Validation of solar-cycle changes in low-degree helioseismic parameters from the Birmingham Solar-Oscillations Network

Howe, R.; Davies, G. R.; Chaplin, W. J.; Elsworth, Y. P.; Hale, S. J.

DOI:
[10.1093/mnras/stv2210](https://doi.org/10.1093/mnras/stv2210)

License:
None: All rights reserved

Document Version
Publisher's PDF, also known as Version of record

Citation for published version (Harvard):
Howe, R, Davies, GR, Chaplin, WJ, Elsworth, YP & Hale, SJ 2015, 'Validation of solar-cycle changes in low-degree helioseismic parameters from the Birmingham Solar-Oscillations Network', *Royal Astronomical Society. Monthly Notices*, vol. 454, no. 4, pp. 4120-4141. <https://doi.org/10.1093/mnras/stv2210>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

Checked for eligibility: 22/01/2016. This article has been accepted for publication in *Monthly Notices of the Royal Astronomical Society* © 2015 The Authors. Published by Oxford University Press on behalf of the Royal Astronomical Society. All rights reserved.

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Validation of solar-cycle changes in low-degree helioseismic parameters from the Birmingham Solar-Oscillations Network

R. Howe,^{*} G. R. Davies, W. J. Chaplin, Y. P. Elsworth and S. J. Hale

School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK

Accepted 2015 September 22. Received 2015 August 26; in original form 2015 May 22

ABSTRACT

We present a new and up-to-date analysis of the solar low-degree p-mode parameter shifts from the Birmingham Solar-Oscillations Network over the past 22 years, up to the end of 2014. We aim to demonstrate that they are not dominated by changes in the asymmetry of the resonant peak profiles of the modes and that the previously published results on the solar-cycle variations of mode parameters are reliable. We compare the results obtained using a conventional maximum-likelihood estimation algorithm and a new one based on the Markov Chain Monte Carlo (MCMC) technique, both taking into account mode asymmetry. We assess the reliability of the solar-cycle trends seen in the data by applying the same analysis to artificially generated spectra. We find that the two methods are in good agreement. Both methods accurately reproduce the input frequency shifts in the artificial data and underestimate the amplitude and width changes by a small amount, around 10 per cent. We confirm earlier findings that the frequency and line width are positively correlated, and the mode amplitude anticorrelated, with the level of solar activity, with the energy supplied to the modes remaining essentially unchanged. For the mode asymmetry the correlation with activity is marginal, but the MCMC algorithm gives more robust results than the MLE (Maximum-Likelihood Estimate). The magnitude of the parameter shifts is consistent with earlier work. There is no evidence that the frequency changes we see arise from changes in the asymmetry, which would need to be much larger than those observed in order to give the observed frequency shift.

Key words: Sun: helioseismology – Sun: oscillations.

1 INTRODUCTION

The variation of solar acoustic mode (p-mode) parameters with solar activity is one of the best-known results in helioseismology, and has been studied at a wide range of spatial and temporal scales over the last few decades. In the light of recent developments in mode parameter estimation techniques it seems timely to revisit these findings. In this work we re-analyse the Sun-as-a-star data from the Birmingham Solar-Oscillations Network (BiSON) over the last two solar cycles and verify our ability to detect subtle variations using two different algorithms, one that uses conventional maximum-likelihood estimation and a new one based on Bayesian principles

1.1 Historical background

1.1.1 Frequency variation

Small changes in the frequencies of the low-degree modes, positively correlated with proxies for the solar activity such as sunspot

number and 10.7-cm radio flux (RF), were reported for example by Woodard & Noyes (1985) using data from the *Active Cavity Radiation Monitor* (ACRIM), and later confirmed in ground-based observations by Pallé, Regulo & Roca Cortes (1989) and Elsworth et al. (1990), while Libbrecht & Woodard (1990) first reported solar-cycle changes in the frequencies of medium-degree modes from resolved-Sun observations at the Big Bear Solar Observatory. Libbrecht & Woodard (1990) also pointed out that the variation increased with frequency, following the inverse of the so-called ‘mode inertia.’ This indicates that the mechanism responsible for the shifts is located close to the solar surface. The frequency dependence of the variation is harder to detect in low-degree data due to the smaller number of modes available, but within a few years it was reported by Elsworth et al. (1994).

In resolved-Sun data the frequency shifts follow the activity variation in location as well as in time (e.g. Howe, Komm & Hill 2002, and references therein); modes of the same radial order n and degree l but different azimuthal order m have different distributions of power with latitude and so are differently affected as the magnetic activity belts move in latitude over the solar cycle. Even for low-degree modes, Chaplin et al. (2004), Jiménez-Reyes et al. (2004), Chaplin et al. (2007), and Salabert, García & Turck-Chieze (2015)

^{*}E-mail: rhowe@nso.edu

have reported marginally significant differences in the response of the frequency to activity for modes of different degree, consistent with the different latitudinal distribution of power for the different spherical harmonics.

1.1.2 Power and width variations

Magnetic activity influences not only the frequency of the acoustic modes but also their power and lifetime. Elsworth et al. (1993) reported an anticorrelation between solar activity and the amplitude of low-degree modes, with about a 30 per cent change between solar minimum and solar maximum. The positive correlation between mode width and activity took longer to establish, due to the difficulties introduced by systematic effects such as the broadening of the modes when the duty cycle is low. However, Chaplin et al. (2000), after a careful analysis of BiSON data, reported a change of around 24 per cent in the line width and a similar decrease in the mode power (amplitude squared) between solar minimum and maximum. The combination of these two results suggested that the solar-cycle change affects the mode damping rather than the excitation, with the rate of energy supplied to the modes remaining constant. Komm, Howe & Hill (2000a,b) found activity-correlated changes in the mode width for medium-degree Global Oscillation Network Group (GONG) data, which could also be localized to the latitudes where magnetic activity is present. (Komm, Howe & Hill 2002). Again, the energy supply rate appears to remain constant.

Salabert et al. (2007) studied the line-width variation in 9.5 years of BiSON data and found marginal evidence for a degree dependence of the sensitivity, with the $l = 0$ shifts about half the size of those for $l \geq 1$, consistent with what would be expected from the latitudinal distribution of the modes in relation to the latitudinal migration of the activity belts during the solar cycle. They concluded that previous estimates of Sun-as-a-star line-width changes may have been overestimated by about 50 per cent because the shifts were averaged over all the modes.

1.1.3 Asymmetry

As discussed below in Section 2, the mode peaks show a small asymmetry due to the correlation of mode excitation with the noise that drives it (Duvall et al. 1993). This term is negative for velocity observations and generally positive for measurements made in intensity; it is smallest where the modes are strongest and for low-degree observations can reach a value of a few per cent at the extremes of the p-mode band and a fraction of 1 per cent at 3 mHz. Jiménez-Reyes et al. (2007) reported a fractional change of about 15 per cent between solar maximum and solar minimum in the asymmetry of low-degree modes in data from BiSON and the Global Oscillations at Low Frequencies (GOLF) instrument on the *Solar and Heliospheric Observatory* (SOHO) spacecraft, with the strongest (most negative) asymmetry at solar maximum. As the uncertainty in the asymmetry depends strongly on the signal-to-noise (S/N) ratio of the modes and also on the duty cycle, the result was most clearly seen in the GOLF data and was described as ‘marginally significant’ for BiSON.

1.2 Mechanisms

There is still some uncertainty about the precise mechanisms responsible for the parameter shifts. Attempts have been made to model the frequency shifts by invoking the direct effects of magnetic

field at the tachocline (Roberts & Campbell 1986), the sunspot anchoring zone around 50 Mm below the surface (Foullon & Roberts 2005), the photosphere (Bogdan & Zweibel 1985), or the chromosphere (Campbell & Roberts 1989; Jain & Roberts 1994), but none of these have predicted shifts of the observed magnitude. The shifts have also been considered as an indirect effect of temperature changes associated with the activity belts (Kuhn 1988) and as an effect of a change in acoustic cavity size (Dziembowski & Goode 2005). None of these models is universally accepted. In any case, it is possible to consider the frequency shifts as a solar-cycle proxy in themselves, and one of the few that has some sensitivity to layers below the photosphere.

1.3 Motivation

Two relatively recent developments have prompted us to revisit the mode parameter variations over the last two decades of BiSON observations.

First, although there is broad consensus on the size and sign of the different solar-cycle effects on mode parameters, Korzennik (2013) has reported results from a re-analysis of GONG, Michelson Doppler Imager, and Helioseismic and Magnetic Imager medium-degree data showing frequency changes smaller by about a factor of 2 than have been seen in the standard analysis of these data sets. The author attributes the discrepancy to the use of asymmetric peak profiles in the fitting, and also reports variations in the peak width and asymmetry, with the asymmetry changes possibly accounting for the ‘missing’ frequency shifts. These results make it timely to reinvestigate the frequency variations over the past two solar cycles of BiSON observations and assess the validity of our mode-parameter variations using fits that take into account the asymmetry of the modes.

Secondly, there is growing evidence of, and interest in, temporal variations in the mode frequency that cannot be described simply as a linear function of the global activity level. These may take the form of fluctuations superimposed on the linear trend with activity proxy or of changes in the strength of the trend over time; the two are not necessarily clearly distinguished, as the presence of an additional short-term variation may reduce the strength of the correlation with activity. Following the long, anomalous solar minimum after Solar Cycle 23 and the relatively weak Cycle 24, there has been particular interest in any changes in frequency shifts that might indicate changes in the structure of the outer solar layers. Basu et al. (2012) found that the shifts in Cycle 23 BiSON data did not follow the trend with activity extrapolated from the Cycle 22 data, with the mismatch showing up first in the low-frequency shifts and only later at higher frequencies. On the other hand, Salabert et al. (2015), using GOLF data, reported a frequency increase in Cycle 24, particularly for $l = 1$ modes, that was larger than that expected from the increase in activity level when compared to the results from Cycle 23. It seems that the interpretation of such results is quite sensitive to which modes are averaged and what period is taken as the reference.

There have also been reports of variations in the mode frequencies on time-scales that do not correspond to the 11-yr solar cycle. Howe et al. (2006) reported aperiodic fluctuations in individual low-degree mode frequencies after the subtraction of the activity effect, which were correlated between different instruments but not with any activity index and not correlated between different modes. These variations were attributed to stochastic fluctuations in the solar spectrum. Fletcher et al. (2010), Broomhall et al. (2011), and Simoniello et al. (2012, 2013) later reported a quasi-biennial

variation in the mode frequencies, superimposed on the 11-yr variation, which they interpreted as evidence of a ‘second solar dynamo’ located in the near-surface layers. This variation appears to be correlated with other activity proxies in high-activity periods but persists during solar minima when the proxies are flat. Along similar lines, Jain, Tripathy & Hill (2011) found that the GONG medium-degree frequencies showed hints of a ‘double minimum’ between Solar Cycles 23 and 24, with the frequencies first reaching their lowest value ahead of the activity minimum; this may reflect a similar phenomenon to that seen in the low-degree data.

All of these results make it timely to re-examine the parameter variations and assess their reliability.

1.4 Arrangement

The arrangement of the paper is as follows: in Section 2 we discuss the issues related to correlated and uncorrelated noise and peak asymmetry and go on to describe the specifics of the data simulation. In Section 3 we describe the algorithms used for extracting the mode parameters. In Section 4 we describe the tests carried out on the artificial data and discuss the results and their implications for the analysis of BiSON data. In Section 5 we present the results from fitting the BiSON data, and in Section 6 we discuss our conclusions.

2 GENERATION OF ARTIFICIAL DATA

In order to test for possible biases introduced into the mode parameters and their variations we have carried out a series of tests on artificial data. The generation of the artificial data was based on the solarFLAG simulator, which was used to make data for the hare-and-hounds study on rotational frequency splittings discussed by Chaplin et al. (2006) and also to test peak-bagging of Sun-as-a-star helioseismic data in Jiménez-Reyes et al. (2007). The original solarFLAG simulator did not include effects of correlated excitation, or of correlations of the excitation with background noise. An important consequence was that the artificial mode peaks in the frequency power spectrum showed no asymmetry, unlike their real solar counterparts. Since any analysis which seeks to extract accurate estimates of the frequencies must cope with this asymmetry it was felt we needed a simulator that could provide such a test. The departure of the mode shape from the pure Lorentzian predicted by a simple harmonic oscillator has three main causes. They are the localization of the noise source, the impact of correlated noise on mode excitation, and finally the more subtle effect of the presence of the wings of nearby correlated modes. We have used a simple

but very powerful method to introduce in the time domain the effects of asymmetry, which is based on the framework proposed by Toutain, Elsworth & Chaplin (2006). The influence of source localization (Chaplin & Appourchaux 1999) does not need to be explicitly considered as it produces a line shape that is the same as from the correlated background. We include it implicitly by our choice of correlation coefficient. Thus, there are just two factors in our method that contribute to the asymmetry of the artificial mode peaks. First, background noise is correlated with the excitation of the modes, and second, overtones of the same angular degree and azimuthal order have excitation functions that are correlated in time. This correlated mode excitation is based on the description given in Chaplin, Elsworth & Toutain (2008) and the method is described in detail in the following sections. Further details, and analytical descriptions of the examples in Section 2.1 below, may be found in Chaplin et al. (2008).

2.1 Correlated noise and mode excitation

In Toutain et al. (2006) it was hypothesized that the excitation function of a mode of angular degree l , azimuthal degree m , and frequency ν , is the same as that component of the solar background (granulation) noise that has the same spherical harmonic projection, Y_{lm} , in the corresponding range in frequency in the Fourier domain. Let us start by considering the impact on the line shape of this correlated background noise. We do this in the context of how the solarFLAG simulator generates artificial data on single p modes.

The basis of the solarFLAG simulator is the method discussed by Chaplin et al. (1997) for generating time series of individual p modes. The method uses the Laplace transform solution of the equation of a forced, damped harmonic oscillator to make the output velocity of each artificial mode. Oscillators are re-excited at each time sample – the chosen cadence is typically 40 or 60 s for Sun-as-a-star data – with small ‘kicks’ from a time series of random noise. This procedure mimics the stochastic excitation of the solar p modes. For the moment we shall assume the random noise data are drawn from a normal (i.e. white) distribution having zero mean. We shall see later that in the final version of the simulator we used ‘granulation-like’ noise – to mimic the spectrum of solar granulation – which has a spectrum that looks white only locally in the vicinity of the oscillator resonance.

The left-hand panel in Fig. 1 shows, on a logarithmic scale, the limit power spectral density of two scenarios, which illustrate the impact of correlated noise. In both scenarios the power spectral density has contributions from a mode – of frequency $\nu = 2990 \mu\text{Hz}$,

Figure 1. Limit frequency power spectra for scenarios with one or two modes. Left-hand panel: a single mode, with background noise. The dashed red line shows the spectrum for no correlation; the solid black line for when the mode is correlated with the background noise. Centre panel: two modes. The dashed red line shows the limit spectrum when the excitation of the modes is uncorrelated; the solid black line when the excitation is correlated. Right-hand panel: two modes, with background noise. The dashed red line shows the limit spectrum when there is no correlation of the excitation, or with the background noise; the solid black line when the excitation is correlated, and there is correlation with the background noise.

and line width $\Gamma = 1 \mu\text{Hz}$ – and from white background noise. The thin dashed red line shows the result for a scenario where the excitation of the mode is *uncorrelated* with the background noise. The solarFLAG simulator would generate the data for this scenario by using one time series of random noise to excite the mode; to the velocity output of the mode would then be added another, completely uncorrelated time series of appropriately scaled random noise, to represent the background (here the height-to-background ratio is 100). Since there is no correlation, the limit frequency power spectrum is given simply by the incoherent addition of the limit spectrum of the mode (a Lorentzian) and the limit spectrum of the noise (here, a flat offset for white noise).

Now, what if we use the excitation time series as the background noise? The excitation and the background will now be 100 per cent correlated, and the power spectral density will have a peak that is asymmetric. This situation, for a large (about 10 per cent) value of the asymmetry, is shown with the thick black line in the left-hand panel of Fig. 1. If only the effect of the background correlation had been considered the peak asymmetry would have been positive with an excess of power on the high-frequency side of the peak. However, the asymmetry seen in low- l solar p modes observed in Doppler velocity is negative and hence we show the effect of negative correlation with excess power on the low-frequency side of the resonance. The observed asymmetry is much lower than illustrated here.

Next, we consider the impact of correlated mode excitation.

An important implication of the framework proposed by Toutain et al. (2006) is that overtones with the same (l, m) should have excitation functions that are correlated in time. (Note that the Y_{lm} for (l, m) and $(l, -m)$ are orthogonal, and are therefore assumed to have independent, i.e. uncorrelated, excitation.) To illustrate the impact of this correlation, we look at the simplest possible scenarios, which have just two modes in the frequency power spectrum. Consecutive overtones of the low- l solar p modes are separated in frequency by $\sim 135 \mu\text{Hz}$. Here, we consider two modes separated in frequency by $20 \mu\text{Hz}$. This smaller frequency spacing exaggerates the impact of the correlated excitation on the observed power spectral density, and therefore allows us to show more clearly the effect of the correlation. We again assume each mode has a line width of $1 \mu\text{Hz}$. The impact of the actual $\sim 135 \mu\text{Hz}$ spacing is considered in Section 2.2 below.

The thin dashed red line in the middle panel of Fig. 1 shows the limit spectrum for two modes whose excitation is uncorrelated in time. There is no background noise. The solarFLAG simulator would generate the data for this scenario by using independent time series of random noise to excite each oscillator. Now, what happens if the simulator excites both oscillators with the same time series of random noise? The excitation is now 100 per cent correlated in time, and the power spectral density shows clearly that the peaks are asymmetric (thick dark line). This asymmetry comes from the interaction of the tails of the mode peaks. We draw an important conclusion from this example: correlated excitation of modes will give a contribution to the observed asymmetry that is dependent on how the tails of the individual modes overlap. It turns out that this effect is important even for quite well-separated modes.

In our final example, we consider both sources of asymmetry together and we add background noise to each two-mode scenario above. The thin dashed red line in the right-hand panel of Fig. 1 is for a scenario where the background noise and the mode excitation are all uncorrelated. The thick dark line instead shows what happens if we add correlated background noise to the two correlated modes. This means the excitation and background noise are all 100 per cent correlated, and we see that addition of the back-

ground further modifies the shape of the power spectral density, relative to the correlated-mode example with no background noise. As such, there are now two factors which contribute to the peak asymmetry: there is a contribution from the correlated excitation of individual modes (and also potentially from source localization), and a contribution from the correlated background.

One other important point to note is that correlation of the excitation in time does *not* imply correlation of the mode amplitudes in time. This can be understood by considering the analogy of damped, stochastically forced oscillators. Modes of different frequencies will be ‘kicked’ by the common excitation at different phases in their oscillation cycles, and provided the frequencies differ by more than a few line widths (see Chaplin et al. 2008) – a condition easily met by consecutive overtones of the low- l modes, which are separated by $\sim 135 \mu\text{Hz}$ – there will be significant differences in how the amplitudes vary in time, due to the excitation.

2.2 The solarFLAG simulator

2.2.1 General overview

The solarFLAG data sets simulate full-disc ‘Sun-as-a-star’ Doppler velocity observations of the Sun, such as those made by the BiSON and GOLF. The solarFLAG data sets are made with a full cohort of simulated low- l modes, covering the ranges $0 \leq l \leq 5$. The frequencies of the modes come from a standard solar model of the user’s choice. A surface term is also added to these frequencies, based on polynomial fits to the differences between the standard model frequencies and frequencies from analysis of BiSON and GOLF data. For data presented in this paper the frequencies came from model BS05(OP) of Bahcall, Serenelli & Basu (2005).

A data base of p-mode power and line width and asymmetry estimates, obtained from analyses of GOLF and BiSON data, was used to guide the choice of the other mode input parameters. The hypothetical solarFLAG instrument was assumed to make its observations from a location in, or close to, the ecliptic plane. This is the perspective from which BiSON (ground-based network) and GOLF view the Sun. The rotation axis of the star is then always nearly perpendicular to the line-of-sight direction, and only a subset of the $2l + 1$ components of the non-radial modes are clearly visible: those having even $l + m$. These components are represented explicitly in the solarFLAG time series. The visibility for given (l, m) also depends, although to a lesser extent, on the spatial filter of the instrument (e.g. Christensen-Dalsgaard 1989; Broomhall et al. 2009). Here, we adopted BiSON-like visibility ratios.

Fig. 2 shows a schematic representation of the solarFLAG simulator, the details of which we discuss next.

Before the simulator can begin its calculations, all the input parameters must be specified. The p-mode input parameters are held in a control file (with mode powers having been suitably scaled to reflect the visibility filter of the Sun-as-a-star observations). The other input parameters relate to the granulation-like noise, and other background noise, and are specified at run time by command-line inputs. Time series of granulation-like kicks are used to excite the modes, and are also used to give correlated background noise. The time-scale and standard deviation of this noise must be chosen. A single constant, ρ , is also set and this fixes the coefficient of correlation for the correlated excitation, and the correlation with the background noise. This gives the user the flexibility to ‘tune’ the asymmetry of the mode peaks (as well as other effects arising from the correlations), since higher correlation gives larger peak asymmetry. The input parameters of other sources of noise, such as the

Figure 2. A schematic representation of the solarFLAG simulator.

standard deviation, σ_{psn} , of the photon-shot noise, are also specified at this point.

The main block of the solarFLAG simulator code generates output time series for each (l, m) . It makes the time series of all the

overtones, n , for that (l, m) . The modes are excited by time series of granulation-like noise. For each (l, m) there is a noise time series that gets used again and again to excite the overtones: this is the *correlated gr. noise* time series. When the user fixes the correlation

at less than 100 per cent, a suitable fraction of *uncorrelated gr. noise* (see schematic in Fig. 2) must be mixed in when the modes are excited. These uncorrelated noise time series must be seeded with completely independent random noise and made afresh for each mode. After all overtones have been made, a suitably scaled mixture of correlated and uncorrelated noise is added to the time series. This gives the cumulative time series for each (l, m) . Execution of the main code block is then repeated for a total of 14 different (l, m) combinations.¹

After time series for all the (l, m) have been made they are added together. The final stage of processing adds in the simulated time series of other uncorrelated noise sources. The schematic includes a contribution from photon-shot noise. Other sources may also be included at this point (e.g. active-region noise and instrumental noise) but this was not done for the data used in this work.

A complete description of the different elements that make up the complex frequency amplitude (and frequency power) spectrum of a solarFLAG time series is presented in Appendix A.

We now turn to specifics on how the data are made, and how the effects of correlations are quantified.

2.2.2 Achieving correlated excitation of overtones

As noted above, we use time series of granulation-like noise to excite the overtones, n , of the same l and m . Later we will refer to this as ‘correlated noise’. The granulation-like noise is made using white-noise input to a low-order, autoregressive model of the AR[1] type (e.g. see Koen 2005), i.e.

$$u(t) = u(t - \Delta t) \exp -\Delta t/\tau + \delta(t). \quad (1)$$

Here, $u(t)$ is the output time series of granulation-like noise; τ is the time constant of the model, which should be given a value to mimic the lifetime of the solar granulation (see below); Δt is the cadence at which samples are generated by the autoregressive process; and finally the $\delta(t)$ are random numbers drawn from a normal distribution having zero mean and sample standard deviation $\sqrt{\sigma^2 \Delta t/\tau}$, where σ fixes the amplitude of the granulation-like noise. The power spectral density, $n(\nu)$, of this granulation-like noise follows the approximate but useful Harvey (1985) power-law model i.e.

$$n(\nu) = \frac{4\sigma^2 \tau}{1 + (2\pi\nu\tau)^2}. \quad (2)$$

The default option would be to excite all overtones of the same (l, m) with the same time series $u(t)$, meaning the excitation of the overtones would be 100 per cent correlated (as in Section 2.1). As indicated earlier, we decided to add to the solarFLAG simulator the option to let the user choose a coefficient of correlation, ρ , common to all modes. We describe in Section 2.2.2 below how the asymmetry of a mode is fixed by the combination of ρ , the noise background, and the input mode parameters.

Overtones of the same (l, m) are excited by a composite time series of kicks, $u(t)$, where

$$u(t) = |\rho|^{1/2} u_c(t) + (1 - |\rho|)^{1/2} u_u(t). \quad (3)$$

¹ Data on modes cover the range $0 \leq l \leq 5$; and there are $l+1$ m -components visible in the Sun-as-a-star data at each l . This actually sums to give 21 combinations of (l, m) . However, components at $l=4$ and 5 that do not have $l = \pm m$ are so weak that they are not included in the time-series generation. Discounting these modes reduces the number of combinations to 14.

In the above, the $u_c(t)$ and $u_u(t)$ are both time series of granulation-like noise. However, the $u_c(t)$ are kicks that are *common* to all the overtones (the *correlated gr. noise* time series shown in Fig. 2), while the $u_u(t)$ are completely *independent* (shown as *uncorrelated gr. noise* in Fig. 2).

So far in this section we have not said anything about tuning the correlation with the background noise. Correlated background noise is provided by the correlated part of the $u(t)$. After all overtones of the same (l, m) have been made, the $u(t)$ is added to the time series to give background noise (having first been appropriately scaled). In the formulation presented by Chaplin et al. (2008), another coefficient, α , was used to describe the correlation with the background noise. Here, we assume this correlation has the same size as the correlation between the excitation of different modes, so that we set $\alpha = \rho$. The a priori choice of ρ therefore also fixes the correlation of the overtones with the background noise.

How do we juggle all the choices implied by the descriptions above to give a time series where both the S/N in the modes and the asymmetry of the modes are realistic? That is the question we turn to next. We use two main factors to contribute to the peak asymmetry: correlation of the modes with the background noise, and correlated excitation of the overtones. We will now consider each of these in turn as it relates to controlling the asymmetry of the artificially generated modes.

2.2.3 Asymmetry due to correlated noise

The contribution to the asymmetry from correlations with the noise background may be dealt with analytically in a fairly straightforward manner. Let b_{nlm} be the asymmetry, due to correlated noise, of the n th overtone of a given (l, m) . The frequency of this mode is ν_{nlm} . The asymmetry (as it appears in the mode parametrization function in equation 12) is given by (Toutain et al. 2006)

$$b_{nlm} = \rho \sqrt{\frac{n_{lm}(\nu_{nlm})}{H_{nlm}}}, \quad (4)$$

where H_{nlm} is the height of the mode peak in the frequency power spectrum. Notice the presence of the factor ρ in equation (4): it is needed because it is only the *correlated* part of $n_{lm}(\nu_{nlm})$ that contributes to the asymmetry. $n_{lm}(\nu_{nlm})$ is the granulation-noise background at the frequency of the resonance, which is just (cf. equation 2)

$$n_{lm}(\nu_{nlm}) = \frac{4\sigma_{lm}^2 \tau}{1 + (2\pi\nu_{nlm}\tau)^2}. \quad (5)$$

We give σ_{lm} in equation (5) an explicit dependence on the combination (l, m) . This gives us the ability to choose different σ_{lm} for different (l, m) , to take proper account of the different mode visibilities. This plays a key part in tuning the asymmetry, as we now go on to explain.

We have assumed that at a given frequency the relative sizes of the granulation noise and the mode amplitudes on the Sun are independent of degrees l and m . An important consequence of this assumption is that the asymmetries from the correlated noise will be the same for all (l, m) . In the solarFLAG simulator, after making overtones of a given (l, m) , the time series of mode amplitudes is multiplied by a visibility factor, S_{lm} , to mimic the visibility filter of the Sun-as-a-star observations. To preserve the independence of

Figure 3. Impact of choice of ρ on mode peak asymmetry. The left-hand panel shows the asymmetry due to correlation with background noise (Section 2.2.3), the centre panel the asymmetry due to correlated mode excitation (Section 2.2.4), and the right-hand panel the total asymmetry, due to the complex interaction of the two contributions. All panels show results for $\sigma = 0.2 \text{ m s}^{-1}$ and $\tau = 260 \text{ s}$. Different line-styles show results for different ρ : dotted for $\rho = -0.20$; solid for $\rho = -0.36$, and dashed for $\rho = -1.00$.

Figure 4. Impact of choice of τ on mode peak asymmetry. The left-hand panel shows the asymmetry due to correlation with background noise (Section 2.2.3), the centre panel the asymmetry due to correlated mode excitation (Section 2.2.4), and the right-hand panel the total asymmetry, due to the complex interaction of the two contributions. All panels show results for $\sigma = 0.2 \text{ m s}^{-1}$ and $\rho = -0.36$. Different line-styles show results for different τ : dotted for $\tau = 130 \text{ s}$; solid for $\tau = 260 \text{ s}$, and dashed for $\tau = 520 \text{ s}$.

asymmetry on (l, m) we must therefore also modify the σ_{lm} of the granulation noise, according to

$$\sigma_{lm} = \sigma \frac{S_{lm}}{\sum_{\text{all}(l,m)} S_{lm}}. \quad (6)$$

Here, σ is the equivalent standard deviation of the cumulative granulation noise background, $n(\nu)$, so that

$$n(\nu) = \frac{4\sigma^2\tau}{1 + (2\pi\nu\tau)^2} = \sum_{\text{all}(l,m)} \frac{4\sigma_{lm}^2\tau}{1 + (2\pi\nu\tau)^2}. \quad (7)$$

We have now covered all the steps needed to see how the asymmetry is tuned.

The p-mode parameters (i.e. frequencies, heights, line widths) are fully specified on input. The asymmetry is then tuned by three free parameters: the coefficient of correlation ρ ; the equivalent standard deviation of the cumulative granulation noise background, σ ; and the time-scale of the granulation, τ . Once σ and τ have been chosen, we use equation (6) to determine the σ_{lm} for each combination of (l, m) . This in turn determines the asymmetry, due to the correlated noise, according to equations (4) and (5).

The left-hand panel of Fig. 3 shows resulting asymmetries given by the parameters $\sigma = 0.2 \text{ m s}^{-1}$ and $\tau = 260 \text{ s}$. The dotted line shows results for $\rho = -0.20$, the solid line for $\rho = -0.36$, and the dashed line for $\rho = -1.00$. Here, we used negative ρ to give the negative asymmetries observed in Doppler velocity data. The magnitude of the asymmetry increases in direct proportion as the magnitude of ρ is increased (this follows trivially from inspection of equation 4). When $\rho = -0.36$ we get asymmetries that resemble quite closely those seen in the real observations. The asymmetries

are largest at low and high frequency, because the ratio of the granulation background to mode amplitude is highest there. Very similar looking plots are given by varying σ ; inspection of equations (4) and (5) indicates that the magnitude of the asymmetry is proportional to σ .

The left-hand panel of Fig. 4 shows instead the effect of varying τ (for fixed $\sigma = 0.2 \text{ m s}^{-1}$ and $\rho = -0.36$). The dotted line shows results for $\tau = 130 \text{ s}$, the solid line for $\tau = 260 \text{ s}$, and the dashed line for $\tau = 520 \text{ s}$. Changes in τ clearly have a less pronounced effect on the asymmetries than do similar relative changes in ρ .

2.2.4 Asymmetry due to correlated excitation

The contribution to the asymmetry overtones of the same (l, m) from the overlapping wings of the overtones of the same (l, m) that have been excited with correlated noise is fixed by the frequency separations, line widths and relative heights of those overtones, and the choice of ρ . We later refer to this as ‘correlated excitation.’ This contribution is far less amenable to a neat analytical description than the correlated noise contribution described above. For a given overtone it must describe, in the complex plane, the contribution of power from all other overtones to which it is correlated. A full description of the complex frequency amplitude spectrum, and frequency power spectrum, is given in Appendix A. We may use the analytical descriptions there to calculate numerically the asymmetries given to the mode peaks by the correlated excitation. The resulting estimates are shown in the central panels of Figs 3 and 4. Note that the correlated excitation contribution does not depend on the sign of ρ (only the magnitude), and is independent of the choice of τ (and σ).

At lower frequencies the asymmetry given to modes by the correlated excitation is negative. Here, there are a larger number of other overtones at frequencies above a mode than there are below it. The reverse is true at higher frequencies, where the asymmetry is positive. To understand the sign of the asymmetry in each region, we refer back to the central panel of Fig. 1. There, we had two correlated modes. The lower-frequency mode had negative asymmetry, due to the correlated impact of its higher-frequency counterpart; the higher-frequency mode showed asymmetry of opposite sign, due to the lower-frequency mode. We see this pattern repeated in the full spectrum of overtones with the crossover in behaviour located where p-mode power is a maximum.

The right-hand panels of Figs 3 and 4 show the full asymmetry given to the modes. It is worth pointing out that this is not simply the sum of the asymmetries given by correlated noise (left-hand panels) and correlated excitation (central panels). Matters are a bit more complicated, because the correlated noise and correlated excitation act together in a non-trivial manner in the complex plane.

One final thing to mention with regard to the asymmetries is that there is actually a third contribution, which comes from the fact that the frequency response of the granulation-like noise used to excite modes is not white (i.e. flat). The response in the vicinity of each resonance of course rises with decreasing frequency, meaning there will be a small negative asymmetry contribution. The effect is, however, negligible. (The impact of the non-white response of the excitation is modelled in Appendix A.)

We have seen that the input frequencies, powers and line widths of overtones, and the correlation coefficient ρ , fix the asymmetry contribution from the correlated excitation. When in addition the σ and τ of the granulation-like noise are chosen the asymmetry contribution from the correlated noise is also completely specified. Once any additional uncorrelated background has been specified, these choices in principle also fix the observed S/N ratio of the mode peaks. For the data analysed below, we chose $\tau = 260$ s – the characteristic solar time-scale – and $\rho = 0.36$, chosen to give a good match to the asymmetry measured in previous work. These values correspond to the solid curves in Figs 3 and 4. We also chose $\sigma = 0.2 \text{ m s}^{-1}$ for correlated and $\sigma = 0.25 \text{ m s}^{-1}$ for uncorrelated noise. We now go on to describe the S/N, in terms of the commonly used background-to-height ratio, β .

2.2.5 Background-to-height ratio in mode peaks

The intrinsic background-to-height ratio, β_{nlm} , in a mode peak is given by

$$\beta_{nlm} = \frac{N(v_{nlm})}{H_{nlm}}, \quad (8)$$

where $N(v_{nlm})$ is the total background at the resonance. It has a part due to the cumulative granulation background of all (l, m) that appear in the time series. There will also be parts due to other sources of uncorrelated noise. In what follows we specify just one of the possible components: photon-shot noise, N_{psn} . This shot noise is specified by its variance, σ_{psn}^2 . The power spectral density due to this noise term is then

$$N_{\text{psn}} = 2\sigma_{\text{psn}}^2 \Delta t. \quad (9)$$

The total background may therefore be written as

$$N(v_{nlm}) = N_{\text{psn}} + \sum_{\text{all}(l, m)} n_{lm}(v_{nlm}), \quad (10)$$

or explicitly

$$N(v_{nlm}) = 2\sigma_{\text{psn}}^2 \Delta t + \frac{4\sigma^2 \tau}{1 + (2\pi v_{nlm} \tau)^2}. \quad (11)$$

The above implies that the intrinsic background-to-height ratio is in principle specified by the choice of four parameters (when no instrumental or other sources of noise are specified): the height of the mode, H_{nlm} ; the σ and τ of the granulation-like noise background; and the σ_{psn} of the uncorrelated shot-noise background.

2.2.6 Solar-cycle effects

It is a fairly straightforward matter to include solar-cycle-like variations in the artificial time-series data. Changes in frequency may be introduced by varying in time the natural frequency of the damped, driven oscillator used to simulate each mode component; whilst changes in amplitude and damping may be introduced by varying in time the damping constant for the oscillator. In all cases we require changes that mimic those seen in the real data. We base these changes on an artificial proxy of the real 10.7-cm RF variations observed over solar activity cycle 23. Variations in the frequencies and damping rates are programmed to follow the artificial proxy, with appropriate calibration constants used to fix the absolute scale of the variation for each parameter.

Here, we used a high-order polynomial fit to the 10.7-cm RF versus time – for the 11 yr period beginning 1997 February 3 – to provide the proxy activity, which we refer to as X in later sections of the paper. The purpose of the fit was to provide a smooth proxy that captures the main 11 yr cycle while removing shorter-term variations. Calibration constants for the frequencies and damping rates of every simulated mode were based on results from previous analyses of real BiSON data. The calibration constants for the frequencies depend not only on frequency – the higher the frequency, the larger is the solar-cycle frequency shift – but also the angular degree, l , and the azimuthal order, m , of each simulated component (e.g. Chaplin et al. 2004). The dependence on (l, m) is due to the strong latitudinal dependence of the real near-surface activity, which drives the changes in the mode parameters.

As noted earlier, due to the relatively large uncertainties in the measurements little information is available on the frequency and (l, m) dependence of detected solar-cycle changes in the mode amplitudes and damping rates for low-degree modes, but we know from medium- and high-degree studies (e.g. Komm et al. 2000a; Howe et al. 2004) that the changes are strongest in the middle of the five-minute band. However, the relative sizes of the variations are consistent with the hypothesis that both are the result of net changes to the damping only (e.g. Chaplin et al. 2000). Here, we have introduced solar-cycle variations that adhere to this finding, so that in our simulations we change only the damping constant. Further information on the introduced variations is included later in Section 4.

Finally, we note that variations in amplitude lead to solar-cycle-like changes in the mode peak asymmetries of our artificial data. This is because the height-to-background ratios of the mode peaks changes with the artificial proxy, due to variations in the intrinsic amplitudes (the simulated granulation and shot-noise components were stationary in time). Again, information on the programmed variations is in Section 4.

3 FITTING METHODS

Part of the object of the current work is to compare the results from two different fitting algorithms. Both of these are conventional in the sense that they use the ‘pairwise’ fitting approach in which each $l = 0/2$ and $1/3$ pair of peaks is fitted independently with its own background and parameter set.

The two methods we use are labelled as MLE (Maximum-Likelihood Estimate) and MCMC (Markov Chain Monte Carlo) but are not solely described by the method of optimization. There are differences in the power spectrum model that we detail below.

3.1 Mode parametrization

Common to both fitting methods, each mode component of degree l , radial order n and azimuthal order m is represented by a peak profile described by the equation

$$P(\xi) = \left(\frac{h}{1 + \xi^2} \right) \times [(1 + b\xi)^2], \quad (12)$$

where

$$\xi = 2(\nu - \nu_0)/\Gamma, \quad (13)$$

ν_0 is the frequency of the Lorentzian component, Γ its width, h its height, and b a fractional parameter characterizing the asymmetry. This expression simplifies to the normal Lorentzian for $b = 0$. This is the profile of Nigam & Kosovichev (1998), with the quadratic term in b suppressed for greater stability (Fletcher et al. 2009). For this work, we have chosen to parametrize the peak in terms of the amplitude, A , where

$$h = A^2/(\pi\Gamma). \quad (14)$$

Neglecting the small asymmetry term, $A^2 \equiv \pi\Gamma h$ is proportional to the integrated energy of the mode (see, for example, Komm et al. 2000b, but note that their A is our h), and the energy supply rate dE/dt is given by

$$dE/dt \propto A^2\Gamma \propto \Gamma^2 h. \quad (15)$$

The likelihood spaces (or the posterior probability distributions) are lognormal for the mode width and mode amplitude parameters, so we follow the common practice of varying the logarithms of the amplitude and width parameters rather than their raw values.

The background term is a constant for each $l = 0/2$ or $1/3$ pair, and in both algorithms the asymmetry term α is also common to all of the peaks in each pair. In the code used for the MLE algorithm there is a separate Γ for each n and l , while the MCMC code uses the same Γ for every peak within an $l = 0/2$ or $1/3$ group. There is a separate amplitude parameter for each n and l , but the relative heights of the rotationally split components of different m within an n, l group are fixed in the MLE code at 1:0.41 for the $m = 2$: $m = 0$ component ratio for $l = 2$ and 1: 0.19 for the $m = 3$: $m = 1$ ratio for $l = 3$. In the MCMC code they are allowed to vary; the priors are set at a normal distribution of width 0.2 centred on 0.4 for the $l = 2, m = 0$: $l = 2, m = 2$ ratio and the $l = 3, m = 1$: $l = 3, m = 3$ ratio while for the $l = 3, m = 1$: $l = 3, m = 3$ ratio we use a flat-topped Gaussian centred on 0.2 that is flat between 0.1 and 0.3 and has Gaussian sides that drop off with width 0.1. As we will see, the differences have no obvious impact on the final results.

The rotational splitting parameter was fixed at $0.4 \mu\text{Hz}$ for the MLE algorithm, while for the MCMC it was left as a free parameter with a Gaussian prior with centre 0.4 and width $0.05 \mu\text{Hz}$.

3.2 Optimization

The Fourier power spectrum has the statistics of χ^2 with two degrees of freedom rather than Gaussian statistics. The quantity to be minimized is the negative log of the likelihood function (Anderson, Duvall & Jefferies 1990),

$$S \equiv -\ln L = \sum_i \left\{ \ln M_i + \frac{O_i}{M_i} \right\}, \quad (16)$$

where M is the model, O is the observed spectrum, and the sum is over the data points (frequency bins) in the fitting window.

In the MLE approach we simply find the values of the model parameters that give the minimum value of S and estimate the errors by inverting the Hessian matrix. For MCMC, on the other hand, we take the standard deviation of the posterior distribution for each parameter.

The MLE code used here was a slightly modified version of the one developed for use with BiSON data and described by Chaplin et al. (1999). The MCMC code is a vanilla Markov Chain Monte Carlo approach (see Gelman et al. 2003; Davies et al. 2014a, for more details).

3.3 Treatment of window function

It is a well-known problem in ground-based helioseismology that the Earth’s rotation causes the observations to be interrupted with a 24 hr periodicity, which gives rise to ‘sidelobe’ peaks at multiples of $11.57 \mu\text{Hz}$ from each mode. At low degrees, where $l = 0/2$ and $1/3$ pairs are separated in frequency by a similar amount, this is particularly problematic. The effect is greatly reduced by the use of multisite observations but cannot be entirely eliminated, and it must therefore be taken into account in the data analysis.

As is conventional, we describe the modulation of the time series by a ‘window function’ – a time series corresponding to the observations in which each time sample is represented by a 1 for an observation and a 0 for missing data. The observed spectrum can be considered as the convolution of the power spectrum of this window function with that of the uninterrupted observations.

In the MLE code used here, the sidelobe peaks are represented in the fitting model by peaks of the same width and asymmetry as each main peak, at $\pm 11.57 \mu\text{Hz}$ from its central frequency. The initial guess for the relative amplitude of this sidelobe peak was estimated from a spectrum in which a single sine wave was convolved with the window function of the time series. For modes in the main five-minute band (between 2.0 and 3.5 mHz) this relative height was then allowed to vary in the first pass of fitting and held fixed for modes outside this frequency range. The mean value of the fitted sidelobe after the first pass was then used as the fixed value for all modes in the final pass. This allowed a larger number of successful fits to be recovered than if the sidelobe height had been left as a free parameter throughout.

The MCMC algorithm uses a different, more accurate but more computationally expensive approach in which the model spectrum is convolved with the window function spectrum at each optimization step. This method can also account for the broadening of the peak as the duty cycle of the observations is decreased; however, neither this approach nor the one used with the MLE algorithm take into account the correlations between frequency bins introduced by the missing data. Note that the treatment of the sidelobes is related to the specifics of the codes used here and is not intrinsic to the MLE or MCMC approach; in principle it would be possible to perform

MLE with the convolution, as was indeed done for example by Jiménez-Reyes et al. (2007) and Fletcher et al. (2009).

4 ARTIFICIAL DATA TESTS

The object of the tests described here is to check that the MLE and MCMC algorithms give consistent and reliable results for the variation of the different mode parameters with the activity index.

The 11-yr time series generated as described in Section 2 above was divided into contiguous 365 d segments and each segment was fitted using both the MLE and MCMC algorithms. The series were fitted both with 100 per cent duty cycle and with a ‘realistic’ duty cycle based on that of the BiSON time series for the 11 years starting with the notional start date of the artificial series, 1997 February 3. For each mode we then take an error-weighted mean over all the data sets and subtract this from the values to obtain the shifts, finally using a linear least-squares fit to the activity index to derive the sensitivity for each mode for comparison with the input values.

For the MLE method only, three additional, more extensive tests were carried out. In one test, 25 additional realizations of the 11-yr artificial data series were fitted in 1 yr segments, while in another each 1 yr segment of a single realization of the artificial data was fitted with the duty cycle corresponding to each of the 21 one-year BiSON time series. The term ‘realization’ here refers to a realization of the noise exciting the simulated modes. The extra fits make it easier to distinguish the systematic errors from the random noise; by using 25 realizations we reduce the random errors by a factor of 5, which is sufficient to clarify the trends. In the real observations we have 22 years of data, so the random errors on the mean parameters and sensitivities are a factor of $\sqrt{2}$ smaller than for a single realization of the simulated data, or approximately three times larger than for 25 realizations. Finally, to help rule out the possibility that the frequency shifts are caused by asymmetry changes, we prepared another set of artificial data realizations where the input amplitude and width changes – and hence the asymmetry changes – were as before but the input frequency change was set to zero. We note that to generate frequency shifts as large as those observed, the asymmetry shift would need to be extremely and unrealistically large, as discussed below in Section 4.2.

4.1 Results

4.1.1 Frequencies

Figs 5(a)–(d) show the input and fitted values of $d\nu_{nl}/dX$, the frequency shift per unit of the activity index X for each fitting method. Note that the differential notation we use throughout is a convenient shorthand for the regression slope. In the case of the artificial data the linear relationship between activity and mode parameters is designed to be exact; the real data may exhibit small deviations from this but it is still a useful measure of the sensitivity to first order. Except for a few modes at the higher end of the frequency range, the results from both of the algorithms agree within errors with the input values; the mean ratio between the fitted and input $d\nu_{nl}/dX$ for $0 \leq l \leq 2$ and $2000 \mu\text{Hz} \leq \nu_{nl} \leq 3500 \mu\text{Hz}$ is consistent with unity in each of the four cases.

Panels (e) and (f) of Fig. 5 illustrate the sensitivity of the frequency to activity for the tests with 25 realizations of the artificial data, with panel (f) showing the results for the test with no input frequency variation. Apart from a few modes at frequencies above

3.7 mHz where the S/N ratio is low, we see no frequency change in the data set where none was explicitly introduced. Fig. 5(g) shows the sensitivity of the frequency to the duty cycle. These results were obtained by performing for each mode a multiple linear regression in which the independent variables were the activity index X and the duty cycle F and the dependent variable was the mode frequency. The results in Fig. 5(e) clearly show the slightly lower $d\nu/dX$ for the $l = 0$ modes – a difference that is lost in the noise for the single-realization test.

In another test of the MLE fitting we applied each of the 21 one-year window functions from the BiSON network for 1993–2013 to each of the 11 years of the solarFLAG time series. These results were used to check for any bias in the sensitivity of the frequency to activity level due to the different window function, as follows. The results were divided into 21 sets in which the 11 one-year spectra all had the same duty cycle but different activity indices. For each set a regression was performed for each mode, with the frequency shift as the dependent and the activity index as the independent variable, to obtain a set of slopes $d\nu_{nl}/dX$ for each value of F . Finally, a regression was performed for each mode with the duty cycle F as the independent variable and the $d\nu/dX$ value as the dependent one. The results are shown in Fig. 5(h). The sensitivity of the frequencies to activity index shows no systematic bias due to the window function.

4.1.2 Amplitudes

For mode amplitude A (defined as in equation 14 above) and width Γ we work with the natural logarithm of the parameter; a shift in the logarithmic parameter is equivalent to a fractional shift in the raw parameter.

Figs 6(a)–(d) show the sensitivity, $d \log A_{nl}/dX$, of the natural logarithm of the amplitude parameter A_{nl} to activity index, obtained using the two methods for artificial data with 100 per cent and realistic duty cycles, compared with the input values. The individual $d \log A_{nl}/dX$ for the two methods are well correlated, with correlation coefficients of 0.77 for the 100 per cent duty cycle and 0.70 for the realistic one over the 47 modes common to both sets; the threshold for 0.1 per cent probability of the same result arising for two unrelated populations of 47 samples is 0.47.

The mean ratios between fitted and input dA_{nl}/dX values are 0.84 ± 0.08 for MCMC and 0.85 ± 0.07 for MLE in the 100 per cent duty cycle case and 0.83 ± 0.09 for MCMC and 0.79 ± 0.10 for MLE in the realistic duty cycle case. This result suggests that the sensitivity of the amplitude to the activity index is being systematically underestimated by both methods. However, as we shall see below, when we average overall realizations of the artificial data we find that the average bias is not as severe as for this case.

In panels (e) and (f) of Fig. 6 we show the sensitivity of the amplitude measurements to activity index for the multiple-realization tests with MLE fitting; panel (f) shows the results for the artificial data with no input frequency variation. The lack of frequency variation does not significantly affect the amplitude variation. The results with multiple realizations clearly reveal the frequency dependence of the shifts, and also confirm the systematic underestimation of the sensitivity. The sensitivity to duty cycle (panel g) has a small non-zero value that does not depend strongly on frequency; this is most likely due to the power in higher order daily sidelobes that is not accounted for in the fitting.

Fig. 6(h) shows the variation with duty cycle of the amplitude shift per unit activity, from the test with MLE fitting where the window

Figure 5. Sensitivity of fitted frequency to activity level and fill for artificial data. Panels (a)–(d) show fitted (symbols) and input (lines) $d\nu_{nl}/dX$ for MLE (a, c) and MCMC (b, d) fits to 11 one-year sets of artificial data at different activity levels, mimicking the solar cycle. The results are for 100 percent (a, b) and realistic (c, d) duty cycle. The lower two rows, for MLE only, show the results of the tests with multiple realizations of the artificial data. Panel (e) shows the sensitivity of the frequency to activity index, X ; panel (f) is similar to panel (e) but for a set of artificial time series where there is no input frequency variation. Panel (g) shows the sensitivity to duty cycle, F , for the test with 25 realizations of the artificial data. Panel (h) shows the rate of change with duty cycle F of the sensitivity $d\nu/dX$ of mode frequency to activity index, for the test where each of the 21 one-year BiSON window functions was applied to each of the 11 years of the solarFLAG time series. Black circles represent $l = 0$, blue diamonds $l = 1$, red squares $l = 2$, and green open triangles $l = 3$.

Figure 6. Sensitivity of mode amplitude to activity index and duty cycle. Panels (a)–(d) show fitted (symbols) and input (lines) $d \log A_{nl}/dX$ for MLE (a, c) and MCMC (b, d) fits to 11 one-year sets of artificial data at different activity levels, mimicking the solar cycle. The results are for 100 percent (a, b) and realistic (c, d) duty cycle. The bottom two rows, for MLE only, show the results of the tests with multiple realizations of the artificial data. Panels (e) and (f) show the sensitivity of the amplitude to activity index, X , with panel (f) showing the results for the test with no input frequency variation. Panel (g) shows the sensitivity to duty cycle, F , for the test with 25 realizations of the artificial data. Panel (h) shows the rate of change with duty cycle F of the sensitivity $d \log A_{nl}/dX$ of mode amplitude to activity index, for the test where each of 21 one-year BiSON window functions was applied to each of the 11 years of the solarFLAG time series. Black circles represent $l = 0$, blue diamonds $l = 1$, red squares $l = 2$, and green open triangles $l = 3$.

Figure 7. Fitted amplitude (left) and line-width (right) shifts averaged over 25 realizations of the artificial data with realistic duty cycle for each mode between 2.0 and 3.5 mHz, plotted against the input value. Black circles represent $l = 0$, blue diamonds $l = 1$, and red squares $l = 2$. The correspondingly coloured lines represent least-squares fits for each l .

function was varied. Apart from a few low-frequency modes and the $l = 3$ modes, the results show no systematic bias of the sensitivity measurement due to the window function.

The apparent underestimate of the amplitude shifts is of some concern. However, when we consider the MLE fits to the larger cohort of artificial time series, we obtain a more favourable result. Fig. 7(a) shows the fitted amplitude shifts for data with realistic fill, averaged over 25 realizations of the artificial data for each mode and plotted as a function of the expected shift. In this case we find, for modes between 2.0 and 3.5 mHz, that the slope of a linear fit of fitted versus input shifts is 0.92 ± 0.03 , 0.95 ± 0.02 , and 0.89 ± 0.02 for $l = 0, 1$, and 2 , respectively, which suggests that the realization chosen for the main plots was a particularly ‘unfortunate’ example and the real systematic underestimate of the sensitivity is likely to be less than 10 per cent. We emphasize that the size of the fractional shifts is not affected by any systematic errors in the underlying values as long as these are not activity-dependent.

4.1.3 Line width

Figs 8(a)–(d) show the input and fitted values of $d \log \Gamma_{nl}/dX$, the rate of change of line width with the activity index X , for the two methods. As for the amplitude measurements, the individual-mode sensitivity results for the two methods are well correlated with one another, with a correlation coefficient of 0.59 in the 100 per cent duty cycle case and 0.71 for the realistic duty cycle, over 27 values. The slope of a fit of MCMC versus MLE values is 0.96 ± 0.48 for 100 per cent duty cycle and 1.08 ± 0.70 for realistic duty cycle, which indicates that the values are consistent; furthermore, for each method the $d\Gamma/dX$ values for the 100 per cent and realistic duty cycle cases are consistent. The results of the tests with multiple realizations are shown in Figs 8(e) and (f). Again, the frequency variation of the sensitivity is clearly seen in Figs 8(e) and (f), with panel (f) being for the case with no input frequency variation. Again, the lack of frequency variation has no significant effect on the width variations. Fig. 8(g) shows that there is a small, more or less frequency-independent, effect of the duty cycle on the fitted line-width, which is to be expected for the MLE algorithm due to the crude handling of the window function effects. Fig. 8(h) shows

that the results of the MLE test with multiple duty cycles applied to the same realization of the artificial data; the duty cycle has no significant effect on the sensitivity of the line width to activity-related variations, so as long as the two factors are not correlated we can treat the effects as independent of one another.

As in the case of the amplitude, $d\Gamma/dX$ appears systematically underestimated, with a mean ratio between fitted and input values of 0.81 ± 0.09 for 100 per cent and 0.79 ± 0.10 for realistic duty cycle for MLE fits and 0.84 ± 0.08 for 100 per cent fill and 0.83 ± 0.09 for realistic duty cycle from MCMC. For the test with MLE fits to 25 realizations of artificial data with realistic duty cycle (Fig. 7b), we obtain slopes of 0.87 ± 0.04 , 0.88 ± 0.03 , and 0.77 ± 0.03 for $l = 0, 1$, and 2 , respectively; as with the amplitude, this suggests that our main sample realization was particularly unfavourable, but we might still be failing to recover the $l = 2$ shifts completely.

4.1.4 Energy supply rate

We recall from equation (15) that the energy of a mode is proportional to the square of its amplitude, and the rate of energy supply to the mode dE_{nl}/dt is proportional to $h_{nl}\Gamma_{nl}^2$; so for the logarithmic quantity we have

$$\log dE_{nl}/dt = 2 \log \Gamma_{nl} + \log A_{nl} + \text{const}, \quad (17)$$

where const includes a contribution from the mode mass and visibility function.

The energy supply rate is believed not to vary with the activity level, and the solarFLAG data were designed to reflect this. In other words, we expect $2dA/dX + d\Gamma/dX = 0$. As the $d \log A_{nl}/dX$ and $d \log \Gamma_{nl}/dX$ values are noisy, we check for this by performing a weighted least-squares fit with the individual $\delta \log A_{nl}$ and $\delta \log \Gamma_{nl}$ as x - and y -values. For the $l = 0, 1$ values between 2.0 and 3.5 mHz we obtain slopes of -1.8 ± 0.58 and -1.47 ± 0.90 for the MLE fitting with 100 per cent and realistic duty cycle, respectively, and -1.53 ± 0.48 and -1.34 ± 0.58 for MCMC. Three of these values are within 1σ of the expected value of -2 , and the other is just outside the 1σ range. A fit with $d \log A_{nl}/dX$ as the abscissa and $d \log \Gamma_{nl}/dX$ as the ordinate gives agreement within the (large) error in all cases, and the mean value of $2d \log A_{nl}/dX + d \log \Gamma_{nl}/dX$ is

Figure 8. Sensitivity of fitted line-width to activity index and fill, for artificial data. Panels (a)–(d) show fitted (symbols) and input (lines) $d \log \Gamma_{nl}/dX$ for MLE (a, c) and MCMC (b, d) fits to 11 one-year sets of artificial data at different activity levels, mimicking the solar cycle. The results are for 100 per cent (a, b) and realistic (c, d) duty cycle. The bottom two rows, for MLE only, show the results of the tests with multiple realizations of the artificial data. Panels (e) and (f) show the sensitivity of the line width to activity index, X , with the results in panel (f) being for the test with no input frequency variation. Panel (g) shows the sensitivity to duty cycle, F , for the test with 25 realizations of the artificial data. Panel (h) shows the rate of change with duty cycle F of the sensitivity $d \log \Gamma_{nl}/dX$ of mode width to activity index, for the test where each of the 21 one-year BiSON window functions was applied to each of the 11 years of the solarFLAG time series. Black circles represent $l = 0$, blue diamonds $l = 1$, red squares $l = 2$, and green open triangles $l = 3$; for MCMC black circles represent $l = 0/2$ pairs and blue diamonds $l = 1/3$.

Figure 9. Fitted (symbols) and input (lines) mean fractional asymmetry shift $\delta b/b$ as a function of time for MLE (a, c) and MCMC (b, d) fits to 11 one-year sets of artificial data, for 100 per cent (a, b) and realistic (c, d) duty cycle. The dashed lines represent the activity index scaled to the best linear least-squares fit to the asymmetry-shift data.

consistent with zero in all four cases. We can therefore say that the fitting is not significantly biasing the variation of the energy supply rate, in spite of the small bias on the $d \log A_{nl}/dX$ and $d \log \Gamma_{nl}/dX$ values.

4.1.5 Asymmetry

The design of the solarFLAG data is such that the fractional shift in the asymmetry is expected to be the same as that in the line width. Because the asymmetry determinations are noisier, we do not show the results per mode. Instead, we look at the error-weighted average fractional shift relative to the (also error-weighted) temporal mean value for each mode. Fig. 9 shows the temporal variation of the mean fractional asymmetry shift, averaged over all modes between 2 and 3.5 mHz, compared with that which would have been obtained if the fitting returned exactly the input values with the same uncertainties. The correlation coefficients between the expected and measured mean asymmetry shifts are as follows: 0.79 and 0.65 for MLE with 100 per cent and realistic duty cycle, and 0.78 and 0.75 for MCMC. The slopes of linear least-squares fits with the input average fractional shift as the independent variable and the measured values as the dependent one are 1.15 ± 0.33 , 2.35 ± 0.68 for MLE with 100 per cent and realistic duty cycle, and 0.82 ± 0.25 , 1.0 ± 0.49 for MCMC. These results suggest that we can (marginally) detect a shift in the asymmetry of about the right magnitude and sign, but even when averaging over many mode pairs the random errors are too large to make a statement

about any systematic over- or underestimate of its size, particularly in the realistic-duty-cycle case. We can also infer, with caution, that the MCMC fitting is performing somewhat better than the MLE in this respect. This is also true in the case (not shown) where the tests were repeated using MLE with the window function handled by convolution rather than sidelobe fitting, which suggests that the advantage is specifically in the MCMC approach rather than the window function convolution.

4.2 Discussion of the test results

The results of the simulated-data exercise give us confidence that with 11 years of observations we can determine the activity-related shifts in the frequencies without significant bias with either fitting method. For the amplitudes and line widths both methods may underestimate the shifts by around 10 per cent (or more for the $l = 2$ widths). This may be the result of a peak profile model that does not fully capture the subtleties of the mode asymmetry. We also have marginal sensitivity to changes in the asymmetry, even with realistic duty cycle. As we are fitting data with asymmetric peaks using an asymmetric peak profile, we would not expect in this case to see any bias in the frequency shifts due to neglect of the asymmetry, and indeed, even though we only make a weak detection of the asymmetry variation, we are very successful in recovering the input frequency shift – and also in recovering a zero input frequency shift. A brief consideration of the numbers makes this understandable. A frequency shift due to neglect of the asymmetry could not exceed

Figure 10. Yearly duty cycle for the BiSON network over the period covered by our analysis.

the product of the line width and the absolute asymmetry parameter. At 3 mHz the line width is about 1 μHz and the maximum low-degree frequency variation is around 0.31 μHz or one-third of the line width, which would require around a 30 per cent change in the absolute value of the asymmetry – far in excess of what is observed, and so large that the asymmetry would be obvious to the eye rather than requiring subtle analysis.

Given the good agreement between the two techniques it seems unlikely that the use of MCMC fitting will invalidate the major results from the previous two decades or more of MLE fitting of Sun-as-a-star data.

A comparison of the error estimates returned by the two algorithms shows that the uncertainties for the frequency are very similar, as are those for the asymmetry parameter. For amplitude, we find that the uncertainties from MLE are slightly (about 10 per cent) higher than those for MCMC for the $l = 0, 1$ modes and about 15 per cent smaller for $l = 2, 3$ where the signal-to-noise is lower and the modes are made up of more components that are handled differently by the two algorithms, with the MLE fit having fewer independent parameters. For line width, the $l = 0$ and 1 uncertainties from MLE are, not surprisingly, somewhat (20 to 50 per cent) larger than those for the $l = 0/2$ and $1/3$ pairs provided by the MCMC algorithm.

These findings will guide our interpretation of the results of the fitting of observational data.

5 BISON DATA

5.1 Data

Early observations by what was to become the BiSON group have been made since 1976, but until the early 1990s the observations were relatively sparse. The final BiSON network was deployed in 1990–1992 and has been operating ever since, although there have been some changes to the instruments over time. Fig. 10 shows the yearly duty cycle of the network over time. For the purposes of this analysis, we use 22 non-overlapping spectra corresponding to the calendar years 1993–2014. We are fortunate in that over this period the BiSON duty cycle is essentially uncorrelated with the activity level. The data for 1992 and earlier years have much lower duty cycle and were excluded from this analysis.

5.2 Analysis

5.2.1 Data preparation

The BiSON data were prepared as described in Davies et al. (2014b) and divided into 1-yr, non-overlapping time series.

5.2.2 Choice of activity proxy

It is common to express activity changes in terms of a linear relationship to an activity measure, although sometimes additional terms may be used, especially for higher-degree modes. Such analyses have been carried out using, for example, the sunspot number, the 10.7-cm RF, and the global magnetic flux. Bachmann & Brown (1993) found that, for observations over the period 1984–1990, the RF and the Mg II index gave better correlation with medium-degree frequency shifts than the global Kitt Peak magnetic field strength, and Chaplin et al. (2007) also found better agreement with proxies such as the RF flux that have a greater sensitivity to weak flux. On the other hand, for low-degree modes Chaplin et al. (2001) found that the correlations were essentially the same for all six proxies they examined, although there was a slight difference between the sensitivity to the Kitt Peak magnetic index in the falling phase of Cycle 22 and the rising phase of Cycle 23. For local and latitudinally resolved measurements a spatially-resolved proxy is needed, and for this purpose a localized measure of the photospheric magnetic field strength (often called a magnetic activity index, or MAI), tends to be preferred because the irradiance-based measures are not spatially resolved. The relationships between these proxies are not necessarily linear, and may not be consistent over time as instruments are upgraded or observing protocols change. Even magnetic measurements from different instruments may show the effects of small differences in calibration; also, many magnetic indices are only available as Carrington maps that cannot easily be translated into daily indices. These considerations need to be borne in mind when using the measures to compare with long sequences of helioseismic data. For the current work, as we are dealing with Sun-as-a-star data, we have chosen to use the daily RF index averaged over the exact period covered by each spectrum.

5.3 Results

For all variables, we first find the modes that are common to all the data sets being compared and then calculate the shifts for each mode relative to an error-weighted mean over all the data sets. These individual-mode shifts can then be combined in an error-weighted mean to show the overall variation.

5.3.1 Frequency

Figs 11(a) and (b) shows the mean frequency variation with time and with RF index for the two algorithms, and Figs 11(c) and (d) shows $dv_{nl}/d\text{RF}$ for each mode as a function of frequency for each method. The methods agree well, even in the small deviations from the linear fit; in fact, the correlation coefficient between the MCMC and MLE mean frequency shifts is 0.992, better than the correlation between the shifts and the RF index, which is 0.975 in each case. The threshold for significance at the 0.1 per cent level is 0.65 for 22 data points, so all of these are highly significant correlations. It is noticeable in Fig. 11(a) that in both solar minima the frequency points tend to lie below the activity trend-line. The frequency dependence of the shifts is very clearly seen, and there is

Figure 11. Solar-cycle frequency changes from BiSON data 1993–2014. The top row shows frequency variation from MLE (filled symbols) and MCMC (open symbols) fits, as an error-weighted mean over all modes with $l \leq 2$ and frequencies between 2 and 3.5 mHz, as a function of time (a) and RF index (b). The solid and dashed lines represent linear fits to the RF index for MLE (solid) and MCMC (dashed). The bottom row shows $d\nu_{nl}/dRF$ as a function of frequency for MLE (c) and MCMC (d) fits, for $l = 0$ (black circles), $l = 1$ (blue diamonds), $l = 2$ (red squares), and $l = 3$ (green open triangles).

a hint that the $l = 0$ shifts are lower than those for $l > 0$, consistent with the expected behaviour due to the latitudinal distribution of the activity bands.

5.3.2 Amplitude

Figs 12(a) and (b) show the amplitude variation for the two algorithms as a function of time and as a function of RF index, and Figs 12(c) and (d) show $d \log A_{nl}/dRF$ as a function of frequency for each algorithm. Again, the two methods show good agreement, with a correlation coefficient of 0.98 between the two sets of mean shifts, while the correlation with the RF index is -0.88 in each case. The data plotted in Figs 12(a) and (b) have been corrected for the window-function effect on the amplitude, but this makes very little difference to the result. The mean shift changes by (-0.078 ± 0.005) per cent per Radio Flux Units (RFU) from MLE and (-0.071 ± 0.005) per cent per RFU from MCMC, giving a change of about (-8.8 ± 0.6) per cent for MLE or -8.0 ± 0.5 per cent for MCMC from the highest to the lowest RF value. The scatter of the points around the linear fit, together with the agreement between the methods, suggests that effects other than a simple linear dependence on activity may be involved. Curiously, the points for the years 2011–2013 in Fig. 12(a) all fall below the fit to the RF index, suggesting that the modes may have been more strongly suppressed in the rising phase of Cycle 24.

The uncertainties are still too large to allow us to usefully quantify the frequency dependence of the individual-mode shifts, but there

is a visible tendency for the modes in the middle of the five-minute band to be more strongly suppressed by activity. The reduction in χ^2 obtained by fitting a quadratic function of frequency to the values with $l \leq 2$, rather than a constant, corresponds to a probability of 12 per cent for MLE and 23 per cent for MCMC that the result could have been obtained from a random distribution – better than a 1σ result but not at the 2σ level. The shape of the variation, with the greatest sensitivity around 3 mHz where the modes are strongest, is consistent with the findings from resolved helioseismology.

5.3.3 Line width

Figs 13(a) and (b) show the mean shift in $\log \Gamma$ for the two algorithms as a function of time and RF, averaged over the modes between 2 and 3.5 mHz. Figs 13(c) and (d) show $d \log \Gamma/dRF$ as a function of frequency for each method. A direct comparison between the two data sets in this case is more difficult because the MCMC algorithm returns only one width per mode pair while the MLE gives two; also the different treatment of the window function affects the line-width estimates. However, the two sets of mean shifts have a correlation coefficient of 0.946, while the correlation coefficient with the RF index is 0.922 for MLE and 0.951 for MCMC. Although the width shifts are larger than those in amplitude they also have larger uncertainties. Given the uncertainties, we do not see any significant deviation from the linear trend of the mean shifts with RF index. The mean shift is (0.146 ± 0.016) per cent per RFU for MLE and (0.142 ± 0.013) per cent per RFU for MCMC,

Figure 12. Solar-cycle mode amplitude changes from fits to 1-yr BiSON spectra from 1993–2013. The top row shows amplitude variation from MLE (filled symbols) and MCMC (open symbols) fits, averaged over all modes with $l \leq 2$ and frequencies between 2 and 3.5 mHz, as a function of time (a) and RF index (b). The solid and dashed lines represent linear fits to the RF index for MLE (solid) and MCMC (dashed). The bottom row shows $d \log A_{nl}/dB$ as a function of frequency for MLE (c) and MCMC (d), for $l = 0$ (black circles), $l = 1$ (blue diamonds), $l = 2$ (red squares), and $l = 3$ (green open triangles). The dashed curve represents a quadratic function in frequency fitted to the modes with $l \leq 2$.

corresponding to a change of (16.3 ± 1.8) per cent and (15.9 ± 1.5) per cent, respectively, between minimum and maximum RF values.

With the available data, the statistical significance of any frequency dependence of $d \log \Gamma_{nl}/dRF$ is not high. The decrease in χ^2 obtained by fitting a quadratic function of frequency to the values, rather than a constant value, corresponds to a 14 per cent probability for MLE and 10 per cent for MCMC of this result being obtained from random data – again better than a 1σ result but not at the 2σ level.

5.3.4 Energy supply rate

We plot in Fig. 14 the variation with date and RF index of the mean change in the natural log of the energy supply rate E_{nl} , formed by taking the sum of the mean line-width shift and twice the mean amplitude shift. Following Chaplin et al. (2000), the ≈ 90 per cent correlation between the A and Γ errors has been taken into account in the error propagation. The results imply a fractional change in the energy supply rate between minimum and maximum RF of (-1.1 ± 2.2) per cent from MLE and (0.17 ± 1.9) per cent from MCMC. The previous result that there is no significant change in the energy supply rate with activity level is thus confirmed. On the other hand, there is a hint of a decrease in the energy supply rate from 2011–2013, which we see as an unexpectedly large decrease in the mode amplitude in this period.

5.3.5 Asymmetry

Fig. 15 shows the mean fractional shift in the asymmetry parameter b as a function of time and of RF index, for the two algorithms. As the uncertainties are large, we do not show the shifts for individual mode pairs. The coefficient of correlation between mean shift and the RF index is 0.36 for MLE (just short of being significant at the 10 per cent level) and 0.66 for MCMC (significant at the 0.5 per cent level), and the correlation between the two sets of mean shifts is 0.46. The results of the artificial data tests lead us to place more reliance on the MCMC estimate. Even though the result is marginally significant, it is consistent with the results in Jiménez-Reyes et al. (2007), which only included data on the declining phase of the last cycle, where the asymmetry change is most pronounced.

6 DISCUSSION AND CONCLUSIONS

We have analysed 22 years of BiSON data for activity-related changes in the mode parameters using two different algorithms. The analysis was validated by the use of an improved process that allowed for the incorporation of the effects of line asymmetry. The newer MCMC algorithm represents a small improvement over the older MLE, particularly in the search for changes in the asymmetry, but in general there is good agreement between the two algorithms. Both methods accurately reproduce the input frequency shifts in artificial data and slightly underestimate the shifts in mode amplitude

Figure 13. Width variation with activity. Top row shows mean shift in $\log \Gamma_{nl}$ from MLE (filled symbols) and MCMC (open symbols) fits, averaged over all modes with $l \leq 2$ and frequencies between 2 and 3.5 mHz, as a function of time (a) and RF index (b). The solid and dashed lines represent linear fits to the RF index for MLE (solid) and MCMC (dashed). The bottom row shows $d \log \Gamma_{nl} / d RF$ as a function of frequency for MLE (c) and MCMC (d) for $l = 0$ (black circles), $l = 1$ (blue diamonds), $l = 2$ (red squares), and $l = 3$ (green open triangles). For MCMC fits the $l = 0$ value is for the $l = 0/2$ pair and the $l = 1$ for the $l = 1/3$ pair. The dashed curve shows a quadratic fit to the $l \leq 2$ values with frequency as the independent variable.

Figure 14. Mean change in log energy supply rate from MCMC (open symbols) and MLE (filled symbols) fits, averaged over all modes with $l \leq 2$ and frequencies between 2 and 3.5 mHz, as a function of time (left) and RF index (right).

and width, which is perhaps not surprising given the limitations of the model peak profile.

One of the objectives of this work was to check that the previous reports of frequency shifts were not an artefact of underlying changes in the mode asymmetry, as suggested by Korzenik (2013). Contrary to this suggestion, we find that our fitting accurately reproduces the input frequency shifts in artificial data and finds shifts in the real data of the same magnitude as those reported in earlier work.

Furthermore, the MCMC parameter estimation explicitly favours a frequency change over a simple change in underlying asymmetry (with the model we apply) and in any case the asymmetry change required to account for the frequency shift would be unrealistically large.

We have not attempted in this work to reproduce the kind of analysis of frequency shifts used by for example Basu et al. (2012), in which spectra from heavily overlapped time periods were used

Figure 15. Asymmetry variation (mean fractional shift) from MLE (filled symbols) and MCMC (open symbols) fits, averaged over all modes with $l \leq 2$ and frequencies between 2 and 3.5 mHz, as a function of time (left) and RF index (right). The solid and dashed lines represent linear fits to the RF index for MLE (dashed) and MCMC (solid).

and the shifts were averaged in frequency bands; our objective was instead to look for systematic biases in long-term trends. Having eliminated the fear of such biases in the frequency shifts, we will be able to proceed with confidence to a more detailed analysis of the short-term variations in further work. We do see deviations from the linear relationship between frequency and activity that are more evident at low activity and appear consistent with those reported by Fletcher et al. (2010), Basu et al. (2012), and others.

Our amplitude and width changes also agree well with earlier work. Elsworth et al. (1993), who fitted using a symmetric Lorentzian peak profile, report a change of -35 per cent in mode area (height times width) between the 1986 minimum and the 1990 maximum, but based on their linear fit to sunspot number a change of -20 per cent would give a more realistic comparison with our result. As our amplitude is proportional to the square root of their mode area, that would translate to a roughly -10 per cent change in amplitude, which compares reasonably well with our estimate of around -8 per cent. Chaplin et al. (2000), in an analysis covering the period 1991–1997 and averaging over modes between 2.6 and 3.6 mHz, report changes of (24 ± 3) per cent in line width and $-(22 \pm 3)$ per cent in power between maximum and minimum, [corresponding to (-11 ± 2) per cent in amplitude]. The 100 d periods analysed by Chaplin et al. (2000) correspond to a range in RF index of about 150 RFU, compared with the range of 110 in the 365 d average values in our analysis, which would make their 24 per cent change in line-width equivalent to 18 per cent, within a 1σ difference from our 15.5 per cent value. For amplitude, their change of -11 per cent would be equivalent to (8.1 ± 1.1) per cent, in good agreement with our analysis. It is not clear whether the ≈ 10 per cent underestimate uncovered by our simulations applies to both results; the biases for symmetric fits may well be different.

We do not reproduce the finding of Salabert et al. (2007) that the $l = 0$ mode width is less sensitive to global activity than the higher-degree modes. It is not possible to check this using the MCMC fitting, as the $l = 0$ and 2 mode components are assumed to have the same width.

Our results on the asymmetry shifts are somewhat disappointing. However, we should note that the work of Jiménez-Reyes et al. (2007) was more narrowly focused on the asymmetry measurement and they carried out a more elaborate analysis to obtain the best possible results. Also, if we look only at the eight years of BiSON data that they used in the frequency range that they considered, we

do reproduce their ‘marginally’ significant result for the BiSON data. The much weaker correlation that we find for the whole data set suggests that their result for BiSON may have been a statistical fluke; measuring the asymmetry in ground-based data with a duty cycle in the 70 per cent range is still challenging. In particular, the intimate relationship between background and asymmetry may make it difficult to improve very much on these results using our current approach of pairwise fitting with a flat background for each pair. We note that the asymmetry shifts – and the absolute asymmetry values – observed for low-degree modes are too small to have a significant influence on the frequency shifts.

In general, both MLE and MCMC methods perform well and give consistent results for the solar-cycle variations of the mode parameters, as well as very similar uncertainties. This has implications for future peak-finding strategies in both solar and stellar applications. For high-quality data the considerable extra computational expense of the Bayesian approach may not be justified unless we are specifically looking for asymmetry changes. We have not addressed the issue of the advantages that might be gained by using global or pseudo-global instead of pairwise fitting for solar data. Such methods are, again, more computationally intensive than the conventional approach, and when combined with the Bayesian method this expense becomes nearly prohibitive for large data sets. Our results suggest that it might be worthwhile instead to concentrate on developing MLE-based global or pseudo-global methods for use with solar data; such methods are already in use for asteroseismology of solar-like oscillators but may require refinement to deal with longer and higher-quality solar data sets. Any such methods would need to be cross-checked against earlier results. Our findings confirm that solar-cycle variations of mode frequency, lifetime, and amplitude from earlier Sun-as-a-star work are reliable. If, as in the case of Korzennik (2013), the results obtained by novel methods differ substantially from those obtained by more conventional means, they should be treated with caution.

The data we have analysed cover the period up to the end of 2014, well into the weak but drawn out and double- or possibly multiple-humped maximum of Solar Cycle 24; we note that the RF value for 2014 is higher than that for 2013. Some of the issues around intercycle differences may become clearer once Cycle 24 definitely enters its declining phase. We await the next few years of measurements with interest.

ACKNOWLEDGEMENTS

We would like to thank all those who are, or have been, associated with BiSON. BiSON is funded by the Science and Technology Facilities Council (STFC). WJC thanks Sarbani Basu and Aldo Serenelli for providing the model BS05(OP) frequencies.

RH thanks the National Solar Observatory for computing support.

REFERENCES

- Anderson E. R., Duvall T. L., Jr, Jefferies S. M., 1990, *ApJ*, 364, 699
 Bachmann K. T., Brown T. M., 1993, *ApJ*, 411, L45
 Bahcall J. N., Serenelli A. M., Basu S., 2005, *ApJ*, 621, L85
 Basu S., Broomhall A.-M., Chaplin W. J., Elsworth Y., 2012, *ApJ*, 758, 43
 Bogdan T. J., Zweibel E. G., 1985, *ApJ*, 298, 867
 Broomhall A. M., Chaplin W. J., Elsworth Y., New R., 2009, *MNRAS*, 397, 793
 Broomhall A.-M. et al., 2011, *J. Phys. Conf. Ser.*, 271, 012025
 Campbell W. R., Roberts B., 1989, *ApJ*, 338, 538
 Chaplin W. J., Appourchaux T., 1999, *MNRAS*, 309, 761
 Chaplin W. J., Elsworth Y., Howe R., Isaak G. R., McLeod C. P., Miller B. A., New R., 1997, *MNRAS*, 287, 51
 Chaplin W. J., Elsworth Y., Isaak G. R., Miller B. A., New R., 1999, *MNRAS*, 308, 424
 Chaplin W. J., Elsworth Y., Isaak G. R., Miller B. A., New R., 2000, *MNRAS*, 313, 32
 Chaplin W. J., Elsworth Y., Isaak G. R., Marchenkov K. I., Miller B. A., New R., 2001, *MNRAS*, 322, 22
 Chaplin W. J., Elsworth Y., Isaak G. R., Miller B. A., New R., 2004, *MNRAS*, 352, 1102
 Chaplin W. J. et al., 2006, *MNRAS*, 369, 985
 Chaplin W. J., Elsworth Y., Miller B. A., Verner G. A., New R., 2007, *ApJ*, 659, 1749
 Chaplin W. J., Elsworth Y., Toutain T., 2008, *Astron. Nachr.*, 329, 440
 Christensen-Dalsgaard J., 1989, *MNRAS*, 239, 977
 Davies G. R., Broomhall A. M., Chaplin W. J., Elsworth Y., Hale S. J., 2014a, *MNRAS*, 439, 2025
 Davies G. R., Chaplin W. J., Elsworth Y., Hale S. J., 2014b, *MNRAS*, 441, 3009
 Duvall, Jr, T. L., Jefferies S. M., Harvey J. W., Osaki Y., Pomerantz M. A., 1993, *ApJ*, 410, 829
 Dziembowski W. A., Goode P. R., 2005, *ApJ*, 625, 548
 Elsworth Y., Howe R., Isaak G. R., McLeod C. P., New R., 1990, *Nature*, 345, 322
 Elsworth Y., Howe R., Isaak G. R., McLeod C. P., Miller B. A., Speake C. C., Wheeler S. J., New R., 1993, *MNRAS*, 265, 888
 Elsworth Y., Howe R., Isaak G. R., McLeod C. P., Miller B. A., New R., Speake C. C., Wheeler S. J., 1994, *ApJ*, 434, 801
 Fletcher S. T., Chaplin W. J., Elsworth Y., New R., 2009, *ApJ*, 694, 144
 Fletcher S. T., Broomhall A.-M., Salabert D., Basu S., Chaplin W. J., Elsworth Y., Garcia R. A., New R., 2010, *ApJ*, 718, L19
 Foullon C., Roberts B., 2005, *A&A*, 439, 713
 Gelman A., Carlin J. B., Stern H. S., Rubin D. B., 2003, *Bayesian Data Analysis*. Chapman and Hall, London
 Harvey J., 1985, in Rolfe E., Battrock B., eds, *ESA SP-235: Future Missions in Solar, Heliospheric and Space Plasma Physics*. ESA, Noordwijk, p. 199
 Howe R., Komm R. W., Hill F., 2002, *ApJ*, 580, 1172
 Howe R., Komm R. W., Hill F., Haber D. A., Hindman B. W., 2004, *ApJ*, 608, 562
 Howe R., Chaplin W. J., Elsworth Y., Hill F., Komm R. W., Isaak G. R., New R., 2006, *MNRAS*, 369, 933
 Jain R., Roberts B., 1994, *Sol. Phys.*, 152, 261
 Jain K., Tripathy S. C., Hill F., 2011, *ApJ*, 739, 6
 Jiménez-Reyes S. J., García R. A., Chaplin W. J., Korzennik S. G., 2004, *ApJ*, 610, L65

- Jiménez-Reyes S. J., Chaplin W. J., Elsworth Y., García R. A., Howe R., Socas-Navarro H., Toutain T., 2007, *ApJ*, 654, 1135
 Koen C., 2005, *MNRAS*, 361, 887
 Komm R. W., Howe R., Hill F., 2000a, *ApJ*, 531, 1094
 Komm R. W., Howe R., Hill F., 2000b, *ApJ*, 543, 472
 Komm R., Howe R., Hill F., 2002, *ApJ*, 572, 663
 Korzennik S. G., 2013, in Jain K., Tripathy S. C., Hill F., Leibacher J. W., Pevtsov A. A., eds, *ASP Conf. Ser. Vol. 478, Fifty Years of Seismology of the Sun and Stars*. Astron. Soc. Pac., San Francisco, p. 137
 Kuhn J. R., 1988, *ApJ*, 331, L131
 Libbrecht K. G., Woodard M. F., 1990, *Nature*, 345, 779
 Nigam R., Kosovichev A. G., 1998, *ApJ*, 505, L51
 Pallé P. L., Regulo C., Roca Cortes T., 1989, *A&A*, 224, 253
 Roberts B., Campbell W. R., 1986, *Nature*, 323, 603
 Salabert D., Chaplin W. J., Elsworth Y., New R., Verner G. A., 2007, *A&A*, 463, 1181
 Salabert D., Garcia R. A., Turck-Chieze S., 2015, *A&A*, 578, A137
 Simoniello R., Finsterle W., Salabert D., García R. A., Turck-Chieze S., Jiménez A., Roth M., 2012, *A&A*, 539, A135
 Simoniello R., Jain K., Tripathy S. C., Turck-Chieze S., Baldner C., Finsterle W., Hill F., Roth M., 2013, *ApJ*, 765, 100
 Toutain T., Elsworth Y., Chaplin W. J., 2006, *MNRAS*, 371, 1731
 Woodard M. F., Noyes R. W., 1985, *Nature*, 318, 449

APPENDIX A: THE SOLARFLAG COMPLEX FREQUENCY AMPLITUDE, AND FREQUENCY POWER, SPECTRUM

In the following description – which is based on the detailed discussions in Toutain et al. (2006) and Chaplin et al. (2008) – we model the p modes as forced, damped oscillators having a high Q . The frequency response of a given mode (with n , l , and m) is then just a Lorentzian, which may be written in complex amplitude form as

$$L_{nlm}(v) = \Re[A_{nlm}(v)] + i\Im[A_{nlm}(v)], \quad (\text{A1})$$

where

$$\Re[A_{nlm}(v)] = \frac{\xi_{nlm} H_{nlm}^{1/2}}{1 + \xi_{nlm}^2}$$

and

$$\Im[A_{nlm}(v)] = \frac{H_{nlm}^{1/2}}{1 + \xi_{nlm}^2} \quad (\text{A2})$$

are the real and imaginary parts, respectively. In the above, H_{nlm} is the height of the peak and

$$x_{nlm} = (v - v_{nlm})/(\Gamma_{nlm}/2),$$

where v_{nlm} and Γ_{nlm} are the central frequency and line width of the mode, respectively.

The observed response in the frequency domain also depends on the spectral response of the excitation function, which we call $E_{nlm}(v)$. The complex amplitude, $A_{nlm}(v)$, of the full response is therefore actually

$$A_{nlm}(v) = L_{nlm}(v)E_{nlm}(v). \quad (\text{A3})$$

We excite the modes with the granulation-like noise, which has a response that is white only locally in the vicinity of the resonance. The power spectral density of this granulation-like noise follows the Harvey (1985) power-law model given in equation (2). To obtain the excitation response, $E_{nlm}(v)$, we must take the square root of the power spectral density of the granulation-like noise; and we

must also normalize to the value of the granulation-like noise at the frequency of the resonance. This gives

$$E_{nlm}(\nu) = \left(\frac{1 + (2\pi\nu_{nlm}\tau)^2}{1 + (2\pi\nu\tau)^2} \right)^{1/2}. \quad (\text{A4})$$

The power spectral density of this mode is then:

$$\text{PSD}_{nlm} = |A_{nlm}(\nu)|^2 = [L_{nlm}(\nu)E_{nlm}(\nu)][L_{nlm}^*(\nu)E_{nlm}^*(\nu)]. \quad (\text{A5})$$

Now let us write down equations to describe the complete solarFLAG spectrum. This spectrum is comprised of many overtones, whose excitation is correlated, and also correlated and uncorrelated background noise. We begin with a description of the power spectral density of the overtones of a given (l, m) . The excitation of the overtones is correlated, with the coefficient ρ describing this correlation. The modes are also correlated with the granulation-noise background, and this correlation is also set equal to ρ .

First, consider the equations which describe the real and imaginary parts of the complex amplitude given by the overtones. We must sum over all the radial orders n of the chosen (l, m) . We therefore have

$$\Re[A_{lm}(\nu)] = \sum_{\text{overn}} \left(\frac{H_{nlm}E_{nlm}\xi_{nlm}}{1 + \xi_{nlm}^2} \right)^{1/2}, \quad (\text{A6})$$

and

$$\Im[A_{lm}(\nu)] = \sum_{\text{overn}} \left(\frac{H_{nlm}E_{nlm}}{1 + \xi_{nlm}^2} \right)^{1/2}. \quad (\text{A7})$$

Note that to keep things tidy, we have dropped the explicit dependence of E_{nlm} and ξ_{nlm} on ν in the equations above, and in what

follows. The observed power spectral density of the overtones *and* the granulation-like noise for this (l, m) is then given by

$$\begin{aligned} \text{PSD}_{lm}(\nu) = & |\rho| \psi_{lm}(\nu) \psi_{lm}^*(\nu) \\ & + (1 - |\rho|) \left(\sum_{\text{overn}} \left(\frac{H_{nlm}E_{nlm}}{1 + \xi_{nlm}^2} \right) + n_{lm} \right), \end{aligned} \quad (\text{A8})$$

where

$$\psi_{lm}(\nu) = \Re[A_{lm}(\nu)] + i \Im[A_{lm}(\nu)] - n_{lm}^{1/2}. \quad (\text{A9})$$

There are two important things to notice about equations (A8) and (A9). First the description in equation (A8) has two parts: one part, prefixed by the factor $|\rho|$, describes the correlated contributions to the spectrum; while the other part, prefixed by the factor $(1 - |\rho|)$, describes the uncorrelated contributions. Secondly, notice how the background factor $n_{lm}^{1/2}$ in equation (A9) is prefixed by a minus sign. This makes the correlation of the modes and background *negative*, which in turn means we get negative asymmetry as in real Doppler velocity observations.

Finally, the total power spectrum density is given by summing, incoherently, the power spectral densities of all (l, m) in the spectrum, and the uncorrelated photon-shot noise background, n_{psn} (see Section 2.2.5), to give

$$\text{PSD}(\nu) = \sum_{\text{all}(l,m)} \text{PSD}_{lm}(\nu) + n_{\text{psn}}. \quad (\text{A10})$$

This paper has been typeset from a $\text{\LaTeX}$ file prepared by the author.