3.7.4 to test the significance of the signal, obtaining Fig. 6(a). Taking 8.667, 31.56 and 197 d gives a reduced χ^2 of the Keplerian fit with three planets plus a constant (16 parameters) of 1.19, yet the stellar jitter is not included. As a consequence, finding other significant signals is unlikely.

Looking only at Fig. 1(b), whether the signal at 197 d or its alias at 0.9921 d is in the signal is unsure. We perform two fits with the two first planets plus one of the candidates. The reduced χ^2 with 0.9921 d is 1.2548, suggesting that the planet at 197 is indeed the best candidate.

Now that there are arguments in favour of a white noise and three planets, let us examine what happens when using a red noise model. The frequency span is restricted to 0–0.95 cycles d^{-1} to avoid spurious peaks (Fig. 1c). As said above, the star is expected to have a jitter in the m s^{-1} range, so we take for the additional jitter $\sigma_W = 0$, $\sigma_R = 1$ m s^{-1} and try several characteristic correlation time lengths $\tau = 0$, 3, 6, 10 or 20 d with definitions of equation (14). In that case, as said in Section 2.3, the estimation of the power is expected to be biased. Fig. 1(c) shows that the peaks at high and low frequencies are respectively overestimated and underestimated. We suggest the following explanation: the weighting matrix accounts for red noise that has more power at low frequencies. Therefore, the minimization of equation (5) has a tendency to 'explain' the low frequencies by noise and put their corresponding energy in the residuals.

When the signal is more complicated, there might be complex effects due to the sampling resulting in a less simple bias. This issue is not discussed in this work, but we stress that when using different matrices W, the tolerance ϵ must be tightened to avoid being too affected by the bias on the peak amplitudes.

To illustrate the advantages of our method, in Appendix C, we generate signals with the same amplitude as the ones of the present example but with periods and phases randomly selected. We show that the maximum of the GLS periodogram does not correspond to a planet in ≈7 per cent of the cases, while the maximum peak of the ℓ_1-periodogram is spurious in less than 0.5 per cent of the cases.

4.3 HD 10180

Lovis et al. (2011) suggested that the system could contain up to seven planets based on 190 HARPS measurements, whose nominal error bars are between 0.4 and 1.3 m s^{-1}. The star has an activity index $\log R'_{HK} = -5$ which lets suppose an inactive star with low jitter. In Lovis et al. (2011), the presence of the planets at 5.79, 16.35, 49.74, 122.7, 600 and 2222 d is firmly stated. Let us mention that there is a concern on whether a planet at 227 d could be in the signal instead of 600 d, as they both appear on the periodogram of the residuals and $1/227 - 1/600 + 1/365 \leq 1/T_{\rm obs}$, where $T_{\rm obs}$ is the total observation time. The possibility of the presence of a seventh planet is also discussed. After the six previous signals are removed with a Keplerian fit, the tallest peaks on the periodogram of the residuals are at 6.51 and 1.178 d (Lovis et al. 2011). They are such that $1/6.51 + 1/1.178 - 1 \leq 1/T_{\rm obs}$, so one is probably the alias of the other. The dynamical stability of a planet at 1.17 d is discussed in Laskar, Boué & Correia (2012), and its ability to survive is shown. However in our analysis, the statistical significance is too low to claim that the planet is actually in the system.

We compute the ℓ_1-periodogram for a grid span of 0–1.5 and 0–0.95 cycles d^{-1}, giving respectively Figs 2(b) and (c) (blue curve). In Appendix B, we show that when W correctly accounts for the red noise, signals might become apparent. Therefore, on the latter we also test different weight matrices. As explained in Appendix B and previous section, in that case we have to decrease $\epsilon_{\rm noise}$ and here $F_{\chi^2_m}(\epsilon^2_{\rm noise}) = 0.1$ was taken, where $F_{\chi^2_m}$ is the CDF of the χ^2 distribution with m degrees of freedom, m being the number of measurements, in accordance with the notations of Section 3.5. We note that there is a signal appearing at 15.2 d and that there is a small peak at 23 d, which is close to the stellar rotation period estimate of 24 d (Lovis et al. 2011). Whether this is due to random or not is not discussed here.

Alike the case of HD 69830, the aliases are overestimated when the frequency span is 3 cycles d^{-1}. In that case, the highest one at 0.9976 d corresponds to an alias of the 2222 d period. We will see that in the two next systems the aliases are not as disturbing, which is discussed in Section 5.2.

We now need to evaluate the significance of the peaks. The FAP test is performed for the seven highest signals, which are the published planets plus 0.177 d or 15.2 d. The latter appears for a non-diagonal weight matrix $\mathbf{W}$; therefore, when performing a Keplerian fit, the χ^2 we take is $(y(t) - \hat{y}(t))^T W^2 (y(t) - \hat{y}(t))$ with the same W, that is $\sigma_W = 0$, $\sigma_R = 1$ m s^{-1} and τ =25 d [with notations of equation (14)]. This analysis gives Figs 6(c) and (d). In both cases, the signals are below the significance threshold. It is also not clear which seventh signal to choose (Fig. 2c), but doing the analysis with other candidates as 6.51, 23 or 67.5 d does not spot significant signals either. Let us note that when choosing a non-diagonal W, the FAPs of the 16.4 and 600 d planets respectively increase and decrease. We suggest the following explanation: the noise model is compatible with noises that have a greater amplitude at low frequencies. As a consequence, the minimization has a tendency to interpret low frequencies as noise and 'trust' higher frequencies. Deciding if a signal is due to a low-frequency noise or a true planet could be done by fitting the noise and the signal at the same time.

4.4 55 Cancri

4.4.1 Data set analysis

Also known as ρ Cancri, Gl 324, BD +28°1660 or HD 75732, 55 Cancri is a binary system. To date, five planets orbiting 55 Cancri A (or HR 552) have been discovered. The first one, a 0.8 $M_{\rm J}$ minimum mass planet at 14.7 d was reported by Butler et al. (1997). Based on the Hamilton spectrograph measurements, Marcy et al. (2002) found a planet with a period of approximately 5800 d and a possible Jupiter mass companion at 44.3 d. With the same observations and additional ones from the Hobby–Eberly Telescope (HET) and ELODIE, McArthur et al. (2004) suggested that a Neptune-mass planet could be responsible for a 2.8 d period. Wisdom (2005) re-analysed the same data set and found evidence for a Neptune-sized planet at 261 d and suggested that the 2.8 period is spurious. This was confirmed by Dawson & Fabrycky (2010), which showed that the 2.8 d periodicity is an alias and the signal indeed comes from a super-Earth orbiting at 0.7365 d. The transit of this planet was then observed by Winn et al. (2011) and Demory et al. (2011), confirming the claim of Dawson & Fabrycky (2010). In the meantime, using previous measurements and 115 additional ones, Fischer et al. (2008) confirmed the presence of a planet at 261 d of minimum mass $M \sin i = 45.7$ $M_\oplus$. They also point out that in 2004 they observed two weak signals at 260 and 470 d on the periodogram. The constraints on the orbital parameters were improved by Endl et al. (2012) based on 663 measurements: 250 from the Hamilton spectrograph at Lick Observatory, 70 from Keck, 212 from HJST and 131 of the high-resolution spectrograph (HET), giving planets at 0.736 546 $\pm 3 \times 10^{-6}$, 14.651 ± 10^{-4}, 44.38 $\pm 7 \times 10^{-3}$, 261.2 ± 0.4 and 4909 ± 30 d. This is the set of measurements we will work on in this section. Let us mention also that Baluev (2015b) and Nelson et al. (2014) studied respectively 55 Cnc dynamics and noise correlations including additional measurements (Fischer et al. 2008).

Let us consider the set of 663 measurements from four instruments used in Endl et al. (2012). The mean of each of the four data set is subtracted and the method described in Section 2 is applied straightforwardly. Here we only display the figure obtained for a white noise model as it is essentially unchanged when correlated noise is taken into account. Fig. 3(b) shows the ℓ_1-periodogram, and Fig. 3(c) is the same figure with a smaller y-axis range. The published signals appear without ambiguity. This is somewhat surprising, as the data come from four different instruments and their respective mean was subtracted. Such a treatment is rather crude, so it shows that at least in that case the method is not too sensitive to the differences of instrumental offsets. When those are fitted with the planets found and corrected, a 365 d periodicity clearly appears on the ℓ_1-periodogram.

The FAPs computed following the method outlined in Section 3.7.4 are significant (see Fig. 6b). The sixth highest peak is at 470 d, the FAP of which is too low to claim a detection. Interestingly enough, a signal at this period was mentioned by Fischer et al. (2008). We will see in the next section that this one is already seen in 2004, and probably due to the different behaviour of the instruments at Lick and HET. The presence of a signal at 2.8 and 260 d in early measurements is also discussed.

4.4.2 Measurements before 2004: no planet at 2.8 d nor 470 d but visible 55 Cnc e and f

The 55 Cnc system has several features that are interesting to test our method. There has been some false detections at 2.8 d, and among candidate signals, one was confirmed (260 d) and one was not (470 d). We now have at least 663 reliable measurements that are very strongly in favour of five planets. As a consequence, the method can be applied on a shortened real data set with specific questions in mind, while being confident about what really is in the system. We will see that the use of the ℓ_1-periodogram could have helped detecting the true planets based on the 313 measurements considered in McArthur et al. (2004). These ones are from Hamilton spectrograph at the Lick Observatory, the HET and ELODIE (Observatoire de Haute Provence). We also show that the signal at 0.7365 d (55 Cnc e) was detectable on the separate data sets from Lick or from HET available in 2004.

Our method is first applied to the three data sets at once, the means of which were subtracted, which gives the lighter blue curve in Fig. 7(a). The true periods appear, although the 260 period is very small and there are peaks at 470, 1314 and 2000 d (the other features of the figure will be explained later). We then consider the three data sets separately; the figure obtained is displayed in Fig. 7(b). The fact that the ℓ_1-periodograms of each of the three instruments span on different length is due to the fact that they do not have the same observational span. As the moving average on the result of SPGL1 is $2\pi/3T_{\rm obs}$, it is wider when the total observation time $T_{\rm obs}$ is small. The 14.65 and long periods are seen for each data set, but the 0.7365 and 44.34 d periodicities are not seen for the ELODIE data set. Interestingly, HET ℓ_1-periodogram displays a periodicity close to 260 d. However, one cannot claim a detection at this period in HET data, as those only span on 180 d, any period longer than the observation timespan is very poorly constrained. Furthermore, the period at 2.8 d is not seen in any data set. The closest one would be a peak at 2.62 d obtained with ELODIE data, which was checked not to be significant. The 470 d periodicity does not appear either. We show in the next paragraph that this is likely due to the velocity offset between Lick-Hamilton and HET data sets. Let us point that CLEAN (Roberts et al. 1987) or frequency analysis (Laskar 1988; Laskar et al. 1992, see Fig. 8) also allows us to retrieve the 0.7365 periodicity, which basically means that the strongest peak of the residual was already this one in 2004.

To compute the significance, the method of Section 3.7.4 is applied to the Lick and HET data separately. The FAPs are computed for circular models with an increasing number of planets whose periods correspond to the subsequent tallest peaks of the ℓ_1-periodogram. Here, as the data come from different instruments we add to the model three vectors $1_{\rm Lick}(t)$, $1_{\rm Elodie}(t)$ and $1_{\rm HET}(t)$, where $1_I(t) = 1$ if the measurement at time t is made by instrument I, $1_I(t) = 0$ otherwise. In the case of Lick data, there is a peak of 6 m s^{-1} at 1.0701 d, but this one can be discarded as it is an alias of the 14.65 d periodicity. In both HET and Hamilton data, the 0.7365 periodicity is significant (Figs 7c and d). Also, one sees a significant long period in both cases (respectively 8617 and 5212 d). The HET data set spans on 170 d, so in this case one can only guess that there is a long-period signal. Finally, when combining the two data sets, the 470, 2150 and 1314 d periodicities become insignificant.

The difference in zero-points of the three instruments has a signature on the ℓ_1-periodogram. Indeed, in problem (11,$\mathrm{BP}_{\epsilon,W}$), the signal is represented as a sum of sinusoids. The algorithm could then attempt to 'explain' the bumps in velocity that occur when passing from one instrument to the other by sines. The previous analysis ensures the presence of four periodicities in the signal: at ≈14.65 d, 44.34, 5000 and 0.7365 d. The fit with these four periods plus the vectors $\boldsymbol{I_I(t)}$ gives coefficients of the latter $\alpha_{\rm Lick}$, $\alpha_{\rm Elodie}$ and $\alpha_{\rm HET}$. The vector $\alpha_{\rm Lick}\boldsymbol{I}_{\rm Lick}(\boldsymbol{t}) + \alpha_{\rm Elodie}\boldsymbol{I}_{\rm Elodie}(\boldsymbol{t}) + \alpha_{\rm HET}\boldsymbol{I}_{\rm HET}(\boldsymbol{t})$ is subtracted from the raw data. The ℓ_1-periodogram of the residuals

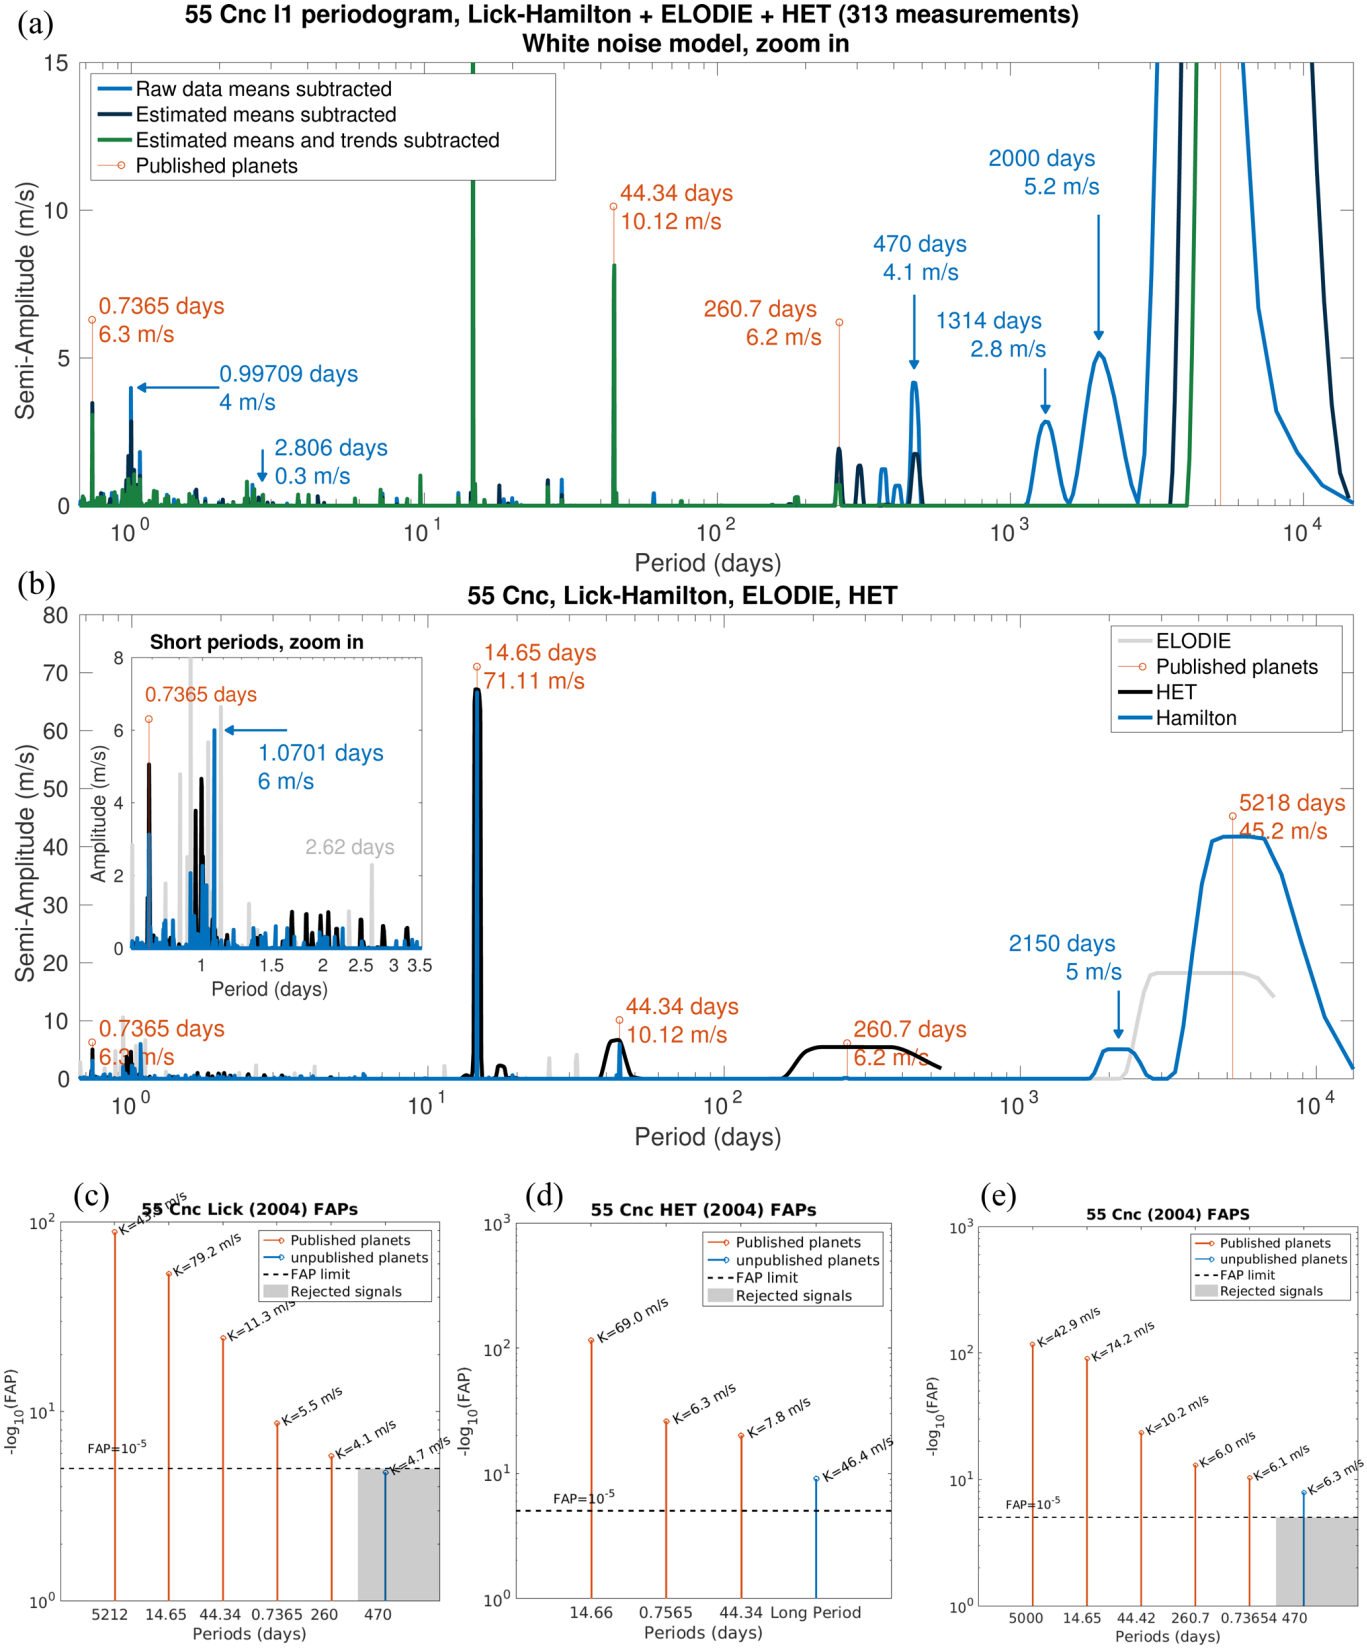

Figure 7. ℓ_1-periodogram of 55 Cnc, using measurements from the Lick-Hamilton, ELODIE spectrograph (Observatoire de Haute Provence) and HET telescope.

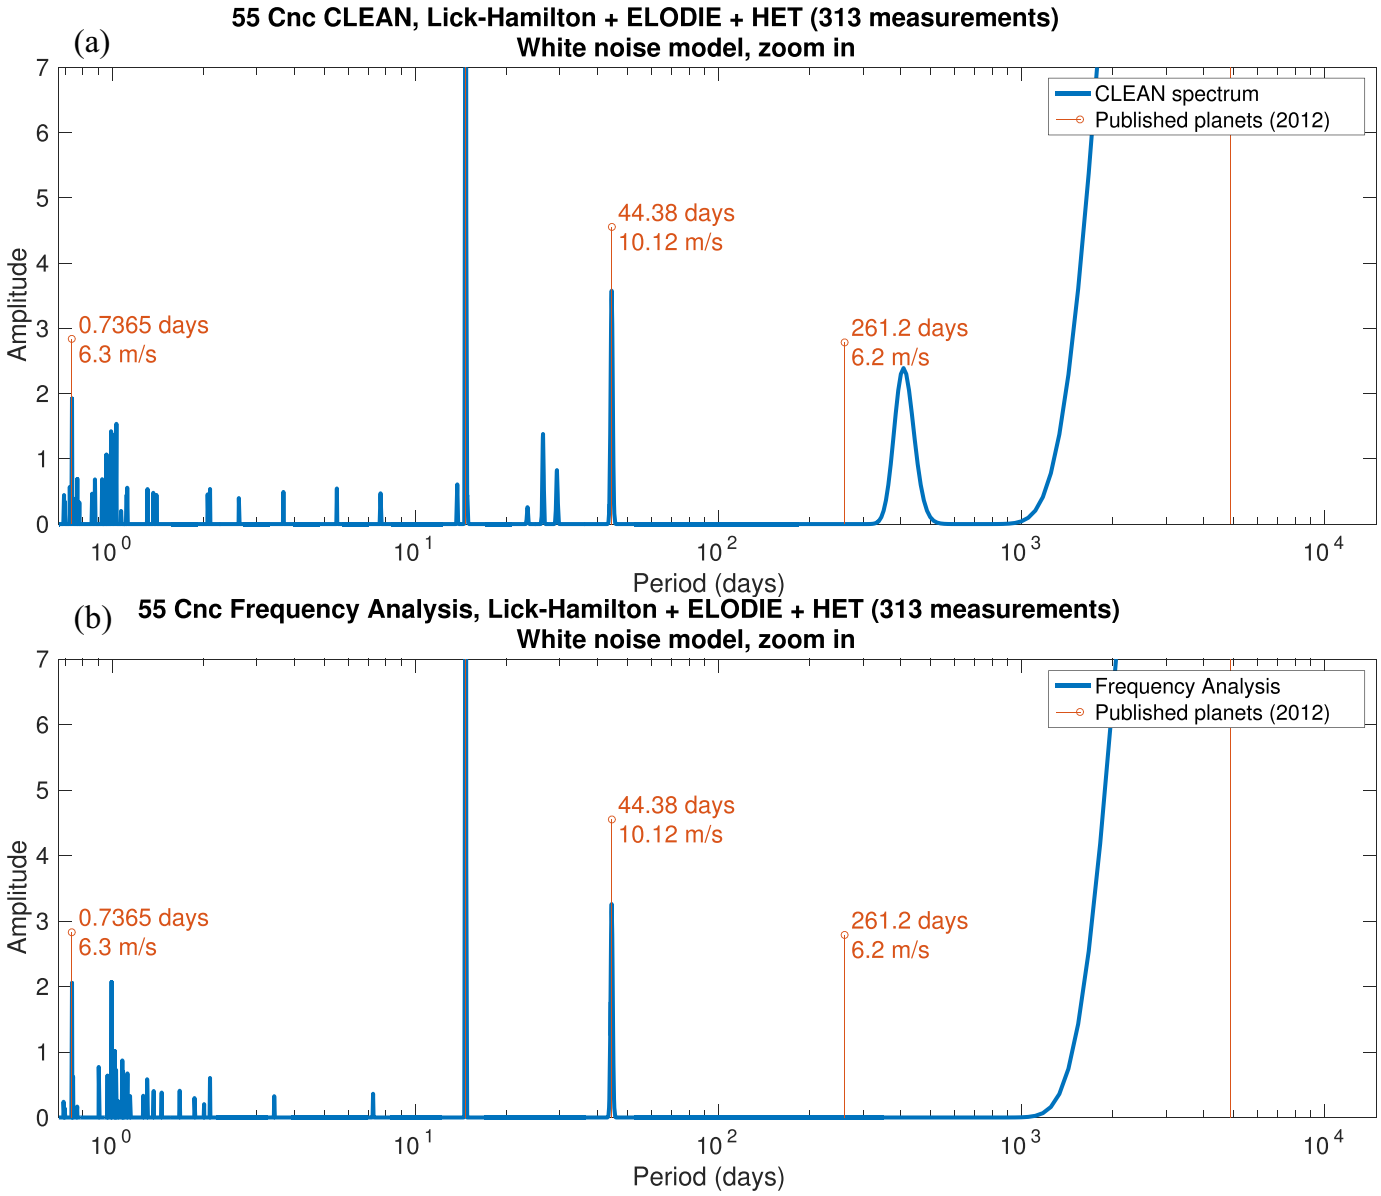

Figure 8. (a) CLEAN spectrum of 55 Cnc with the data available in 2004. (b) Frequency analysis of the same data.

is computed, which gives the dark blue curve shown in Fig. 7(a). The 2000 and 1314 periods disappear and the 470 d peak decreases. Interestingly enough, the fifth tallest peak (except the 0.99 709 d alias) becomes 260 d, which was suggested by Wisdom (2005) and confirmed by Fischer et al. (2008) and Endl et al. (2012), but it does not appear on the CLEAN spectrum nor the frequency analysis (Figs 8a and b).

We now fit the model with five planets along with the $\boldsymbol{1}_I$ vectors and trends for each instrument, which are vectors $\boldsymbol{t}_I$ such that $\boldsymbol{t}_I(t) = t$ and 0 elsewhere if the measurement at time t is done by the instrument I. The vector $\sum \alpha_k 1_{I_k} + \beta_k \boldsymbol{t}_{I_k}$ is subtracted from the raw data, and we compute again the ℓ_1-periodogram (Fig. 7a, green curve). This time, the 470 d periodicity disappears, suggesting – though not proving – that it is due to a difference in behaviour between the instruments. The fact that the 470 d signal disappears just shows its presence depends on the models of the instruments. The same analysis on Lick and HET data altogether shows the same features at 470 d; therefore, we exclude the possibility that it is due to the lesser precision of ELODIE.

The analysis by Wisdom (2005) does not use ℓ_1 minimization to unveil the 260 d periodicity (55 Cnc f). We tried to reproduce a similar analysis 'by hand' on the same data set, namely the one of McArthur et al. (2004). The rationale is to determine if it was easy to make 55 Cnc f appear with an analysis more conventional than the ℓ_1-periodogram. Also, the short-period planet can be injected at 0.7365 d, not ≈2.8 d as it was then. We found that the size of the peak in the residuals at 260 d depends on the initialization of the fits, both with classical and recursive periodograms. While in most cases the 260 periodicity does appear in the residuals, it sometimes coexists with peaks of similar amplitude. Interestingly enough, an analysis of Lick-Hamilton and HET data sets by recursive periodograms suggests that the periods estimated by HET are shifted to longer ones with respect to Lick ones. We found that adding the periods 14.8, 15 000 ($1/14.65 - 1/14.8 \approx 1/5000 - 1/15\,000$) to those of the four planets and a 2500 one (probably due to a harmonic of the 5000 d periodicity) makes the 400 (seen on the CLEAN spectrum in Fig. 8a) and 470 periodicity disappear, and the 260 d peak appears very clearly. As the data come from

an older generation of spectrographs, one could expect complicated systematic errors. Again, this discussion focuses on the possibility of seeing the 55 Cnc f in 2004; we do not raise the question of its existence, well established by the subsequent measurements.

Finally, we perform the FAP test on the data from the three instruments (see Fig. 7e). The model is made of Keplerians plus the 1_I vectors. The four significant signals in each data set are still significant. The 260 d periodicity is significant as well. This analysis shows that both the 0.7365 and 260 d periodicities were already present in the data. Long periods might be due to instrumental effects; therefore, the planetary origin of the 260 period could have been subject to discussion. In contrast, it seems hard to explain a steady 0.7365 d periodicity with a non-planetary effect. Let us also note that our χ^2 minimization algorithm and the one used by Wisdom (2005) are different, the latter being more sophisticated, which might account for discrepancies between our analyses.

4.5 GJ 876

4.5.1 Previous work

The GJ 876 host star is one of the first discovered multiplanetary systems. First, two giant planets at 30 and 61 d were reported by Marcy et al. (1998) and Delfosse et al. (1998). Subsequently, Rivera et al. (2005) find a short-period Neptune at 1.94 d and a Uranus-mass planet at 124 d (Rivera et al. 2010).

The giant planets are close to each other and in 2:1 resonance, therefore we might expect visible dynamical effects. Indeed, Correia et al. (2010), Baluev (2011) and Nelson et al. (2016) perform four-body Newtonian fits which give a χ^2 of the residuals smaller than a Keplerian fit. The dynamical fits also allow one to have constraints on the inclinations, therefore on the true masses of the planets. Furthermore, Baluev (2011) shows that the maximum of a posterior likelihood including a noise model as the one used here (equation 14) occurs at $\sigma_W = 1.31$ m s^{-1}, $\sigma_R = 1.8$ m s^{-1} and $\tau = 3$ d.

Jenkins et al. (2014) take a different approach and search for sine functions in the signal. They claim that six significant sinusoidal signals are in the data. The following discussion first confirms these results. Secondly, we investigate the origins of the additional two signals and find that they are likely to be due to the interactions between the giant planets.

4.5.2 Six significant sines

Jenkins et al. (2014) analyse the GJ 876 data by aiming at solving the problem (1), which they call minimum mean squared error. To do so, the phase space is explored with an iterative arborescent method. They find the following periods: 61.03±3.81, 30.23±0.19, 15.04±0.04, 1.94±0.001, 10.01±0.02 and 124.69±90.04 d. To compare our results with Jenkins et al. (2014), the significance of the signals is tested with FAPs as previously. We use different weight matrix models according to equation (14) and two grid spans: 1.5 and 0.95 cycles d^{-1} (see Figs 4b and c). In Fig. 4(c), we see that the six tallest signals correspond to the periods we expect. Depending on the noise model, the seventh tallest peak varies. We compute the FAP test for 7.748, 1200 or 4200 d as candidate seventh planets, respectively, with the $\mathbf{W}$ matrix yielding their greatest amplitude. In Fig. 6(e), we display the result for 7.748 d but in other cases the signals are not significant. Let us still point out that in the case of $\tau = 6$ d, initializing a 4200 d periodicity, after the non-linear fit we obtain a 4862 d periodicity which has an FAP of 0.0007. This one is close to the total observation timespan (4600 d). Therefore, it is hard to determine what could be its cause.

Before discussing the origin of these signals, we wish to comment on the behaviour of the ℓ_1-periodogram towards the 124 d periodicity. Indeed, in the case of the 1.5 cycles d^{-1}, this one has the same order of magnitude as the tallest alias in the 1 d region (at 0.9812 d, alias of the 61 d periodicity). Furthermore, the peak becomes visible only for non-diagonal weight matrix $\mathbf{W}$, while a white noise model is sufficient to see it when using a shorter grid (Fig. 4c). To understand this feature, we argue as follows. There are three effects against finding the correct planets: the red noise (Baluev 2011), the uncertainties on the two instrumental means and the inner faults of our method. The persistence of aliases at 1 d indeed shows that the recovery of the true signals is more difficult when considering a grid Ω where some of the frequencies are very correlated. We also computed the ℓ_1-periodogram when the mean of each instrument is corrected after the orbital parameters fit, as done in Section 4.4.2. In that case, the 124 d periodicity does appear and the aliases are reduced. We suggest the following explanation: when at least one of the three obstacles is correctly taken into account, the method is sufficient. When the three are ignored, their joint effect is deadly to our ability to recover the correct planets.

4.5.3 Signals at 10 and 15 d

Now that the six sines are seen in the signal, we show that the peaks at 15.06 and 10.01 d are due to the dynamical interactions.

We perform the same four-body fit of GJ 876 with the same method as Correia et al. (2010). This one includes 25 parameters: the mass of the star, a velocity offset, the mass of the planets, for the smallest planets: period, semi-amplitude, eccentricity, argument of periastron and initial mean anomaly. For the giant planets at 30 and 61 d, the inclination is also a free parameter.

A planetary system with the orbital elements found by the least-squares fit is simulated on 100 years for the two giant planets and the four planets at once. The frequency analysis (Laskar 1988, 1993; Laskar et al. 1992) is then performed on the resulting time series of the star velocity along the x-axis. We find that 15.06 and 10.01 periods appear and are a combination of the fundamental frequencies. Denoting by ω_P the frequency of a planet of period P, we have $\omega_{15} = 3\omega_{30} - 2\omega_{60}$ and $\omega_{10} = 5\omega_{30} - 4\omega_{60}$, both in the two-planet and four-planet cases. We also performed another test: if we adjust the two giant planets with a dynamical fit, then the peaks at 15.06 and 10.01 d are not seen on the residuals. This agrees with the analysis of Nelson et al. (2016), where they discuss the possibility that the signals at 10.01 and 15.06 d could be due to additional planets, and find it unlikely. They compute the evidence ratio of Newtonian models with four and five planets, $\Pr\{\boldsymbol{y}|5 \text{ planets}\} / \Pr\{\boldsymbol{y}|4 \text{ planets}\}$, and find that it is not higher than the threshold we chose. The difference between ω_{15} and the first harmonic of the planet gives an estimate of the frequency of precession of the periastron of the inner orbit; we find $2\pi(1/\omega_{15} - 2/\omega_{30}) \approx 8.77$ yr, which is consistent with the estimate of Correia et al. (2010, $g_2 = 8.73$ yr, table 4).

To obtain the expressions of ω_{15} and ω_{10}, we used frequency analysis. This could be puzzling as the present work defines a method to retrieve the frequencies in the signal. The rationale is that we do the frequency analysis on a numerical integration; therefore, we have tens of thousands of points available. Frequency analysis has been used in that situation for years and is known to be fast and robust. We double checked the results by computing the ℓ_1-periodogram on a thousand points from the simulation (handling as many as the frequency map analysis is too long for now); the periods at 15.06 and 10.01 do appear very clearly.

4.6 Very active star (simulated signal)

The examples above concern rather quiet stars, where the noise can be modelled by Gaussian time series. However, in some cases, the stellar activity does not have a known Gaussian signature. The method described here is not yet adapted to handle such situations. In this section, we show that the problem can be circumvented, provided there are enough measurements.

We exploit the fact that stellar noise can be correlated with the bisector span (Queloz et al. 2001), the full width at half-maximum (FWHM) and the $\log R'_{HK}$. This correlation has been used for example in Meunier, Lagrange & De Bondt (2012), which shows that the detection threshold limit improves by an order of magnitude by testing the correlation between the radial velocity and ancillary measurements. They compute the correlation of the periodograms of radial velocity measurements and bisector span, but a correlation in the frequency domain is also visible in the time domain, as the Fourier transform contains the same amount of information as the original time series. Here we take an approach similar to Melo et al. (2007), Boisse et al. (2009) and Gregory (2016) in so far as we use the ancillary measurements as proxies for estimating the activity induced signal. Here, we simply fit and remove the three ancillary measurements from the data and then use the method described above on the residuals. To compute the FAP, we use a model of the form $A\mathrm{FWHM} + B\mathrm{bisector} + C \log R'_{HK} + \mathrm{Circ}(k, h, P, D, E)$, Circ denoting a circular model as defined in Section 3.7.3. The validity of this approach is discussed in Appendix D.

The data set used is taken from the RV Fitting Challenge (Dumusque 2016; Dumusque et al. 2016). In this challenge, 15 systems were simulated with a red noise component taken from observations of real stars plus activity simulated via SOAP 2 (Dumusque, Boisse & Santos 2014). Here we consider the system number two of the challenge. The data set is made of 492 measurements, and the mean precision is 0.67 cm s^{-1}. The first step of the processing is to fit a linear model made of the ancillary measurements, an offset, a linear and a quadratic trend (six parameters). Secondly, we compute the ℓ_1-periodogram for different weight matrices, which gives Fig. 5(c). The GLS periodogram is also computed before and after the fit of the six parameters for comparison (Figs 5a and b).

We find without ambiguity the three planets whose semi-amplitude is above 1 m s^{-1}, and also the 20.16 d periodicity. The planet with the smallest amplitude does not appear clearly, but there is a peak at 5.4 d which seems to be significant. In fact, the spectral window is such that 5.4 d is an alias of $5.32 = 10.64/2$ d, and corresponds to the first harmonic due to eccentricity. This feature seems to be due to an error in the noise model. When accounting for a red noise effect, the relative amplitudes of 5.32 and 5.4 change in favour of 5.32 d. This effect is also observed on the recursive periodograms which are not represented here for the sake of brevity. One can see a peak at 6.25 d which grows stronger as the characteristic correlation time of the noise model increases. This coincides with the fourth harmonic of the rotational period and is therefore not surprising.

5 DISCUSSION

5.1 Summary

The present work was first devised to overcome the distortions in the residual that arise when fitting planets one by one. It is compatible with the assumption that the noise is Gaussian and correlated through the weighting matrix $\mathbf{W}$. One of the main advantages of the method is that, as opposed to global χ^2 minimization, the minimization problem (11,$\mathrm{BP}_{\epsilon,W}$) is convex therefore quicker to solve. On our workstation (Intel Xeon CPU E5-2698 v3 at 2.30 GHz), it takes typically 30 s to 10 min to obtain (resp. for HD 69830, 74 measurements and 55 Cnc, 663 measurements). The speed here depends mainly on three parameters: the number of observations m, the number of columns of matrix $\mathbf{A}$ (see Section 3.3), n, and the precision wanted in output, tol (see Section 3.2). The SPGL1 algorithm used to solve equation (11,$\mathrm{BP}_{\epsilon,W}$) relies on a Newton algorithm; therefore, its complexity is $O(\log(p)F(p))$, where $p = 10^{-\mathrm{tol}}$ is the number of significant digits desired and $F(p)$ the cost of evaluating the objective function to p digits accuracy. The most expensive steps of the evaluation are a matrix vector product and a projection on to a convex set (see van den Berg & Friedlander 2008), which have a respective complexity of $O(mn)$ and a worst case complexity of $O(n \log n)$. The post-processing operation also is in $O(mn)$. This overall should amount asymptotically to complexity $O(mn)$, similarly to the Lomb–Scargle periodogram. Its complexity is in $O(mn)$ if there are m measurements and n frequency scanned. The constants are however different.

Furthermore, our method does not require the number of planets as input parameter and offers a graphic representation of the information content of the signal. However, the statistical properties of the solution are not as easy to interpret as in the case of a global least-squares minimization. Considering that the method presented here is in its infancy, comparing its merits to other techniques is left for future work. Here, we will only stress that the ℓ_1- and GLS periodogram are tools are of different levels, and we do not advocate to give up the latter.

We will confine ourselves to addressing some internal issues of our method. Ultimately, we would like to know if there is a way to determine which peaks are to be associated with planets. As the present paper is concerned with unveiling the periodicities in the signal but not their origins, we will address a simpler question: assuming that the signal is only made of sines plus a Gaussian noise, are there risks to see spurious peaks on the ℓ_1-periodogram?

Unfortunately, the answer is yes, as we have seen in the previous examples. The method is in particular sensitive to the aliases due to the daily repetition of the measurements: spurious peaks are especially present around 1 d periods. To shed some light on this problem, the following questions will be briefly discussed in the two next sections.

(i) Are spurious peaks to be expected from the theoretical properties of the method or from its implementation?

(ii) If they are to appear anyway on the ℓ_1-periodogram, is there a way to spot them?

5.2 Mutual coherence

To test if the algorithm behaves appropriately, we reason as follows. Considering a set of observational times $\boldsymbol{t} = t_1, \ldots, t_m$, a linear combination of p pure sine signals $y(t_k) = a_1 \cos(\omega_1 t_k + \phi_1) + \cdots + a_p \cos(\omega_p t_k + \phi_p)$ is generated with uniformly distributed phases ϕ and various amplitudes. For any tolerance ϵ, the SPGL1 algorithm must give a solution $\boldsymbol{x}^\star$ (see equation (11,$\mathrm{BP}_{\epsilon,W}$)) such that $\|\boldsymbol{x}^\star\|_{\ell_1} \leqslant |a_1| + \cdots + |a_p|$, as obviously $\boldsymbol{y}(\boldsymbol{t})$ belongs to the set of signals u verifying $\|\boldsymbol{u} - \boldsymbol{y}(\boldsymbol{t})\|_{\ell_2} \leqslant \epsilon$. To test if SPGL1 gives the best solution, we take the measurement dates of HD 69830 and generate three pure cosine functions of one amplitude whose frequencies are in the grid. They are fed to the SPGL1 solver for $\epsilon = 0.01$ and W equal to the identity matrix. The solution $x^\star$ to equation (11,$\mathrm{BP}_{\epsilon,W}$) must verify

Table 1. Maximum amplitude of the spectral window in the 1 cycle d^{-1} and 1 cycle yr^{-1} for the examples considered here.

	≈1 cycle d^{-1}	≈1 cycle yr^{-1}
HD 69830	0.926	0.600
HD 10180	0.949	0.703
55 Cnc	0.822	0.557
GJ 876	0.732 46	0.501
RV Challenge 2	0.870	0.800

$\|\boldsymbol{x}^\star\|_{\ell_1} < 3$ as the original signal is not noisy. The test is performed for a thousand set of three frequencies randomly selected on the grid. We find that the average ℓ_1 norm of the solution is 3.26, suggesting that the algorithm could be improved.

Secondly, in the discrete case (problem 5), there are theoretical guarantees on the success of the recovery if the mutual coherence of the dictionary is sufficiently small (Donoho 2006). This one is defined as the maximal correlation between two columns a_j and a_k of the dictionary A,

$$\mu = \max_{\substack{k=1,\dots,n \\ j=1,\dots,n \\ j\neq k}} |\langle \boldsymbol{a_k}, \boldsymbol{a_j}\rangle|. \tag{30}$$

In the case of a dictionary such that $\boldsymbol{a}_k = \mathrm{e}^{\mathrm{i}\omega_k t}$, taking the convention $\langle \boldsymbol{a}_k, \boldsymbol{a}_j\rangle = \boldsymbol{a}_k^* \boldsymbol{a}_j$, where the superscript $*$ denotes the conjugate transpose,

$$|\langle \boldsymbol{a}_k, \boldsymbol{a}_j\rangle| = \left|\sum_{l=1}^{m} \mathrm{e}^{-\mathrm{i}(\omega_k-\omega_j)t_l}\right| \tag{31}$$

that is the spectral window in $\omega_k - \omega_j$. As a consequence, the method cannot resolve very close frequencies due to their high correlation. More importantly, aliases are still a limitation – though not as much as in iterative algorithms in general (Donoho et al. 2006), see also Appendix C. This feature is responsible for the aliases that still appear around 1 d, where there is generally a strong alias due to the sampling constraints. The problem tends to get worse as the maxima of the spectral window increase. Aliases are higher relative to the true peaks for HD 69830, HD 10180 and the separate sets of 55 Cnc than GJ 876 (see Figs 1, 2, 3, 4, 7 and Table 1).

5.3 Spotting spurious peaks

We know that the theoretical obstacle for a good recovery is correlation between the elements of the dictionary. If a frequency ω_0 truly is in the signal, it is expected to cause significant amplitudes at $\omega_0 + \omega_k$, where the ω_k are maxima of the spectral window. So if two peaks at frequencies ω_1 and ω_2 are seen on the ℓ_1-periodogram and the spectral window has a strong local maximum close to $\omega_1 - \omega_2$, one can suspect that one of the two peaks is spurious.

5.4 When to use the method?

We consider the general problem of finding the frequencies of a signal made of several harmonics (the multitone problem). It seems natural – though not mandatory – to try to find the global minimum for a given number of sinusoids, and possibly additional parameters such as the offset or a trend. We do not know a priori the number of sinusoids in the signal. Ideally, we would like to solve the global minimization (1) for any number of sines inferior to the number of measurements and regarding their amplitudes, which seem to be truly in the signal. The approach consisting in using grids has a computational cost growing exponentially with the number of frequency. Therefore, strategies must be found to estimate a reliable solution to this problem. The recursive periodogram (Anglada-Escudé & Tuomi 2012), the treillis approach (Jenkins et al. 2014) or the super-resolution methods (Candès & Fernandez-Granda 2014; Tang et al. 2013b) can be viewed as a way to approximate equation (1) and selecting the relevant number of frequencies at the same time. These have the advantage of not being bothered by the ℓ_1 norm minimization, which biases downwards the amplitude of the signal. Even more, the bias becomes more complicated when using a correlated noise model.

The most interesting use of the ℓ_1-periodogram seems to be as a complement to the classical periodogram: it gives a much clearer idea of the number of spikes and their significance. If the peaks spotted by the ℓ_1-periodogram yield a χ^2 of the residuals consistent with the noise assumptions as in HD 69830, then it is likely that there are not many more signals. To check that there are not very high correlations between signals, one can use the spectral window. Furthermore, we have exhibited in Appendix C1 examples where the main peak of the classical periodogram is spurious while ℓ_1 minimization (5) avoids selecting the first spurious peak. Such an example was also presented in Bourguignon et al. (2007). Those findings are consistent with the claims of Donoho et al. (2006): the ℓ_1 method is more reliable in general than orthogonal matching pursuit. A failure of the ℓ_1-periodogram is also informative, as shown in Fig. C2 of Appendix C1. If there still is a forest of peaks below a certain amplitude, it might indicate that the signal is noisy, possibly that noise is higher than expected or non-Gaussian. This means that the set of observations requires a more careful analysis. To sum up, the ℓ_1-periodogram can yield an estimation of the difficulty of the system; in some cases, it is a shortcut to random searches and its use decreases the chance of being misled by a spurious tallest peak.

6 CONCLUSION

The aim of the present paper was to produce a tool for analysing radial velocity that can be used as the periodogram but without having to estimate the frequencies iteratively. To do so, we used the theory of compressed sensing, adapted for handling correlated noise, and went through the following steps.

(i) Selecting a family of normalized vectors where the signal is represented by a small number of coefficients.

(ii) Approximating a solution to equation (9); for example, by discretizing the dictionary, and ensuring that the grid spacing is consistent with the noise power (see equation 15) then solving equation (11,$\mathrm{BP}_{\epsilon,W}$) with SPGL1 and take the average power. The introduction of the weight matrix $\mathbf{W}$ accounts for correlated Gaussian noises.

(iii) Estimating the detection significance, which we do by computing subsequent FAPs of the models with an increasing number of planets.

We showed that the published planets for each system could be seen directly on the same graph, and that taking into account the possible correlations in the noise could make a signal appear. This was established in the case of radial velocity data but the method could be adapted to other types of measurements, such as astrometric observations.

The use of the basis pursuit/ℓ_1-periodogram we suggest is as follows. This method can be used as a first guess to see if the signal is sparse or not; in that extent, it constitutes an evaluation of the

difficulty of the system and possibly a shortcut to the solution. It can bring attention to signal features that are hidden in the classical periodogram, which can still be used for an analysis 'by hand'. Secondly, for confirming the planetary nature of a system, we advocate to use in a second time statistical hypothesis testing.

The perspective for future work is twofold. First, we saw that the algorithm itself could be improved. Also, there might be significance tests more robust than the FAP, and the effect of introducing a weight matrix $\mathbf{W}$ must be studied into more depth. Secondly, let us recall that our method uses a priori information, that is the sparsity of the signal, but still does not handle all the information we have. To improve the technique, we wish to broaden its field of application by

(i) adapting the method for very eccentric orbits, through the addition of Keplerian vectors to the dictionary for example;

(ii) using precise models of the noise, especially magnetic activity, granulation, p-modes. Possibly include an adaptive estimation of the noise, especially one could extend the dictionary to wavelets;

(iii) handling several types of measurements at once (e.g. radial velocity, astrometry and photometry).

ACKNOWLEDGEMENTS

The authors wish to thank the anonymous referee for his insightful suggestions. NH thanks Evgeni Grishin for pointing out the matched filter technique to him. AC acknowledges support from CIDMA strategic project UID/MAT/04106/2013.

REFERENCES

Aigrain S., Gibson N., Roberts S., Evans T., McQuillan A., Reece S., Osborne M., 2011, in AAS/Division for Extreme Solar Systems Abstracts, Vol. 2, p. 11.05
Anglada-Escudé G., Tuomi M., 2012, A&A, 548, A58
Anglada-Escudé G., López-Morales M., Chambers J. E., 2010, ApJ, 709, 168
Arildsen T., Larsen T., 2014, Signal Process., 98, 275
Babu P., Stoica P., 2010, Digit. Signal Process., 20, 359
Babu P., Stoica P., Li J., Chen Z., Ge J., 2010, AJ, 139, 783
Baluev R. V., 2008, MNRAS, 385, 1279
Baluev R. V., 2009, MNRAS, 393, 969
Baluev R. V., 2011, Celest. Mech. Dyn. Astron., 111, 235
Baluev R. V., 2013a, Astron. Comput., 3, 50
Baluev R. V., 2013b, MNRAS, 436, 807
Baluev R. V., 2015a, MNRAS, 446, 1478
Baluev R. V., 2015b, MNRAS, 446, 1493
Becker S., Bobin J., Candès E. J., 2011, SIAM J. Imaging Sci., 4, 1
Bellmann K., 1975, Biometr. Z., 17, 271
Bobin J., Starck J.-L., Ottensamer R., 2008, IEEE J. Sel. Top. Signal Process., 2, 718
Boisse I. et al., 2009, A&A, 495, 959
Bourguignon S., Carfantan H., Böhm T., 2007, A&A, 462, 379
Butler R. P., Marcy G. W., Williams E., Hauser H., Shirts P., 1997, ApJ, 474, L115
Candès E., Fernandez-Granda C., 2013, J. Fourier Anal. Appl., 19, 1229
Candès E., Fernandez-Granda C., 2014, Commun., Pure Appl. Math., 67, 906
Candès E., Romberg J., Tao T., 2006a, IEEE Trans. Inf. Theory, 52, 489
Candès E. J., Romberg J. K., Tao T., 2006b, Commun. Pure Appl. Math., 59, 1207
Chandrasekaran V., Recht B., Parrilo P. A., Willsky A. S., 2010, preprint (arXiv:1012.0621)
Chen Y., Chi Y., 2014, IEEE Trans. Inf. Theory, 60, 6576
Chen S. S., Donoho D. L., Saunders M. A., 1998, SIAM J. Sci. Comput., 20, 33
Cohen A., Dahmen W., Devore R., 2009, J. Am. Math. Soc., 22, 211
Correia A. C. M. et al., 2010, A&A, 511, A21
Cumming A., 2004, MNRAS, 354, 1165
Cumming A., Marcy G. W., Butler R. P., 1999, ApJ, 526, 890
Daubechies I., DeVore R., Fornasier M., Güntürk C. S., 2010, Commun. Pure Appl. Math., 63, 1
Dawson R. I., Fabrycky D. C., 2010, ApJ, 722, 937
Delfosse X., Forveille T., Mayor M., Perrier C., Naef D., Queloz D., 1998, A&A, 338, L67
Demory B.-O. et al., 2011, A&A, 533, A114
Díaz R. F. et al., 2016, A&A, 585, A134
Donoho D., 2006, IEEE Trans. Inf. Theory, 52, 1289
Donoho D. L., Elad M., Temlyakov V. N., 2006, IEEE Trans. Inf. Theory, 52, 6
Duarte M. F., Baraniuk R. G., 2013, Appl. Comput. Harmon. Anal., 35, 111
Dumusque X., 2016, A&A, 593, A5
Dumusque X., Boisse I., Santos N. C., 2014, ApJ, 796, 132
Dumusque X. et al., 2016, preprint (arXiv:1609.03674)
Endl M. et al., 2012, ApJ, 759, 19
Engelbrecht C. A., 2013, in Guzik J. A., Chaplin W. J., Handler G., Pigulski A., eds, Proc. IAU Symp. 301, Precision Asteroseismology. Cambridge Univ. Press, Cambridge, p. 77
Ferraz-Mello S., 1981, AJ, 86, 619
Fischer D. A. et al., 2008, ApJ, 675, 790
Foster G., 1995, AJ, 109, 1889
Ge D., Jiang X., Ye Y., 2011, Math. Program., 129, 285
Gorodnitsky I. F., Rao B. D., 1997, IEEE Trans. Signal Process., 45, 600
Grant M., Boyd S., 2008, in Blondel V., Boyd S., Kimura H., eds, Lecture Notes in Control and Information Sciences, Recent Advances in Learning and Control. Springer-Verlag, Berlin, p. 95
Gregory P. C., 2011, MNRAS, 410, 94
Gregory P. C., 2016, MNRAS, 458, 2604
Horne J. H., Baliunas S. L., 1986, ApJ, 302, 757
Jenkins J. S., Yoma N. B., Rojo P., Mahu R., Wuth J., 2014, MNRAS, 441, 2253
Kay S. M., 1993, Fundamentals of Statistical Signal Processing: Estimation Theory. Prentice-Hall, Englewood Cliffs, NJ
Kay S., Marple S. L. J., 1981, Proc. IEEE, 69, 1380
Laskar J., 1988, A&A, 198, 341
Laskar J., 1993, Celest. Mech. Dyn. Astron., 56, 191
Laskar J., 2003, preprint (arXiv:math/0305364)
Laskar J., Froeschlé C., Celletti A., 1992, Physica D, 56, 253
Laskar J., Boué G., Correia A. C. M., 2012, A&A, 538, A105
Lomb N. R., 1976, Ap&SS, 39, 447
Lovis C. et al., 2006, Nature, 441, 305
Lovis C. et al., 2011, A&A, 528, A112
McArthur B. E. et al., 2004, ApJ, 614, L81
Mallat S. G., Zhang Z., 1993, IEEE Trans. Signal Process, 41, 3397.
Marcy G. W., Butler R. P., Vogt S. S., Fischer D., Lissauer J. J., 1998, ApJ, 505, L147
Marcy G. W., Butler R. P., Fischer D. A., Laughlin G., Vogt S. S., Henry G. W., Pourbaix D., 2002, ApJ, 581, 1375
Melo C. et al., 2007, A&A, 467, 721
Meunier N., Lagrange A.-M., De Bondt K., 2012, A&A, 545, A87
Mishali M., Eldar Y., Tropp J., 2008, in IEEE 25th Convention of Electrical and Electronics Engineers in Israel, IEEEI 2008. p. 290
Mortier A., Faria J. P., Correia C. M., Santerne A., Santos N. C., 2015, A&A, 573, A101
Nelson B. E., Ford E. B., Wright J. T., Fischer D. A., von Braun K., Howard A. W., Payne M. J., Dindar S., 2014, MNRAS, 441, 442
Nelson B. E., Robertson P. M., Payne M. J., Pritchard S. M., Deck K. M., Ford E. B., Wright J. T., Isaacson H. T., 2016, MNRAS, 455, 2484
O'Toole S. J., Tinney C. G., Jones H. R. A., Butler R. P., Marcy G. W., Carter B., Bailey J., 2009, MNRAS, 392, 641
Pati Y. C., Rezaiifar R., Krishnaprasad P. S., 1993, in Proc. of the 27th Asilomar Conference on Signals, Systems and Computers, p. 40
Pelat D., 2013, Bases et Méthodes pour le Traitement de Données. Observatoire de Paris

Queloz D. et al., 2001, A&A, 379, 279
Rajpaul V., Aigrain S., Osborne M. A., Reece S., Roberts S. J., 2015, MNRAS, 452, 2269
Reegen P., 2007, A&A, 467, 1353
Rivera E. J. et al., 2005, ApJ, 634, 625
Rivera E. J., Laughlin G., Butler R. P., Vogt S. S., Haghighipour N., Meschiari S., 2010, ApJ, 719, 890
Roberts D. H., Lehar J., Dreher J. W., 1987, AJ, 93, 968
Rockafellar R. T., 1970, Convex Analysis. Princeton Univ. Press, Princeton, NJ
Scargle J. D., 1982, ApJ, 263, 835
Schuster A., 1898, Terr. Magn., 3, 13
Schwarzenberg-Czerny A., 1998, Balt. Astron., 7, 43
Ségransan D. et al., 2011, A&A, 535, A54
Starck J.-L., Elad M., Donoho D. L., 2005, IEEE Trans. Image Process., 14, 1570
Stoica P., Babu P., 2012, Signal Process., 92, 1580
Sulis S., Mary D., Bigot L., 2016, preprint (arXiv:1601.07375)
Tang G., Bhaskar B., Recht B., 2013a, in 2013 Asilomar Conference on Signals, Systems and Computers, p. 1043
Tang G., Bhaskar B., Shah P., Recht B., 2013b, IEEE Trans. Inf. Theory, 59, 7465
Tibshirani R., 1994, J. R. Stat. Soc. Ser. B, 58, 267
Tropp J. A., Gilbert A. C., 2007, IEEE Trans. Inf. Theory, 53, 4655
Tropp J. A., Laska J. N., Duarte M. F., Romberg J. K., Baraniuk R. G., 2010, IEEE Trans. Inf. Theory, 56, 520
Tuomi M., 2012, A&A, 543, A52
Tuomi M. et al., 2013, A&A, 551, A79
Tuomi M., Jones H. R. A., Barnes J. R., Anglada-Escudé G., Jenkins J. S., 2014, MNRAS, 441, 1545
van den Berg E., Friedlander M. P., 2008, SIAM J. Sci. Comput., 31, 890
Winn J. N. et al., 2011, ApJ, 737, L18
Wisdom J., 2005, BAAS, 37, 525
Zechmeister M., Kürster M., 2009, A&A, 496, 577
Zucker S., 2015, MNRAS, 449, 2723
Zucker S., 2016, MNRAS, 457, L118

APPENDIX A: MINIMUM GRID SPACING

Let us consider a signal made of p pure harmonics sampled at times $\boldsymbol{t} = (t_k)_{k=1,\dots,m}$, $\boldsymbol{y} = \sum_{j=1}^{p} c_j \mathrm{e}^{\mathrm{i}\omega_j t}$. We denote by ω'_j and $\Delta\omega$ two real numbers such that for each j

$$\Delta\omega < \frac{4}{T} \tag{A1}$$

$$|\omega_j - \omega'_j| < \Delta\omega\,, \tag{A2}$$

where $T = t_m - t_1$. For each t_k and each j,

$$|c_j||e^{\mathrm{i}\omega_j t_k} - e^{\mathrm{i}\omega'_j t_k}| = |c_j| \left| \mathrm{e}^{\mathrm{i}\frac{\omega_j+\omega'_j}{2}t_k} \left(\mathrm{e}^{\mathrm{i}\frac{\omega_j-\omega'_j}{2}t_k} - \mathrm{e}^{-\mathrm{i}\frac{\omega_j-\omega'_j}{2}t_k} \right) \right|$$
$$= 2|c_j| \left| \sin\left(\frac{\omega_j - \omega'_j}{2} t_k \right) \right|.$$

So denoting $\boldsymbol{y}' = \sum_{j=1}^{p} c_j \mathrm{e}^{\mathrm{i}\omega'_j t}$,

$$|y_k - y'_k| = \left| \sum_{j=1}^{p} c_j \left(\mathrm{e}^{\mathrm{i}\omega_j t_k} - \mathrm{e}^{\mathrm{i}\omega'_j t_k} \right) \right|$$
$$\leqslant 2 \sum_{j=1}^{p} \left| c_j \sin\left(\frac{\omega_j - \omega'_j}{2} t_k \right) \right|.$$

Without loss of generality, the origin of time is shifted to $-T/2$; therefore,

$$2 \sum_{j=1}^{p} \left| c_j \sin\left(\frac{\omega_j - \omega'_j}{2} t_k \right) \right| \leqslant \sin\frac{\Delta\omega T}{4} \sqrt{\sum_{j=1}^{p} |c_j|^2}. \tag{A3}$$

Finally, a condition for y' to be an acceptable solution is

$$\|\mathbf{W}(y - \boldsymbol{y}')\|_{\ell_2}^2 \leqslant \|\mathbf{W}\|^2 \|\boldsymbol{y} - \boldsymbol{y}'\|_{\ell_2}^2$$
$$\leqslant \|\mathbf{W}\|^2 \sum_{k=1}^{m} |y_k - y'_k|^2$$
$$\leqslant 4\|\mathbf{W}\|^2 \sum_{k=1}^{m} \left(\sum_{j=1}^{p} \left| c_j \sin\left(\frac{\omega_j - \omega'_j}{2} t_k \right) \right| \right)^2$$

given (equation A3),

$$\leqslant 4m\|\mathbf{W}\|^2 \sin^2\frac{\Delta\omega T}{4} \sum_{j=1}^{p} |c_j|^2$$

where $\|\mathbf{W}\| = \sup_{x\in\mathbb{C}^m} \frac{\|\mathbf{W}x\|_{\ell_2}}{\|x\|_{\ell_2}}$. When the matrix $\mathbf{W}$ is diagonal, the formula can be improved:

$$\|\mathbf{W}(\boldsymbol{y} - \boldsymbol{y}')\|_{\ell_2}^2 = \sum_{k=1}^{m} \frac{|y_k - y'_k|^2}{\sigma_k^2}$$
$$\leqslant 4 \sum_{k=1}^{m} \frac{1}{\sigma_k^2} \left(\sum_{j=1}^{p} \left| c_j \sin\left(\frac{\omega_j - \omega'_j}{2} t_k \right) \right| \right)^2$$

given (equation A3),

$$\leqslant 4 \sin^2\frac{\Delta\omega T}{4} \sum_{j=1}^{p} |c_j|^2 \sum_{k=1}^{m} \frac{1}{\sigma_k^2}.$$

So ϵ_{grid} can be chosen as

$$\epsilon_{\mathrm{grid}} = 2 \sqrt{\sum_{j=1}^{p} |c_j|^2} \sqrt{\sum_{k=1}^{m} \frac{1}{\sigma_k^2}} \sin\frac{\Delta\omega T}{4}. \tag{A4}$$

And conversely given an ϵ, the grid spacing that ensures that there exists a vector that has the correct ℓ_0 norm is

$$\Delta\omega = \frac{4}{T} \arcsin \frac{\epsilon}{2\sqrt{\sum_{j=1}^{p} |c_j|^2} \sqrt{\sum_{k=1}^{m} \frac{1}{\sigma_k^2}}}. \tag{A5}$$

APPENDIX B: DIGGING IN RED NOISE WITH NON-DIAGONAL *W*

B1 Short period buried in the noise

Our method uses the tools of compressed sensing, especially the algorithms to minimize ℓ_1 norms with the constraint that the reconstructed signal is not too far from the observations [see equation (5)]. To the best of our knowledge, the case where the noise is correlated has been considered only in Arildsen & Larsen (2014), and is not specialized for Gaussian processes. Here, we introduce a weight matrix and obtain problem (11,BP$_{\epsilon,W}$), reproduced here:

$$x^\star = \arg\min_{x\in\mathbb{C}^n} \quad \|\boldsymbol{x}\|_{\ell_1} \quad \text{s. t.} \quad \|\mathbf{W}(Ax - \boldsymbol{y})\|_{\ell_2} \leqslant \epsilon. \tag{BP$_{\epsilon,\mathbf{W}}$}$$

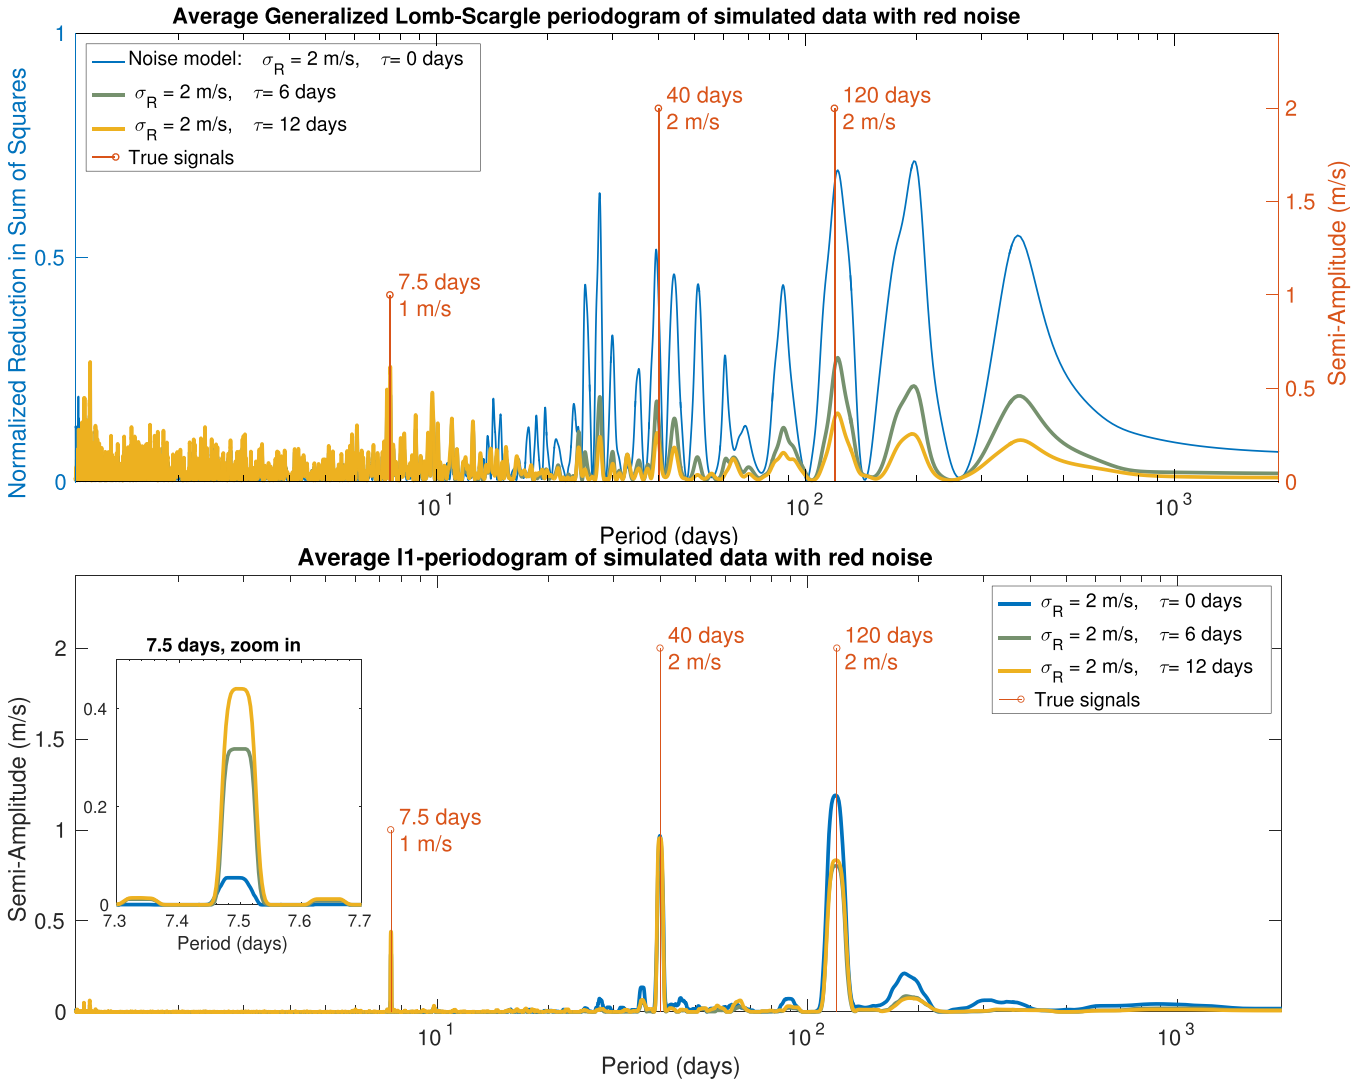

Figure B1. Average ℓ_1-periodogram for 50 data sets generated with red noise of characteristics $\sigma_W = 0$, $\sigma_R = 2$ m s^{-1} and $\tau = 12$ d according to model (14). The curves correspond to the solutions of equation (11,BP$_{\epsilon,W}$) with different weight matrices W whose parameters are $\sigma_W = 0$, $\sigma_R = 2$ m s^{-1} and $\tau = 0$, 6 or 12 d(respectively the blue, green and yellow curves).

To illustrate the interest of choosing an appropriate weight matrix, we will show an example where acknowledging the red noise makes a planet visible. Let us first consider a data set constructed as follows.

(i) The measurement times are those of HD 69830 (74 measurements).

(ii) The true signal is $y(t) = 1\cos(\frac{2\pi}{7.5}t) + 2\cos(\frac{2\pi}{40}t + 2) + 2\cos(\frac{2\pi}{120}t + 1)$ m s^{-1}.

(iii) The noise is red, with parameters $\sigma_W = 0$, $\sigma_R = 2$ m s^{-1} and $\tau = 12$ d, where $\sigma_{\mathbf{W}}$, σ_R and τ are the parameters of the autocorrelation function R defined in equation (14) reproduced here:

$$R(\Delta t) = \sigma_R^2 \mathrm{e}^{-\frac{|\Delta t|}{\tau}}, \quad \Delta t \neq 0$$

$$R(0) = \sigma_W^2 + \sigma_R^2.$$

The noise defined above is such that its correlation with low frequencies is higher than with high frequencies.

We test if changing the weight matrix could allow us to find signals that would not be seen otherwise. To do so, 50 noise time series $(n_k(\boldsymbol{t}))_{k=1,\ldots,50}$ are generated and the method is applied to each $y_k(\boldsymbol{t}) = \boldsymbol{y}(\boldsymbol{t}) + n_k(\boldsymbol{t})$ for three different weight matrices, all other parameters being fixed. In each case, they are defined according to model (14) with $\sigma_W = 0$, $\sigma_R = 2$ m s^{-1} and $\tau = 0$, 6 or 12 d. The grid goes between 0 and 0.95 cycles d^{-1} and ϵ verifies $F_{\chi^2_m}(\epsilon^2_{\text{noise}}) = 0.1$. The resulting ℓ_1-periodograms are averaged (see Fig. B1b).

To compare with a classical approach, we also compute classical periodograms for the same signals $y_k(\boldsymbol{t})$ and average them. For the comparison to be fair, we fit the model parameters A, B, C in $A\cos\omega t + B\sin\omega t + C$ to $\boldsymbol{y}(\boldsymbol{t})$ with the same weight matrices as the ones used above. This gives Fig. B1(a). If the weight matrix is left diagonal, then the low-frequency terms dominate. Using the appropriate noise model gradually reduces the spurious low frequencies.

We stress two features: as the noise model becomes accurate, the short period becomes apparent, which justifies the trial of different noise matrices on real radial velocity data sets to see if a peak appears. Secondly, when $\mathbf{W}$ is defined with an exponential autocorrelation function, the estimation of the peaks becomes biased: some frequencies will have a tendency to be interpreted by the algorithm as noise. The amplitude of the 120 d periodicity is then underestimated. This bias could prevent us from finding small amplitudes when using non-diagonal weight matrices. When the number of frequency in the signal increases, the bias becomes more

complicated. In order to mitigate this effect, we suggest to decrease the value of ϵ when testing different noise model. Thus, the model 'sticks' to the observations, and if a periodicity truly is in the data, the chance of it being too underestimated decreases. This is why, we took $\epsilon_{\rm noise}$ such that $F_{\chi^2_m}(\epsilon^2_{\rm noise}) = 0.1$ and not $F_{\chi^2_m}(\epsilon^2_{\rm noise}) = 0.999$, which would reject more signals in the residual.

B2 No automatic procedure so far

Here the improvement due to an appropriate handle of the noise is seen by eye. One could wonder if a simple criterion could allow one to chose an appropriate weight matrix automatically. In all cases when the algorithm has converged, we have $\|\mathbf{W}(\mathbf{A}\boldsymbol{x} - \boldsymbol{y})\|_{\ell_2} = \epsilon$ to a certain tolerance, or $x = 0$. Looking at the χ^2 of the residuals as usual is then not appropriate.

As in all cases the columns of matrix $\mathbf{W}A$ and the weighted observations Wy are normalized. Therefore, the problem always comes down to minimizing

$$x^\star = \arg\min_{\boldsymbol{x}\in\mathbb{R}^n} \|\boldsymbol{x}\|_{\ell_1} \text{ s. t. } \|\mathbf{A}'\boldsymbol{x} - \boldsymbol{y}'\|_{\ell_2} \leqslant \epsilon, \tag{B1}$$

where $\mathbf{A}'$ has normed columns and $\boldsymbol{y}'$ is a unitary vector. It is then tempting to see if there is a correlation between the ℓ_0 or ℓ_1 norm of $x^\star$ and the success of the method. Unfortunately, this is not the case. Whether there is an automatic way to select the appropriate weight matrix remains an open question.

APPENDIX C: SPURIOUS TALLEST PEAK OF THE GLS PERIODOGRAM

In this section, we show examples where the initial highest peak of the periodogram is spurious due to aliasing. We take the 74 measurement dates of HD 69830 and generate 500 systems with three circular orbits with the following properties.

(i) The amplitudes are those of the three Neptunes of HD 69830 (2.2, 2.66 and 3.51 m s^{-1}).
(ii) The periods P_1, P_2, P_3, are selected uniformly in $\log P$ in the range 1.2–2000 d.
(iii) The phases are uniformly distributed on $[0, 2\pi]$.
(iv) The noise standard deviation is 0.6 m s^{-1}.

We compute the number of times the maximum peaks of the GLS and ℓ_1-periodogram are spurious. The criterion we take for failure is when the frequency of the highest peak and any of the three true frequencies are greater than the inverse of the total observation time, that is $|1/P_{1,2,3} - 1/P_{\rm max}| > 1/T_{\rm obs}$.

Fig. C1 shows the GLS periodogram and ℓ_1-periodogram of representative cases where the highest peak of the GLS periodogram is spurious. In these conditions, when searching for periods in the 1.2–2000 d with the periodogram, we find that the strongest peak is spurious in 33 cases out of 500 simulations, while the tallest peak of the ℓ_1-periodogram only was incorrect in two cases. In those, the GLS periodogram was also failing.

An interesting feature of the cases where the ℓ_1-periodogram fails is that one can see that the solution is not sparse. This is a very useful property we observed empirically: we have not found any occurrence of ℓ_1-periodogram that looks clean, with well-separated clear peaks, where one of the peaks was completely spurious. We display one of the two failures of the ℓ_1-periodogram in Fig. C2. First of all neither the GLS nor the ℓ_1-periodogram leads the observer completely astray. Secondly, we see that as opposed to the ℓ_1-periodogram of the systems studied here, the figure is not clean, which should invite the analyst to a certain suspicion.

APPENDIX D: FITTING THE ANCILLARY MEASUREMENTS

In Section 4.6, we suggest to fit the activity indicators to the radial velocity time series. The present discussion wishes to give a justification to this approach. The idea is to exploit the possible correlations between radial velocity and ancillary measurements when the star is active. For instance, on the first system of the RV Fitting Challenge (Dumusque et al. 2016) where activity dominates the signal, the radial velocity, FWHM, bisector span and $\log R'_{HK}$ exhibit very similar features at low frequency (see Fig. D1).

Let us approximate the error made when fitting an ancillary indicator. We consider the radial velocity signal $\boldsymbol{y}(t) = \boldsymbol{P}(t) + \boldsymbol{a}(t) + \boldsymbol{\epsilon}(t)$, where $\boldsymbol{P}(t)$ is due to a planetary companion, $\boldsymbol{y}(\boldsymbol{a})$ is a deterministic signal due to activity and $\boldsymbol{\epsilon}$ is a Gaussian noise of covariance matrix $\mathbf{V}$. We also consider an ancillary measurement $\boldsymbol{z}(t) = \boldsymbol{a}(t) + \boldsymbol{\epsilon}'$, where $\boldsymbol{\epsilon}'(t)$ is another Gaussian noise of covariance matrix $\mathbf{V}$. If we fit $\boldsymbol{z}(t)$ to $\boldsymbol{y}(t)$, we obtain (dropping the t notation)

$$\boldsymbol{y}_{\rm detrend} = \boldsymbol{y} - \boldsymbol{y}_{\rm fit} = \boldsymbol{y} - \frac{\boldsymbol{z}^T\mathbf{V}^{-1}y}{\boldsymbol{z}^T\mathbf{V}^{-1}\boldsymbol{z}}\boldsymbol{z} \tag{D1}$$

$$\boldsymbol{y}_{\rm detrend} = \boldsymbol{y} - \frac{(\boldsymbol{a}+\boldsymbol{\epsilon}')^T\mathbf{V}^{-1}(\boldsymbol{P}+\boldsymbol{a}+\boldsymbol{\epsilon})^T}{(\boldsymbol{a}+\boldsymbol{\epsilon}')^T\mathbf{V}^{-1}(\boldsymbol{a}+\boldsymbol{\epsilon})}(\boldsymbol{a}+\boldsymbol{\epsilon}'). \tag{D2}$$

We assume that the noise is small compared to $\boldsymbol{a}$, which allows us to develop the denominator at first order in $\boldsymbol{\epsilon}$ and $\boldsymbol{\epsilon}'$

$$y_{\rm fit} \approx \frac{(\boldsymbol{a}+\boldsymbol{\epsilon}')^T\mathbf{V}^{-1}(\boldsymbol{P}+\boldsymbol{a}+\boldsymbol{\epsilon})}{\boldsymbol{a}^T\mathbf{V}^{-1}\boldsymbol{a}} \times \left(1 - \frac{\boldsymbol{\epsilon}'^T\mathbf{V}^{-1}\boldsymbol{a}}{\boldsymbol{a}^T\mathbf{V}^{-1}\boldsymbol{a}} - \frac{\boldsymbol{\epsilon}^T\mathbf{V}^{-1}\boldsymbol{a}}{\boldsymbol{a}^T\mathbf{V}^{-1}\boldsymbol{a}}\right)(\boldsymbol{a}+\boldsymbol{\epsilon}').$$

After developing that expression at first order in $\boldsymbol{\epsilon}$ and ϵ', we compute its mathematical expectancy taking into account only the zero-order, ϵ^2 and ϵ'^2 coefficients. In the simple case where the noise is i.i.d. of variance σ^2, we obtain

$$\mathbb{E}\{\boldsymbol{y}_{\rm fit}\} \approx \frac{\sigma^2}{\|\boldsymbol{a}\|^2}\boldsymbol{P} \tag{D3}$$

$$+ \left(1 + \frac{\boldsymbol{a}^T\boldsymbol{P}}{\|\boldsymbol{a}\|^2_{\ell_2}} - \frac{2\sigma^2}{\|\boldsymbol{a}\|^2_{\ell_2}} - \frac{\|\boldsymbol{P}\|_{\ell_2}\sigma^2}{\|\boldsymbol{a}\|^3_{\ell_2}} - \frac{\boldsymbol{a}^T\boldsymbol{P}\sigma^2}{\|\boldsymbol{a}\|^4_{\ell_2}}\right)\boldsymbol{a}. \tag{D4}$$

We would like $\boldsymbol{y}_{\rm fit}$ to be as close to $\boldsymbol{a}$ as possible. This will be better satisfied as the correlation $\boldsymbol{a}^T\boldsymbol{P}$ and the signal-to-noise $\sigma^2/\|\mathbf{a}\|_{\ell_2}$ decrease. The fact that a term $\boldsymbol{a}^T\boldsymbol{P}$ appears in the equation above should not be surprising. The mutual coherence defined in Section 5.2 grasps that the correlation between the parts of the model is an obstacle to the recovery of the true signals.

For the RV Fitting Challenge, not only have we fitted one activity indicator but several. We point out that this approach is consistent with Rajpaul et al. (2015). Indeed, they consider that the activity-induced variations of the measurements depend linearly on an underlying zero-mean Gaussian process $G(t) = F^2(t)$ and its derivative $\dot{G}(t)$, where $F(t)$ is the fraction of the sphere covered with spots. The evolution of the indicators is modelled by formulae (14)–(16), reproduced below,

$$\Delta{\rm RV} = V_c G(t) + V_r \dot{G}(t), \tag{D5}$$

$$\log R'_{HK} = L_c G(t), \tag{D6}$$

$$\text{BIS} = B_c G(t) + B_r \dot{G}(t), \tag{D7}$$

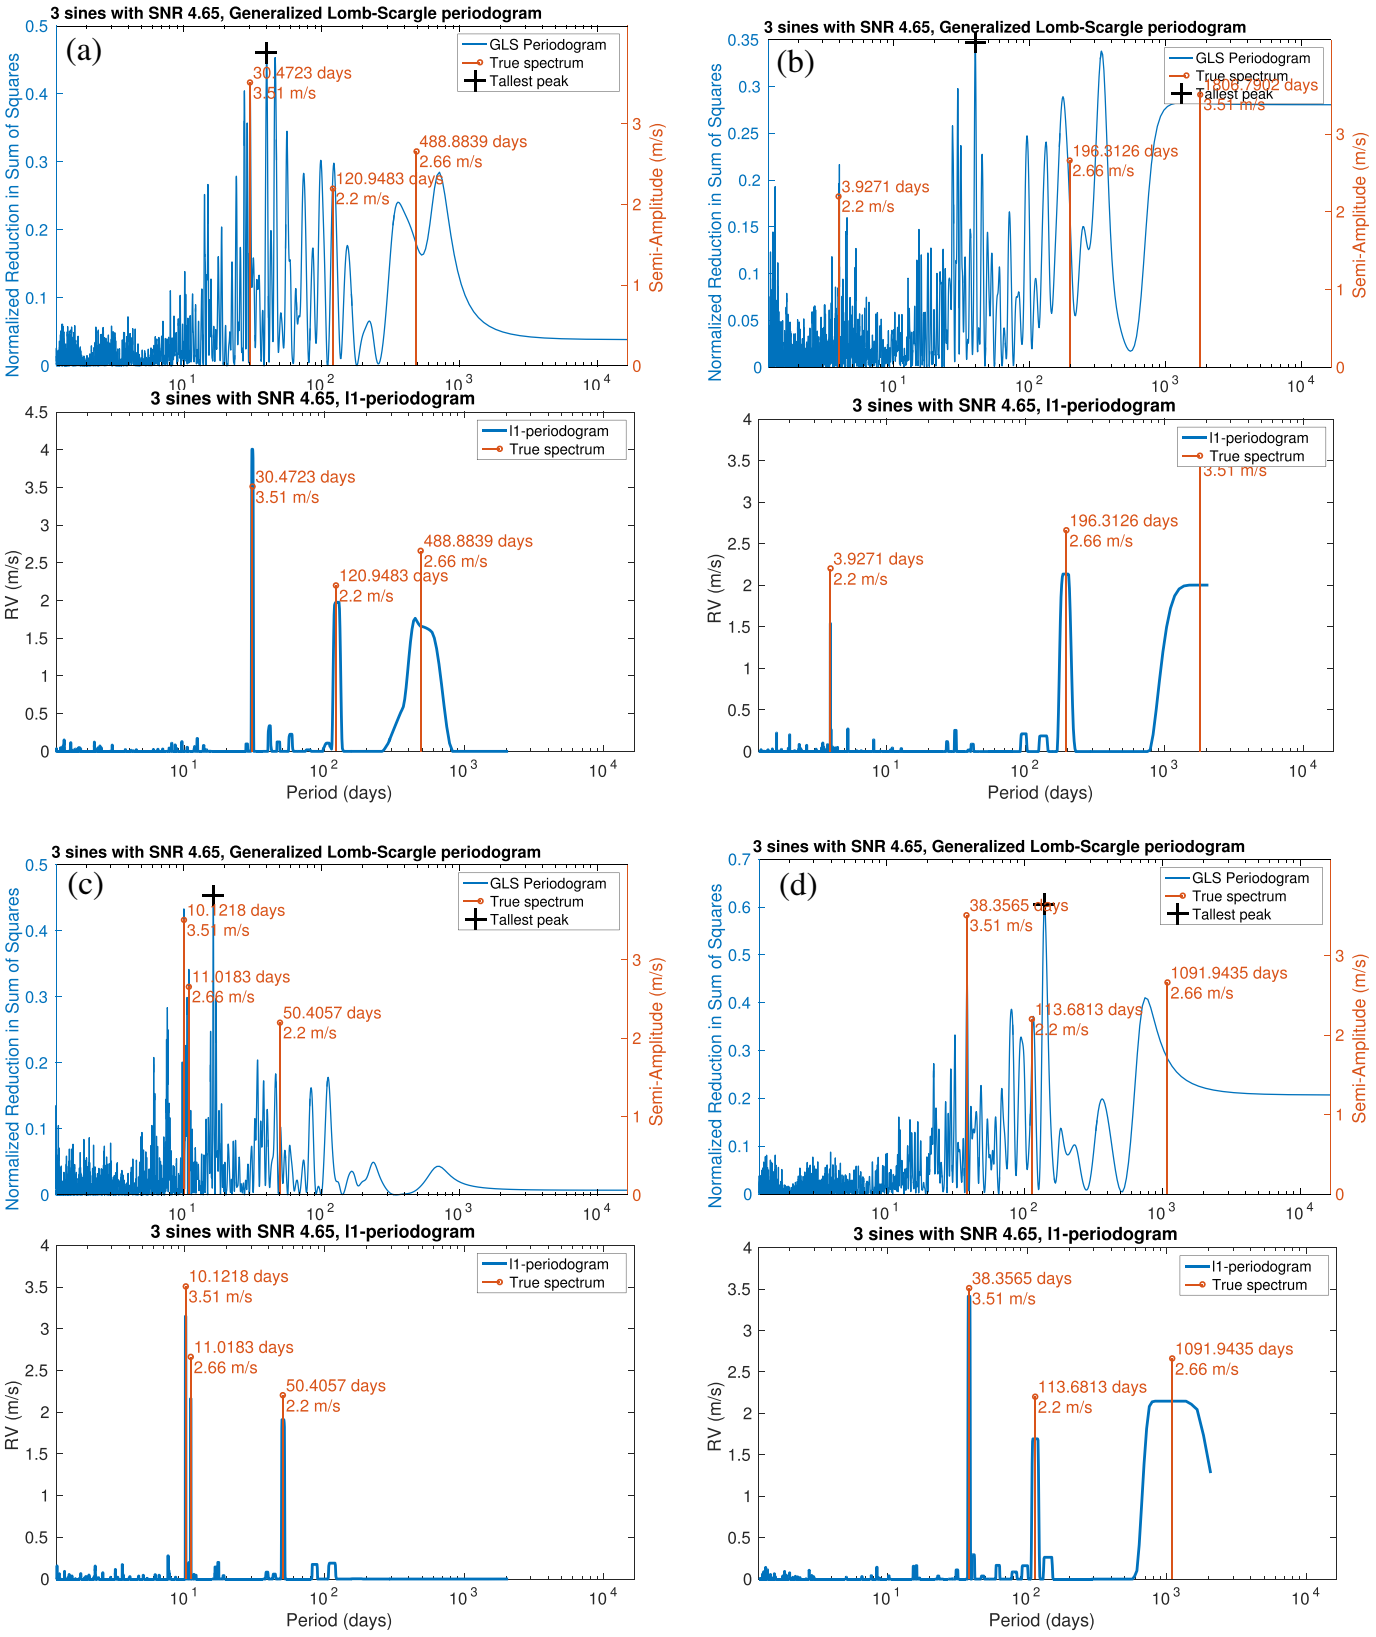

Figure C1. Peak amplitudes and associated FAPs for the four systems analysed.

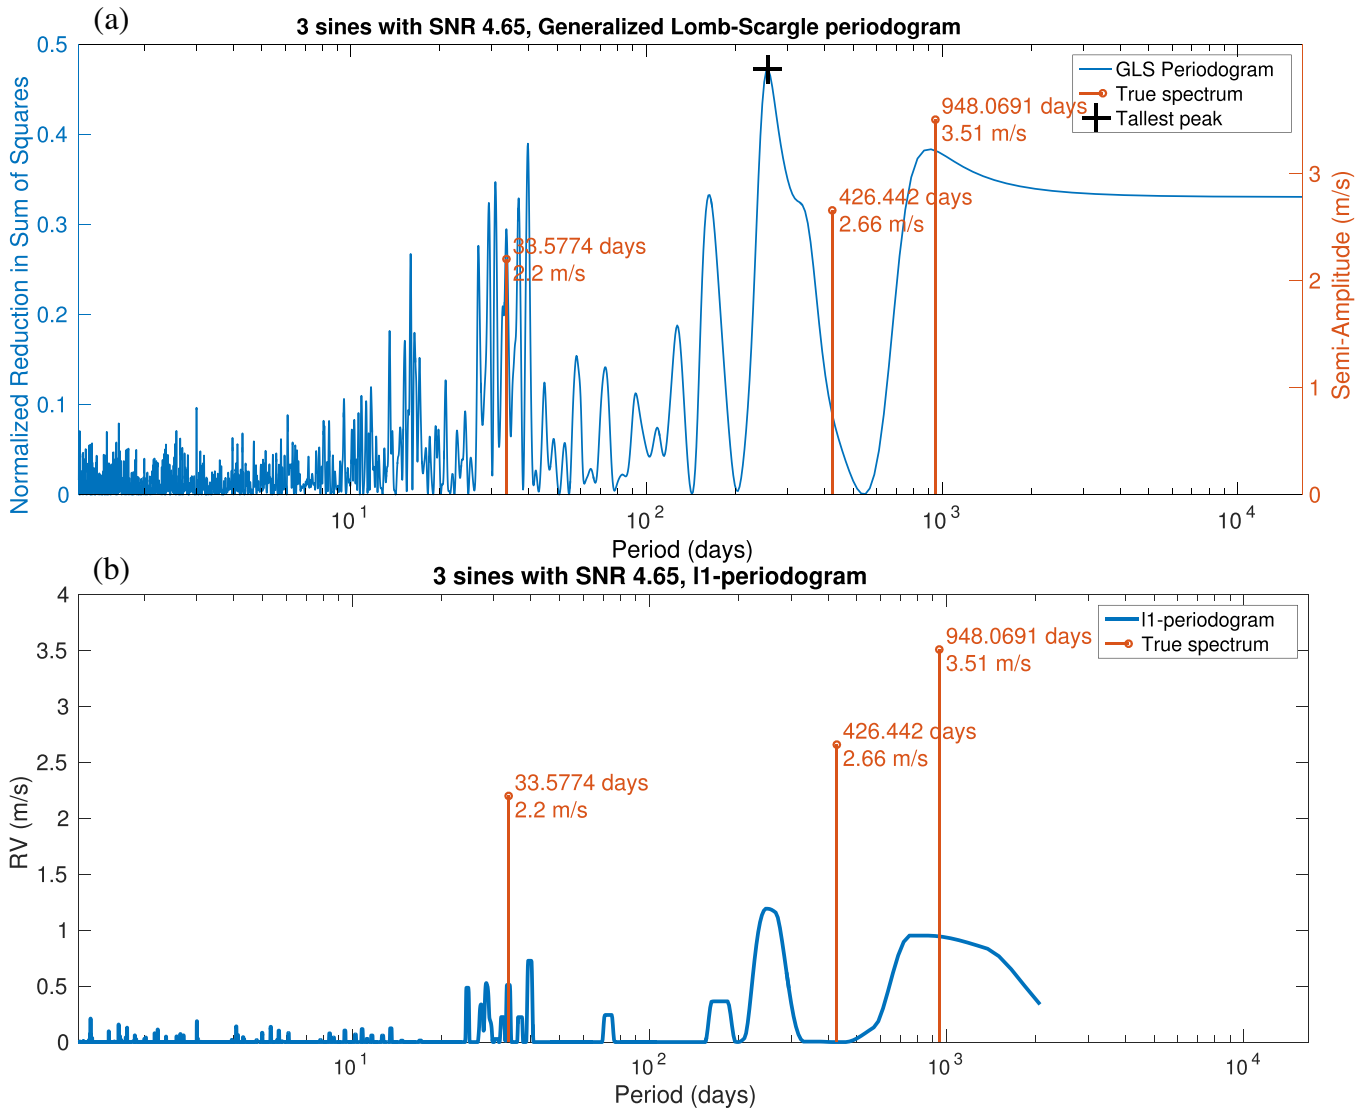

Figure C2. Failure of the GLS (a) and ℓ_1 (b) periodograms.

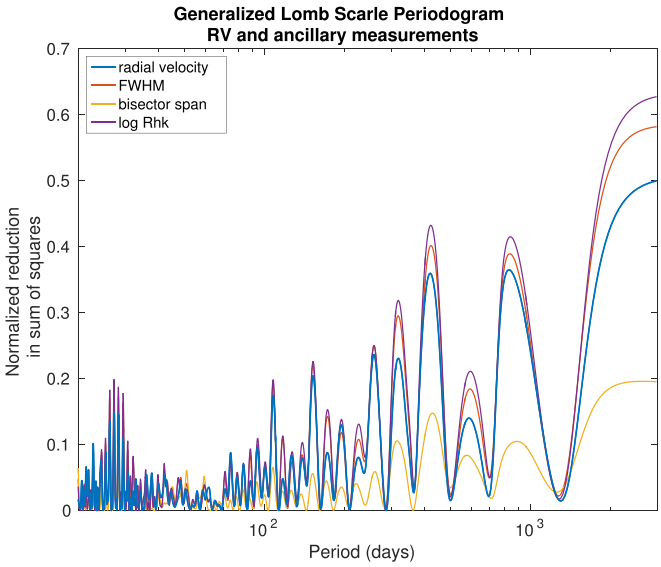

Figure D1. GLS periodogram of radial velocity and ancillary measurements at low frequencies.

for some constants V_c, V_r, L_c, B_c, B_r. This means that for a given realization $(\boldsymbol{g}, \boldsymbol{g}')$ of $(G(t), \dot{G}(t))$, the subspace generated by the log R'_{HK} and the bisector span BIS is the same as the space generated by $\boldsymbol{g}, \boldsymbol{g}'$. So according to that model, projecting the radial velocity on to (log R'_{HK},BIS) is equivalent to projecting on to $(\boldsymbol{g}, \boldsymbol{g}')$.

However, there is an uncertainty on the behaviour of the ancillary measurements and additional noise. We have to decide if fitting an uncertain model is better than working with the raw data. One thing that could happen is that fitting the combination of the three ancillary measurements would greatly change the spectral content of the radial velocity time series by absorbing some frequencies, potentially due to planets. To estimate this risk, we first compute the term $\boldsymbol{a}^T \boldsymbol{P}/\|\boldsymbol{a}\|_{\ell_2}^2$ in equation (D4), assuming that the signal $\boldsymbol{y} = \boldsymbol{P} = \mathrm{e}^{\mathrm{i}\omega t}$ is a pure harmonic of amplitude 1 m s^{-1}. Here $\boldsymbol{a}$ designates the FWHM, bisector span or log R'_{HK}, respectively, shown by the red, yellow and purple curves in Fig. D2. We also compute the fraction of the energy of the signal before and after the fit of the three ancillary measurements simultaneously, that is

$$\text{Fraction}(\omega) = \frac{(\boldsymbol{y}_\omega - \boldsymbol{y}_{\text{fit}})^T \mathbf{V}^{-1} (\boldsymbol{y}_\omega - \boldsymbol{y}_{\text{fit}})}{\boldsymbol{y}_\omega^T \mathbf{V}^{-1} \boldsymbol{y}_\omega}; \tag{D8}$$

this one is represented by the blue curve in Fig. D2.

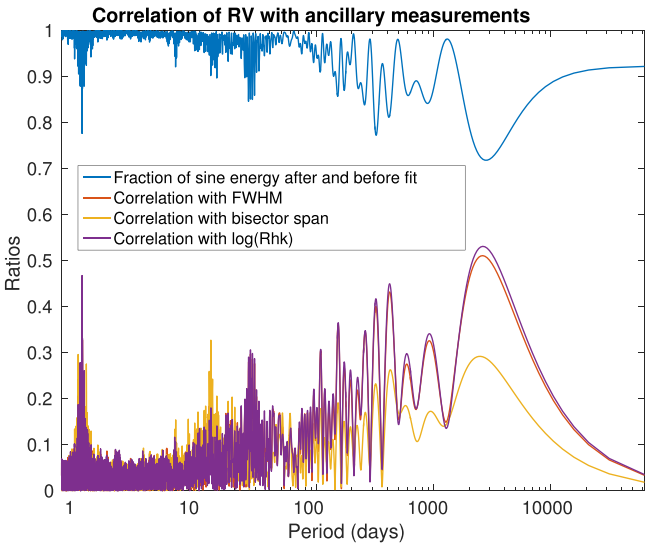

Figure D2. Energy of a cosine function after the fit of the FWHM, bisector span, log R'_{HK} and a constant.

For the system analysed in Section 4.6, only 15 per cent of the energy is absorbed in general, with a maximum of 27 per cent at a period of 2000 d. The peaks at 25 and 12.5 d correspond to the rotation period of the star and its first harmonic, which are expected to be correlated with the radial velocity and ancillary measurements.

This discussion does not intend to provide strong statistical arguments, but rather to show that the spectral content should not be too affected by fitting the FWHM, bisector span and log R'_{HK}.

This paper has been typeset from a TEX/LATEX file prepared by the author.

Résumé

Lorqu'une étoile a des compagnons planétaires, elle décrit un mouvement quasi épicycloïdal autour du centre de masse du système. Si l'orientation du plan de l'orbite le permet, un observateur situé sur la Terre peut détecter la composante de ce mouvement sur la ligne de visée grâce à l'effet Doppler. Il mesure ainsi la "vitesse radiale de l'étoile". Si cette vitesse présente des variations périodiques suffisamment claires, la présence de planètes peut être inférée et leurs orbites contraintes.

Une des difficultés de l'analyse de telles mesures est qu'une combinaison de signaux de plusieurs planètes et de divers bruits peut être confondue avec l'effet d'une planète en réalité inexistante. Après avoir présenté les effets à prendre en compte pour analyser des données de vitesses radiales, nous abordons ce problème. Pour limiter son occurrence, nous utilisons un algorithme de poursuite de base modifié, dont on démontre l'efficacité sur des signaux réels et simulés.

Nous abordons ensuite le problème de l'estimation des paramètres orbitaux pour un système donné ainsi que leur distribution pour une population de planètes. On s'intéresse en particulier à l'excentricité, dont on montre qu'elle est d'autant plus surestimée que le modèle du signal est mauvais. Nous proposons des solutions pour une estimation robuste des paramètres orbitaux.

Mots Clés

Exoplanètes, vitesses radiales, Inférence, Analyse de données, Acquisition Comprimée, Excentricité, Biais, Robustesse

Abstract

When a star is orbited by planetary companions, it describes a nearly epicyclic motion around the center of mass of the system. When the orientation of the orbital plane is appropriate, an observer on Earth can measure the velocity of the star along the line of sight by Doppler effect. If this "radial velocity" presents clear enough periodic variations, the presence of planets can be inferred and their orbit can be constrained.

Detection and estimation of orbits is made difficult by the photon noise, the unpredictable variations of luminosity of the star as well as instrumental faults. In particular, signals from several planets can add coherently with the noises and mimic the effect of a planet absent from the system. After listing the relevant effects to make inference on exoplanets from radial velocity data, we tackle this problem. To limit its rate of occurrence, we use a modified basis pursuit algorithm, allowing to search for several signals simultaneously. The efficiency of the method is demonstrated on real and simulated signals.

We then address the problem of orbital parameters estimation for a given system, as well as the estimation of their distribution on a planet population. We look in detail at the eccentricity, and show that its overestimation increases as the model moves away from the correct one. We suggest methods for robust inference of orbital parameters.

Keywords

Exoplanets, Radial Velocity, Inference, Data analysis, Sparsity, Compressed Sensing, Eccentricity, Bias, Robustness