The Astrophysical Journal, 659:1749–1760, 2007 April 20

SOLAR *p*-MODE FREQUENCIES OVER THREE SOLAR CYCLES

W. J. Chaplin, Y. Elsworth, B. A. Miller, and G. A. Verner
School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK;
wjc@bison.ph.bham.ac.uk, ype@bison.ph.bham.ac.uk, miller@bison.ph.bham.ac.uk, gav@bison.ph.bham.ac.uk

and

R. New
Faculty of Arts, Computing, Engineering and Sciences, Sheffield Hallam University, Sheffield S1 1WB, UK; r.new@shu.ac.uk
Received 2006 November 15; accepted 2007 January 8

ABSTRACT

We analyze thirty years of solar oscillations data collected by the Birmingham Solar Oscillations Network (BiSON). Estimates of the mean frequency shifts of low-degree *p*-modes have been extracted over a period spanning solar cycles 21–23. Two methods of analysis are used to extract the frequency shifts: one method uses results on fitted frequencies of individual modes, which are then averaged to give mean frequency shifts; the other method uses cross-correlations of power frequency spectra made from subsets of the data shifted in time. The frequency shifts are correlated against six proxies of solar activity, which are sensitive to magnetic and irradiance variability at a range of locations from the photosphere to the corona. We find proxies that have good sensitivity to the effects of weak-component magnetic flux—which is more widely distributed in latitude than the strong flux in the active regions—are those that follow the frequency shifts most consistently over the three cycles. This list includes the Mg II H and K core-to-wing data, the 10.7 cm radio flux, and the He I equivalent width data. While the two methods of analysis give consistent results, use of the cross-correlation function to measure mean frequency shifts returns less precise values in cases in which the duty cycle is greater than 30%. Estimation of uncertainties from the cross-correlation method also requires that proper allowance be made for strong correlations in the data.

Subject headings: Sun: activity — Sun: helioseismology — Sun: magnetic fields

1. INTRODUCTION

The frequencies of solar *p*-mode oscillations provide a unique insight into the interior structure of the Sun. These frequencies are known to vary with solar activity (Woodard & Noyes 1985; Elsworth et al. 1990, 1994; Chaplin et al. 1998, 2001, 2004) and may be used to infer the structural changes occurring within the Sun. Each of the 11 yr Schwabe activity cycles has a distinct amplitude, duration, and shape, when seen in a number of solar activity proxy indicators that respond to different physical processes. Solar cycle 23, which reached maximum during 2000 and 2001, has been described as a magnetically simpler cycle when compared with recent cycles (de Toma et al. 2004). The reduction in strong magnetic flux over cycle 23 is seen in several activity proxies with the notable exception of the total solar irradiance (TSI). Despite the largely reduced magnetic activity, the TSI has similar minimum-to-maximum ranges over cycle 22 and cycle 23, reflecting the complex balance between dark sunspots and bright faculae or plages over the two cycles.

The Birmingham Solar Oscillations Network (BiSON) has been in operation, in some capacity, since mid 1976, with steadily increasing coverage until the full six-station network was implemented in late 1992. The BiSON data archive now contains solar oscillations data covering a 30 yr epoch, which includes most of cycles 21–23. In this paper, we analyze this extensive data set to look for correlations and consistencies between low-degree ($l \leq 3$) *p*-mode oscillation frequencies and proxies of solar activity across each of the three solar cycles.

2. DATA

The BiSON instruments use a resonant scattering spectrometer to detect the line shift in the solar 769.9 nm K D1 absorption line (Brookes et al. 1978). These line shifts were calibrated into line-of-site velocities and residual solar oscillation velocities extracted according to Elsworth et al. (1995). An optimized concatenated time series was built based on the criterion of maximum signal-to-noise ratio where data are available from more than one site. This long time series was then divided into 80 independent short series. After the ninth such series, all remaining time series are contiguous and have a duration of 108 days. Early data from the BiSON were collected from just one or two sites, and sometimes only over the summer months. The first nine series are of different durations reflecting the observing runs available at the time. The durations and duty cycles of these early series are shown in Table 1. The duty cycles of all remaining 108 day sets are illustrated in Figure 1.

To obtain estimates of the solar variability, we have used proxies that are sensitive to global magnetic activity and irradiance variability over different locations from the photosphere to the corona. The total solar irradiance composite (TSI), combined from a number of satellite programs by Fröhlich (2000), is affected by both photospheric and chromospheric changes. The NOAA Mg II core-to-wing ratio (Viereck et al. 2001) is derived from space-based UV spectral irradiance measurements and is sensitive to the entire solar disk. The Mg II H and K line cores originate in the chromosphere, and the ratio of the core size to that of the more stable background gives a robust indication of chromospheric activity. The NSO Kitt Peak Magnetic Index (KPMI) was derived from full-disk magnetograms as the mean line-of-sight absolute magnetic field strength. This magnetic index is available covering much of the same period as the BiSON observations but was terminated in late 2003. Also discontinued at this time were the NSO He I 1083 nm line equivalent width observations. The He I 1083 nm multiplet is affected strongly by photoionizing radiation from the upper transition region and corona (Andretta & Jones

TABLE 1
Duty Cycles and Durations of the First Nine Time Series Used in the Analysis

Spectrum	Date	Length (days)	Fill (%)	Site
1.....	1976.56	48.1	7.6	IZ
2.....	1977.53	43.0	13.2	IZ
3.....	1978.58	40.3	15.5	IZ
4.....	1980.55	30.4	20.7	IZ
5.....	1981.41	91.7	48.2	IZ/HA
6.....	1982.29	142.4	37.4	IZ/HA
7.....	1983.35	115.3	43.7	IZ/HA
8.....	1984.30	138.4	47.3	IZ/HA
9.....	1984.84	59.3	18.9	IZ

Note.—The sites codes used are: IZ, Izaña, Tenerife; HA, Haleakala, Hawaii.

1997). The Ottawa/Penticton 10.7 cm radio flux contains a contribution from radio plages in the upper chromosphere and low corona above developing active regions, in addition to a quiet-Sun background contribution from the entire solar disk. The NOAA International Sunspot Number (ISN) provides a measurement of magnetic activity through the number of active sunspot groups and individual spots on the visible solar disk.

In this study, we have not used spatially decomposed proxies of solar activity. Previous work has shown that the Legendre polynomial decomposition of the KPMI is correlated strongly with shifts in the corresponding even a-coefficients obtained from intermediate- and high-degree modes during cycle 23 (Antia et al. 2001; Howe et al. 2002). Low-degree mode frequency shifts have also been shown to match their respective decomposed magnetic activity components (Chaplin et al. 2004). However, the intention here is to look at the behavior of helioseismic frequency shifts with the disk-averaged global proxies over three solar cycles.

3. CALCULATING THE FREQUENCY SHIFTS

In order to determine how the p-mode frequencies are affected by the solar cycle, we compared the frequencies from each short time series with those obtained during a minimum activity reference epoch. This epoch was taken to be the minimum of cycle 23, around 1996, for which high-fill BiSON observations are available from the complete network. Thus, all frequency shifts were calculated relative to the minimum of cycle 23. Two methods were used to extract the frequency shifts: frequency integrated frequency shifts, using the cross-correlation of the power spectrum of each time series with that of an equal length time series centered on the minimum activity epoch; and the shifts of the mode frequencies determined by fitting asymmetric Lorentzian mode profiles to the resonant peaks in the power spectra.

3.1. *Cross-Correlation Fitting*

Use of cross-correlations of power spectra to estimate the mean shifts in frequency has previously proved useful where data have a low duty cycle and individual mode fitting is difficult (Régulo et al. 1994). This method can use all frequency data in the range over which the cross-correlation is taken, which for mean-shift studies is usually restricted to 2500–3700 μHz. Modes below 2500 μHz show very little change in frequency with increasing activity, while modes above 3700 μHz are expected to experience a decrease in frequency shift. The frequency data used in the cross-correlation includes the information from those modes which are poorly resolved and would be discarded in an individual mode-fitting analysis. The interpretation of the frequency shift determined by the cross-correlation over a large frequency range must be made carefully. The 2500 $\mu\text{Hz} \leq \nu \leq 3700$ μHz range in frequency encompasses 34 modes of degree $l \leq 3$, the highest frequency modes experiencing a shift of up to 4 times the magnitude of the lowest frequency modes. In addition, the cross-correlation is influenced more by modes with the largest amplitudes in the power spectrum, i.e., the $l \leq 2$ modes around 3100 μHz. Since the frequencies of the largest amplitude modes lie in the center of the range used to calculate the cross-correlation, we can expect the cross-correlation determined shift to be an estimate of the average frequency shift of modes at $\sim$3100 μHz.

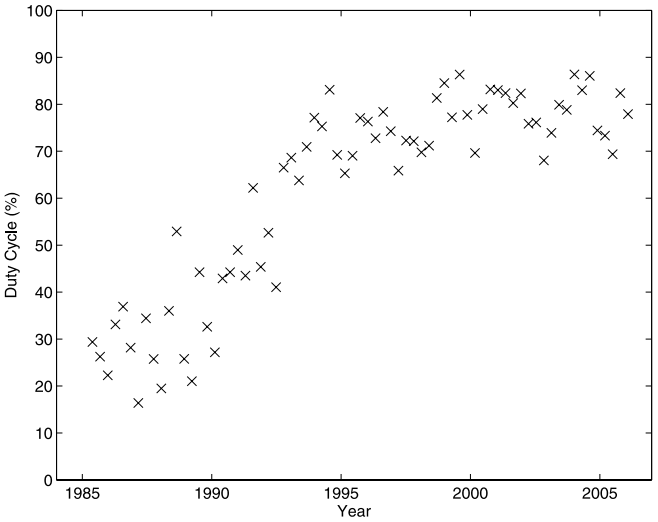

Fig. 1.—Duty cycles of the 71 108 day contiguous time series.

To find the position of the main peak in the cross-correlation, a method has previously been devised to fit a second-order polynomial to the logarithm of the cross-correlation function over a range of ± 5 μHz around zero lag (Régulo et al. 1994; Jiménez-Reyes et al. 1998). This was extended by Jiménez-Reyes et al. (2001) by modeling the cross-correlation peak as a symmetric Lorentzian, following from the assumption that the mode peaks themselves are Lorentzian, and fitting the central peak and diurnal sidebands (± 11.57 μHz) over a range of ± 20 μHz. This functional form for the cross-correlation function is adequate for low duty cycle data, but as the fill increases the diurnal sidebands become less significant and further structure appears due to the close spacing in frequency of $l = 0, 2$ and $l = 1, 3$ modes. For high-fill spectra, cross-correlation peaks appear at those lag frequencies corresponding to $l = 0$ modes overlapping their neighboring $l = 2$ modes, at the average $d_{0,2}(\nu)$ small-separation frequency, and where $l = 1$ modes overlap with the adjacent $l = 3$ modes, at the average $d_{1,3}(\nu)$ small-separation frequency. This is illustrated in Figure 2.

The model used here to fit the cross-correlation function is a superposition of seven symmetric Lorentzian functions representing the central peak, diurnal sidebands, and $l = 0, 2$ and $l = 1, 3$ overlapping pairs. The widths of the central peak and those of the diurnal sidebands are assumed to be the same; however, the widths of the peaks in the cross-correlation caused by overlapping $l = 0, 2$ and $l = 1, 3$ modes are increased due to the difference in the number of resolved components within the individual degree multiplets. For instance, the $l = 2$ modes have three visible, rotationally split components, whereas the $l = 0$ modes are represented by a single peak. The higher effective total multiplet width of the $l = 2$ multiplet, in comparison with the single $l = 0$ width, increases the width of the peak in the cross-correlation function, at

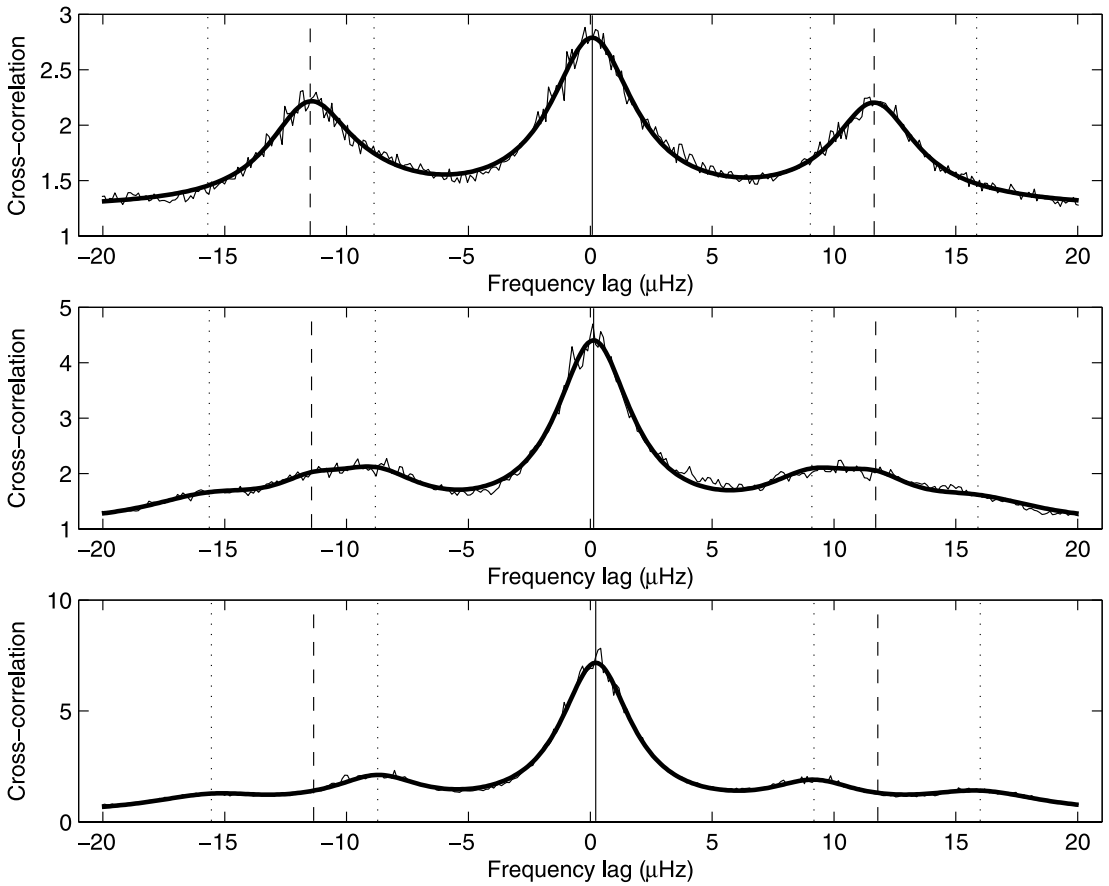

FIG. 2.—Comparison of cross-correlation fitting functions for 108 day spectra. Dotted lines show positions of shifted small-separation peaks, dashed lines show positions of shifted diurnal sideband peaks, and the solid vertical line indicates the shift in the central peak. *Top*, Duty cycle of 16%, fitted shift of 0.08 μHz; *middle*, duty cycle of 26%, fitted shift of 0.14 μHz; *bottom*, duty cycle of 86%, fitted shift of 0.23 μHz.

the frequency lag where these modes overlap, relative to the cross-correlation peak width that would be observed if the overlapping modes had the same number of visible rotationally split components.

The separation of the main peak and the diurnal sidebands is fixed at 11.57 μHz. The separations of the main peak and the small-separation overlapping peaks are determined by the fitting. The initial guess values of the frequency lags of these small-separation peaks are determined from the mean separation between the respective mode pairs across the frequency range used to calculate the cross-correlation. This equates to a spacing of 8.95 μHz for $l = 0, 2$ pairs and 15.77 μHz for $l = 1, 3$ pairs, for a cross-correlation taken over the range of 2500 $\mu\text{Hz} \leq \nu \leq$ 3700 μHz. The total range in cross-correlation frequency lag over which the fitting is performed is ± 20 μHz.

The functional form used for the cross-correlation frequency lag, $\tilde{\nu}$, can be written as

$$M(\tilde{\nu}) = \sum_{k=-3}^{3} \frac{\beta_k A \left[\left(\Gamma + \delta\Gamma_{|k|}\right)/2\right]^2}{\left[\tilde{\nu} - \Delta\tilde{\nu} + (k/|k|)d_{|k|}\right]^2 + \left[\left(\Gamma + \delta\Gamma_{|k|}\right)/2\right]^2} + B, \quad (1)$$

where A is the central peak amplitude, β_k is the relative height of each adjacent cross-correlation peak to that of the central peak, Γ is the width of the central cross-correlation peak, $\delta\Gamma_{|k|}$ is the increase in width of the $l = 0, 2$ and $l = 1, 3$ overlapping peaks ($\delta\Gamma = 0$ for $|k| = 0$ and 2), $\Delta\tilde{\nu}$ is the frequency shift between the two spectra, $d_{|k|}$ is the frequency spacing to the adjacent peaks ($d = 11.57$ μHz for $|k| = 2$ [diurnal sidebands]), where d is fitted as the mean small-separation frequencies for $l = 0, 2$ ($|k| = 1$) and $l = 1, 3$ ($|k| = 3$), and B is the background offset.

Although the statistics of the power spectra are not Gaussian, each point in the cross-correlation is a sum of many points in each power spectrum, so the central limit theorem applies. Thus, the fitting of equation (1) to the cross-correlation data was approached in a standard least-squares sense. It must be stressed that the points in the cross-correlation data are not independent and are, in fact, very highly correlated. This should not introduce any bias into the fitted parameters, but the formal uncertainties, calculated from the Hessian matrix of the fit, are underestimated; i.e., the precision is overestimated and must be scaled accordingly.

To estimate the true uncertainty in each of the fitted parameters, a series of fits were performed using Fitting at Low Angular Degree (solarFLAG) artificial data (Chaplin et al. 2006b). Fifty solarFLAG simulated BiSON-like time series of duration 3456 days were generated, each with the same imposed activity cycle but different stochastic noise. The mode frequencies used in the solarFLAG time series were modified continuously in time in accordance with the imposed activity perturbation. Each solarFLAG series was then modulated by a BiSON-like window function corresponding to a high-fill regime, typically 70%–80% duty cycle. The linear correlation between the mean activity level for each solarFLAG subset spectrum and the mean of the corresponding frequency shifts, calculated using the cross-correlation fitting, gives an indication of any bias introduced by this method. If the method is free from bias then the mean value of the frequency shifts should be very strongly correlated with the mean activity in

TABLE 2
CORRELATION COEFFICIENTS

Duration (days)	No. of Sets	Correlation Coefficient
54	64	0.9950
108	32	0.9980
162	21	0.9985
216	16	0.9993
270	12	0.9995
324	10	0.9996
378	9	0.9998
432	8	0.9996

NOTE.—Shown are the correlation coefficients for the linear relationship between the mean value of frequency shift over all 50 solarFLAG series and the mean imposed activities.

each set. These correlation values are given in Table 2 for simulated spectrum durations from 54 to 432 days.

Table 2 shows that the mean frequency shift for each activity level is strongly correlated with the mean activity over that period. Each individual solarFLAG series displays considerable scatter from the linear relation. The cross-correlation frequency shifts from the solarFLAG time series, split into 108 day short spectra, are shown in Figure 3 with the imposed activity signature scaled according to a linear fit to the mean values. The scatter of points in the sets at minimum activity is noticeably reduced because the cross-correlation function in this case contains a large spike at zero lag due to the high correlation between the minimum reference set and the set with which it overlaps.

An estimate of the true uncertainty in the frequency shift using the cross-correlation method can be obtained from the standard deviation of the fitted solarFLAG frequency shifts for each activity level. The mean of the standard deviations for each duration of simulated time series, excluding the set covering the epoch of minimum activity, can be taken as the uncertainty estimate. The scaling factor which must be applied to the formal error returned assuming uncorrelated data, i.e., as returned by the Hessian of the fit, can then be determined from the mean value of the returned

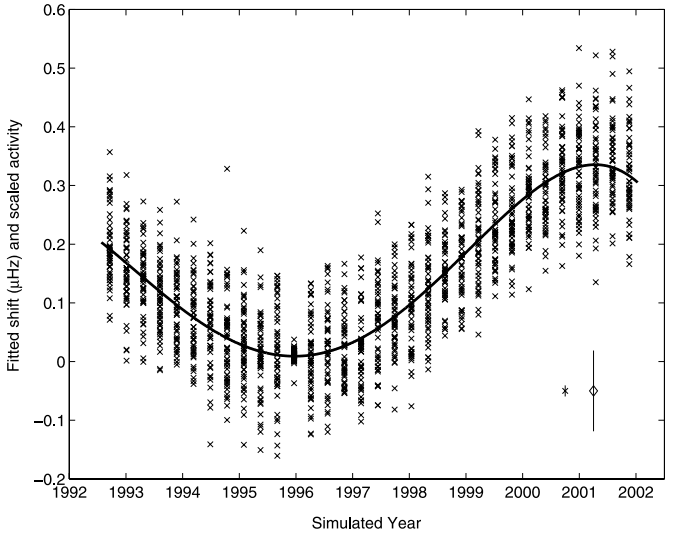

FIG. 3.—Frequency shifts from the cross-correlation fitting of solarFLAG data split into 108 day contiguous high-fill sets. The solid line shows the imposed activity proxy scaled according to a linear fit to the mean frequency shifts. Error bars shown are indicative of the formal errors from the least-squares fit (*cross*) and the standard deviation of the simulated results (*diamond*).

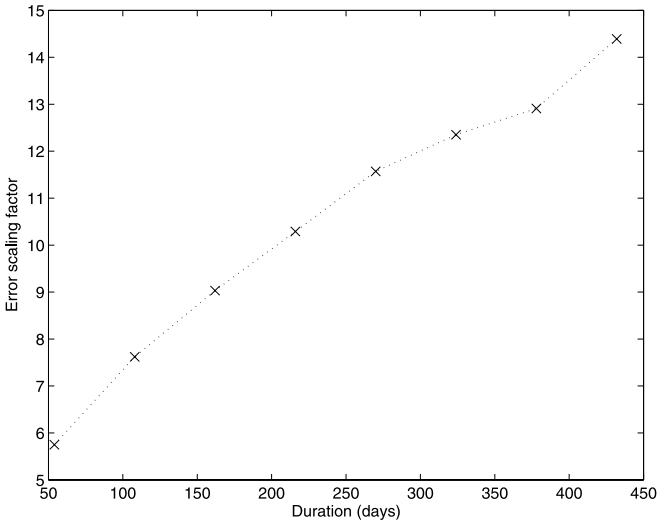

FIG. 4.—Scaling factor necessary to scale the underestimated formal error in mean frequency shift to the correct precision.

formal errors from the least-squares fit of equation (1) to the cross-correlation function. The scaling factors for each duration are shown in Figure 4. It should be noted that these scaling factors are all significantly greater than unity.

3.2. *Mode-Pair Fitting*

For time series of long enough duration to provide sufficient spectral resolution, pairs of modes in the power spectrum ($l = 0, 2$ and $l = 1, 3$), along with their diurnal sidebands, can be fitted to asymmetric Lorentzian profiles (Nigam & Kosovichev 1998). As the fitting is performed in the frequency domain, the negative exponential statistics of the power spectrum must be taken into account, and a more generalized maximum likelihood minimization method is used to fit the mode profiles. The pair-by-pair fitting procedure used here is detailed comprehensively in Chaplin et al. (1999 and references therein). The mode fitting provides frequencies, with formal errors, for modes of degree $l \leq 3$. The number of modes which are readily fitted, and the uncertainties associated with the modes, depends on the duration of the spectrum, the noise levels in the data, and the duty cycle of the observations.

Once the mode frequencies had been determined, a minimum activity reference frequency set was calculated using the average frequencies of modes during the minimum activity epoch. The differences between this minimum reference and the corresponding modes in each spectrum provide the residual frequency shifts from minimum activity. The frequency shift has a well-known frequency and degree dependence that must be removed before an average shift can be found over a large range of frequencies. The method used here was the same as that in Chaplin et al. (2004). The degree dependence, due to differences in mode inertia, is small for $l \leq 3$ but was removed, nonetheless, by scaling each mode residual to the radial ($l = 0$) equivalent using the mode inertia ratio, $Q_{n,l}$, of Christensen-Dalsgaard & Berthomieu (1991). The remaining frequency dependence was removed by scaling to a fiducial frequency of 3000 μHz using a fifth-order polynomial approximation to the inertia-scaled frequency shifts of intermediate- and high-degree modes, using data from the Global Oscillations Network Group (GONG). This was the same frequency scaling function, $F(\nu)$, of Chaplin et al. (2004).

The mean frequency shifts for each spectrum, $\langle \delta\nu(t) \rangle$, were calculated from the shift of each given mode, $\delta\nu_{n,l}(t)$, and the

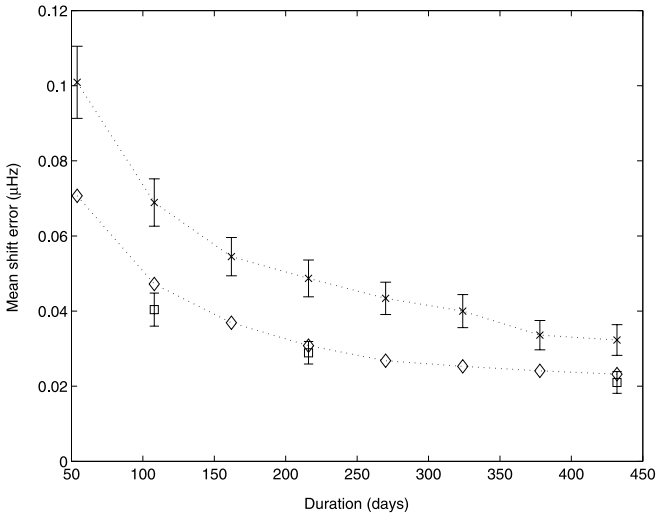

FIG. 5.—Comparison of the uncertainties in the mean frequency shifts in high-fill data: *crosses*, cross-correlation fitted shifts; *diamonds*, theoretical mode-pair fitted shift estimate; *squares*, mean-shift uncertainty from high-fill BiSON data. Error bars signify the 1 σ spread in uncertainties across the data sets used.

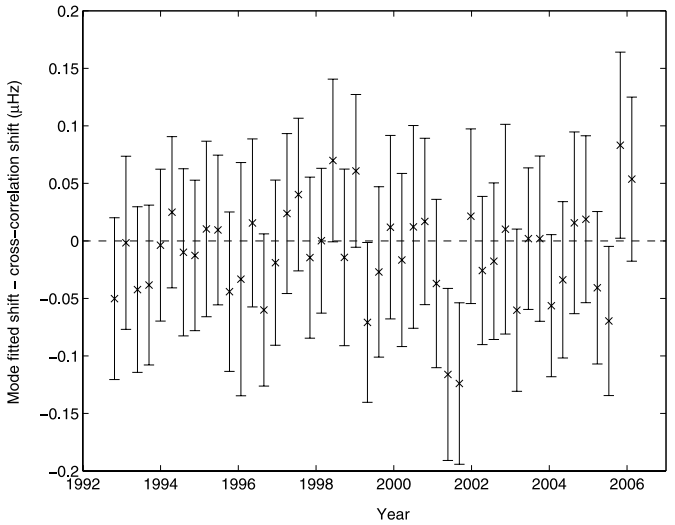

FIG. 6.—Differences between mode-fitted shifts and cross-correlation method (in the sense of mode-fitted minus cross-correlation) for 108 day high-fill BiSON data.

associated uncertainty, $\sigma_{n,l}(t)$, using the weighted mean of the scaled values:

$$\langle \delta\nu(t) \rangle = \left[\sum_{\nu} \frac{\delta\nu_{n,l}(t) F(\nu)}{\sigma^2_{n,l}(t) Q_{n,l}} \right] \left[\sum_{\nu} \frac{F^2(\nu)}{\sigma^2_{n,l}(t) Q^2_{n,l}} \right]^{-1}. \quad (2)$$

The uncertainty on this mean frequency shift is given by

$$\langle \sigma(t) \rangle = \left[\sum_{\nu} \frac{F^2(\nu)}{\sigma^2_{n,l}(t) Q^2_{n,l}} \right]^{-1/2}. \quad (3)$$

To allow for comparison of the two methods used here to determine the frequency shifts, the mean fitted shifts were calculated using the 34 modes of degree $l \leq 3$ for which 2500 μHz $\leq \nu \leq$ 3700 μHz.

The expected uncertainty on the mean mode-fitted frequency shift is dependent on the number of modes fitted, the observational error on the mode frequencies, and the number of spectra used to generate the frequencies in the minimum activity reference set. The duty cycle has a only a small effect on the determined frequency of a given mode, but it does have an impact on the number of modes which return good, robust fits to the asymmetric Lorentzian model. The uncertainty in mode frequencies scales with the duration of observation as $T^{-1/2}$ (Chaplin et al. 2002). The observational error increases with increasing frequency due to the line width of the modes, but it changes only slightly for 2500 μHz $\leq \nu \leq$ 3700 μHz. Over this range it was found empirically that the formal error on each mode frequency, $\sigma_\nu(T)$, depends only on the spectrum length, T (in days), as

$$\sigma_\nu(T) \simeq \frac{2.103}{T^{1/2}} \ \mu\text{Hz}. \quad (4)$$

An estimate of the uncertainty on the mean shift can be obtained by propagating the errors on observational mode frequencies. For frequencies that are *not* used to calculate the mean frequencies during the minimum activity reference epoch, the uncertainty is given by

$$\langle \tilde{\sigma}(T) \rangle \simeq \frac{\sigma_\nu(T)}{M^{1/2}} \left(1 + \frac{1}{N_{\min}} \right)^{1/2}, \quad (5)$$

where M is the mean number of fitted modes in each spectrum and $N_{\min}$ is the number of spectra averaged over the minimum activity epoch to make the reference frequencies. For frequencies that are used in the calculation of the mean minimum activity frequency set, the formula becomes

$$\langle \tilde{\sigma}(T) \rangle \simeq \frac{\sigma_\nu(T)}{M^{1/2}} \left(1 - \frac{1}{N_{\min}} \right)^{1/2}. \quad (6)$$

The change in the formula arises because the mean frequencies are then correlated with the individual frequencies (see W. J. Chaplin et al. 2007, in preparation).

3.3. *Comparison of the Techniques*

The mean mode-fitted frequency shift has been scaled to indicate the shift of a radial ($l = 0$) mode at 3000 μHz. The cross-correlation fitting method is dominated by the large-amplitude $l \leq 2$ modes around 3100 μHz. The difference in sensitivity between the two methods may produce slightly different values for the frequency shift, but these should be strongly correlated.

The precision of the mean-shift estimate is affected by the duration of the spectra used and the fill in the data. It is useful to consider the two cases of high- and low-fill regimes separately. The uncertainty using the cross-correlation method can be estimated from the solarFLAG simulations outlined in § 3.1. The mode-fitted mean-shift uncertainty is estimated using equation (5), with the exact uncertainties given by equation (3) for those power spectra that have been mode-pair fitted. A comparison of the uncertainties in high-fill data (∼80% duty cycle) is shown in Figure 5 for the two methods. Equation (5) slightly overestimates the true uncertainty in the mode-fitting analysis, particularly for sets of low duration. This is a result of outlier rejection criteria and the removal of individual mode frequency shifts that have an unreasonable level of uncertainty, which propagates back to removing those modes with high observational errors.

Figure 5 shows that in the case of high-fill data, the cross-correlation fitting method returns values of the mean frequency shift that are less precise than those produced by direct mode fitting. A comparison of the mode-fitted shifts and cross-correlation

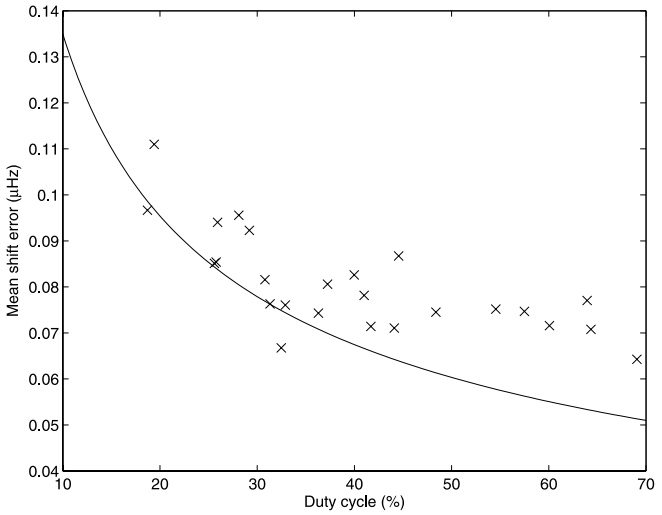

FIG. 7.—Comparison of the uncertainties in the mean frequency shifts for data of reduced fill in 108 day spectra: *solid line*, expected uncertainty in mode-fitted mean frequency shifts; *crosses*, 1 σ spread in fitted cross-correlation shifts from solarFLAG data simulations.

fitted shifts for high-fill 108 day BiSON data is shown in Figure 6. The shift values are generally compatible with each other once the errors in the cross-correlation have been scaled. There is a small variation between the methods at high activity ($\sim$2002) when the cross-correlation gives a slightly higher mean shift than in the fitted case. This may be due to the slightly higher frequencies that the cross-correlation is sensitive to, compared to the scaled 3000 μHz sensitivity of the mode-fitted mean shifts.

As the duty cycle becomes lower for older, nonnetwork BiSON data, the differences in precision between the two techniques become less significant. While the precision of the mode-fitted mean frequency shifts is affected strongly by the smaller number of modes available in low-fill data, the cross-correlation uncertainties, calculated from the solarFLAG simulations, remain similar down to a duty cycle of $\sim$40%. Figure 7 shows the comparison between the expected mode-fitted uncertainties and the cross-correlation uncertainties (from simulations) for 108 day data of reduced fill.

The simulations have shown that data of fill greater than 40%–50% should be mode-fitted to calculate the mean frequency shifts. Data of lower fill may be problematic to fit reliably, and the reduction in the number of well-fitted modes reduces the precision to a level close to that achieved by fitting the cross-correlation function. In these cases both methods were attempted and the optimum method was determined by the precision of the results.

4. RESULTS ON THE BiSON FREQUENCY SHIFTS

The methods outlined in the previous section were used to calculate mean frequency shifts from the extended BiSON observations. These mean shifts were then correlated against each of the six activity proxies outlined in § 2, using data over each cycle separately and also using all data over the entire 30 yr period.

Table 3 shows the values attained by each of the activity proxies and the frequency shifts over the maximum and minimum of each cycle. The minimum of cycle 21 was not included due to the low fill in the single data set available during that epoch. The maximum of cycle 21 was represented solely by spectrum 5, which has a relatively high fill of 48.2% over 91.7 days. This long observing run in 1981 was made from two sites, well separated in longitude (Hawaii and Tenerife), specifically to observe the Sun over maximum activity. The frequency shifts and activities during the maxima and minima of cycles 22 and 23 were averaged over eight 108 day spectra. The activities at the maximum of cycle 21 were calculated using an 864 day mean of each activity proxy centered on the same epoch (1980.8) to remain consistent with the averaged activities over the following cycles. The final column in Table 3 shows the ratio of amplitudes of cycle 22 to cycle 23, in the sense $(\max - \min)_{22}/(\max - \min)_{23}$. This gives an indication of the different response to the magnetically weaker cycle 23 in each of the activity proxies. Proxies that are sensitive to effects at the photosphere or in the lower atmosphere, and are dominated by changes in strong flux, are affected more by the disparity in magnetism between the two cycles. The chromospheric and coronal activity proxies show a lower amplitude change between the two cycles. The amplitudes of the frequency shifts over each maximum are greater in the cross-correlation fitted analysis than in the mode-fitted case. This is to be expected if the cross-correlation is an indication of the average frequency shift at a frequency higher than 3000 μHz. However, the ratio of amplitudes over the two cycles is in very good agreement between the two methods.

Table 4 shows the shift per unit activity gradients for each of the proxies over successive solar cycles with the corresponding linear correlation coefficients. The relative change in these gradients over the three cycles is plotted in Figure 8 for the mode-pair fitted results, clearly showing the inconsistency of the KPMI and TSI proxies when used to describe the frequency shifts. The mode-pair fitted and cross-correlation fitted frequency shifts correlated against each of the activity proxies are shown in Figures 9–12. The linear regressions in Figures 10 and 12, and the

TABLE 3
VALUES OF ACTIVITY PROXIES AND FREQUENCY SHIFTS

	Cycle 21	Cycle 22		Cycle 23		
Parameter	Maximum	Minimum	Maximum	Minimum	Maximum	$(\mathrm{Max} - \mathrm{Min})_{22}/(\mathrm{Max} - \mathrm{Min})_{23}$
KPMI (G)	18.41	7.10	21.65	7.07	17.39	1.410
ISN	153.49	16.49	150.91	11.23	113.09	1.320
$F_{10.7}$	202.1	74.7	204.0	73.6	182.5	1.187
TSI (W m^{-2})	1366.456	1365.574	1366.438	1365.553	1366.348	1.087
Mg II	0.2812	0.2645	0.2803	0.2647	0.2795	1.066
He I EW (10^{-13} m)	76.10	46.77	76.47	47.37	77.58	0.983
Mode-fitted shifts (μHz)	0.321 $\pm$ 0.047	$-$0.017 $\pm$ 0.014	0.352 $\pm$ 0.016	0.001 $\pm$ 0.011	0.325 $\pm$ 0.013	1.137 $\pm$ 0.089
Cross-corr. shifts (μHz)	0.363 $\pm$ 0.090	0.005 $\pm$ 0.036	0.416 $\pm$ 0.032	0.017 $\pm$ 0.022	0.380 $\pm$ 0.021	1.131 $\pm$ 0.162

NOTES.—Shown are the values of the activity proxies and frequency shifts over the maxima and minima of the solar cycles and the ratios comparing the minimum-to-maximum range of cycles 22 and 23. The 10.7 cm radio flux has units of 10^{-22} J s^{-1} m^{-2} Hz^{-1}; the ISN and Mg II core-to-wing ratio are dimensionless.

TABLE 4

Gradients and Linear Correlation Coefficients of Frequency Shifts as a Function of Each Activity Proxy

Parameter	Cycle 21		Cycle 22		Cycle 23		All Data	
	Gradient	Corr.	Gradient	Corr.	Gradient	Corr.	Gradient	Corr.
Mode Fitted								
KPMI	0.0330 ± 0.0041	0.949	0.0221 ± 0.0020	0.899	0.0315 ± 0.0026	0.942	0.0258 ± 0.0015	0.903
ISN	0.00311 ± 0.00045	0.935	0.00258 ± 0.00017	0.948	0.00283 ± 0.00031	0.875	0.00271 ± 0.00015	0.913
$F_{10.7}$	0.00327 ± 0.00040	0.951	0.00268 ± 0.00018	0.939	0.00280 ± 0.00024	0.917	0.00279 ± 0.00013	0.930
TSI	0.340 ± 0.086	0.851	0.376 ± 0.025	0.941	0.254 ± 0.040	0.773	0.326 ± 0.022	0.855
Mg II	24.3 ± 3.5	0.952	22.5 ± 1.3	0.950	22.9 ± 1.6	0.944	22.3 ± 0.9	0.943
He I EW	0.0126 ± 0.0017	0.933	0.0122 ± 0.0007	0.951	0.0108 ± 0.0008	0.956	0.0111 ± 0.0005	0.942
Cross-Correlation Fitted								
KPMI	0.0326 ± 0.0047	0.933	0.0242 ± 0.0026	0.862	0.0349 ± 0.0033	0.923	0.0294 ± 0.0019	0.878
ISN	0.00324 ± 0.00067	0.885	0.00281 ± 0.00027	0.913	0.00332 ± 0.00037	0.882	0.00304 ± 0.00019	0.884
$F_{10.7}$	0.00320 ± 0.00058	0.904	0.00290 ± 0.00027	0.898	0.00315 ± 0.00028	0.909	0.00305 ± 0.00018	0.900
TSI	0.398 ± 0.107	0.835	0.447 ± 0.037	0.911	0.284 ± 0.050	0.744	0.349 ± 0.031	0.806
Mg II	24.4 ± 5.3	0.892	25.2 ± 2.1	0.919	25.7 ± 2.0	0.926	24.5 ± 1.3	0.914
He I EW	0.0127 ± 0.0031	0.844	0.0131 ± 0.0011	0.913	0.0119 ± 0.0011	0.929	0.0119 ± 0.0007	0.905

Note.—Units of the gradients are microhertz per unit activity, where the activity units are defined in Table 3.

linear correlation coefficients in Table 4, suggest that some of the activity proxies display significant deviations from a simple linear relation with the frequency shifts. Figure 13 shows the cyclic "hysteresis" curves for each activity proxy over the three cycles, where each point is the mean of three sets.

The level of correlation seen in each proxy over all three cycles is largely in agreement with the difference between the cycle 22 to cycle 23 ratios of the frequency shifts and respective proxies, given in Table 3, showing that the three-cycle correlation is dominated by differences in response to the reduced magnetic activity of cycle 23. The TSI, however, does not follow this trend. Despite having a similar amplitude ratio over the cycles as the frequency shifts, the TSI produces the poorest three-cycle correlation with the frequency shifts. The TSI correlates very well with the frequency shifts over cycle 22, but it shows the poorest linear correlation of all proxies during cycle 23. The cyclic curves for the TSI (Fig. 13, *bottom left*) show that while it follows a reasonable linear relationship with frequency shifts for cycles 21 and 22, there is a clear divergence from this during cycle 23. In particular, the recent declining phase of the current cycle shows a large variation from the usual TSI-shift relationship observed over the last three cycles.

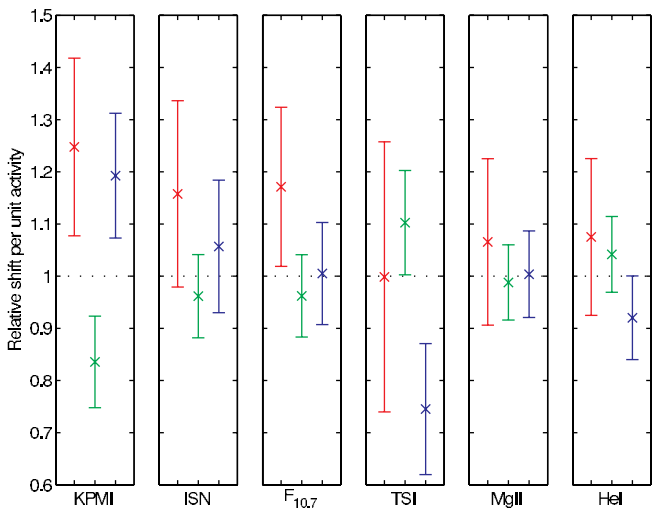

Fig. 8.—Relative shift per unit activity for all proxies over cycles 21 (*red*), 22 (*green*), and 23 (*blue*), normalized to the mean shift per unit activity over all three cycles.

The KPMI displays the greatest relative difference in amplitude between cycles 22 and 23 and shows a correspondingly poor correlation with frequency shift over three cycles. The individual cycle correlations with this magnetic index are notably higher for cycles 21 and 23 than for the magnetically stronger cycle 22. The correlations with magnetic activity in the top left panels of Figures 10 and 12 show that for the high-activity period of cycle 22, KPMI values of higher than 18 G correspond to a small additional increase in mean frequency shift. By comparing the variation in KPMI with the frequency shifts over cycle 22 (9 and 11, *top left*), it is clear that the maximum frequency shift lags the maximum KPMI by approximately 1 yr. For the KPMI, the hysteresis behavior is more pronounced over cycle 22 than it is for cycles 21 or 23 and accounts for the nonlinear relationship between KPMI and frequency shifts in cycle 22. We note that the instrumentation at Kitt Peak was upgraded in 1992, reducing the noise level in the data substantially (Jones et al. 1992). The newer observations have been calibrated to put them on the same scale as the pre-1992 results. Any systematic offsets remaining are at a level far below that which would account for the changes in KPMI-shift behavior observed over each of the solar cycles.

The chromospheric and coronal activity indicators produce the highest correlations and are most consistent with the frequency shifts over successive cycles. They also show very little nonlinear behavior over the successive cycles. The sunspot number is in reasonable correlation with the frequency shifts over cycles 21 and 22 but falls out of agreement during cycle 23. It is the only proxy not to show the largest increase in the activity of cycle 23 during 2002 (see Figs. 9 and 11, *top middle*), which corresponds to the period of greatest frequency shift.

The cross-correlation results show slightly higher values of shift per unit activity than in the mode-fitted case. This is to be expected from the higher value of mean shift at high activity returned using the cross-correlation method. The ranking of the linear correlation coefficients is consistent over both methods.

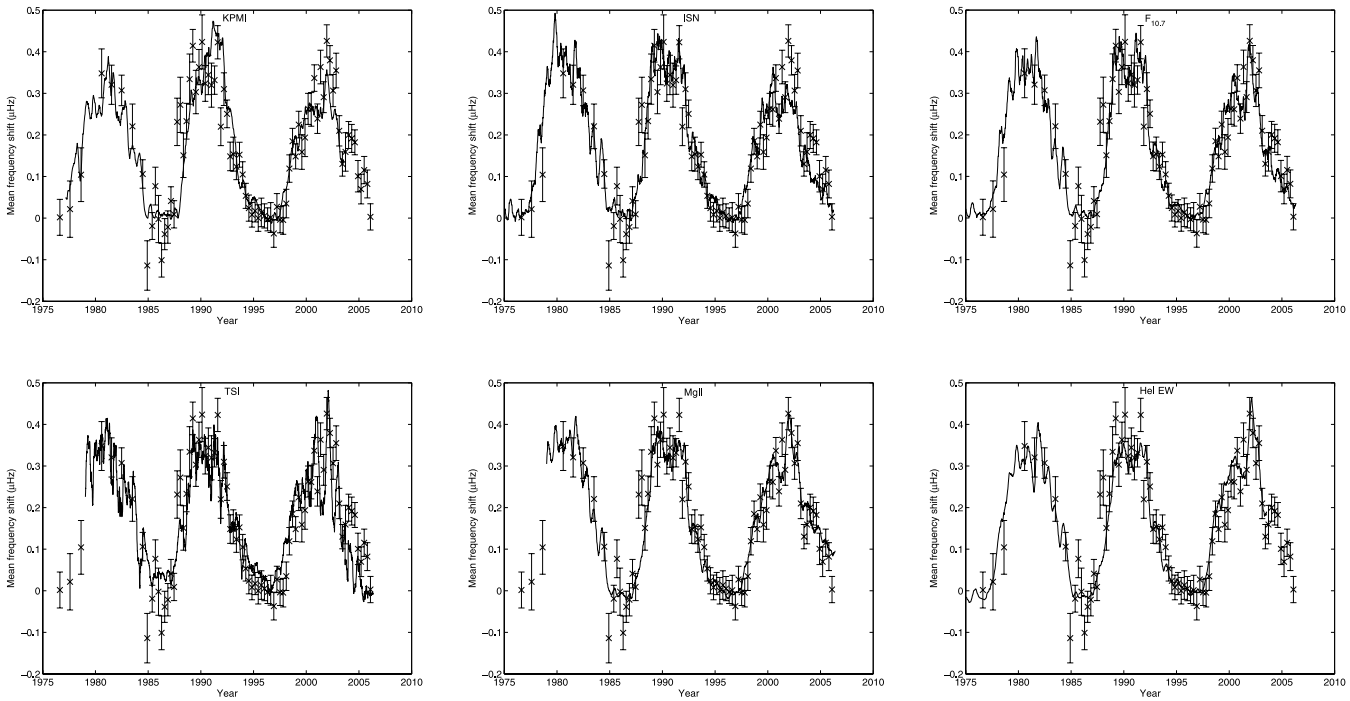

FIG. 9.—Mean mode-fitted frequency shifts (*crosses*) with 108 day smoothed activity levels scaled according to a linear regression over all three cycles (*solid line*). The activity proxy is indicated in the title of each plot.

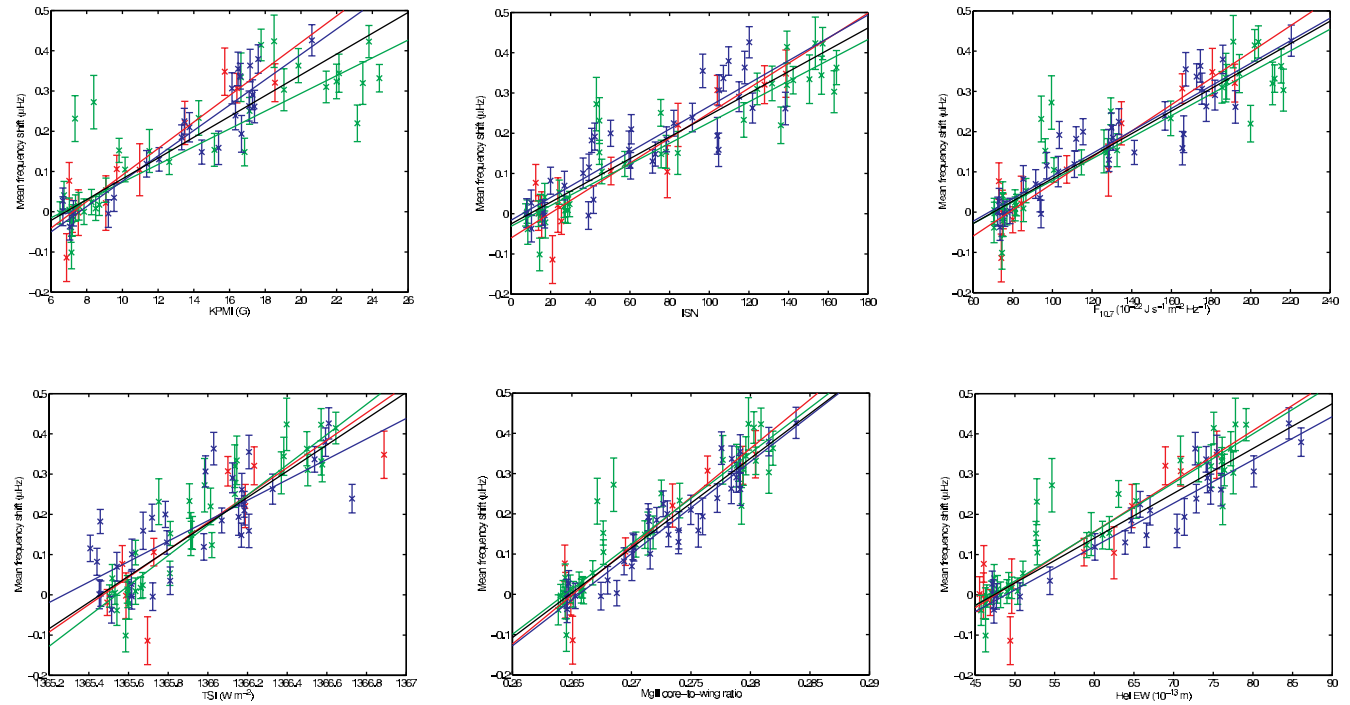

FIG. 10.—Correlations of the average activity proxies corresponding to the epochs used to calculate the mean frequency shifts using mode-pair fitting. The activity proxies are given in the abscissae of the plots. Linear regressions are indicated for cycle 21 (*red*), cycle 22 (*green*), cycle 23 (*blue*), and each complete data set (*black*).

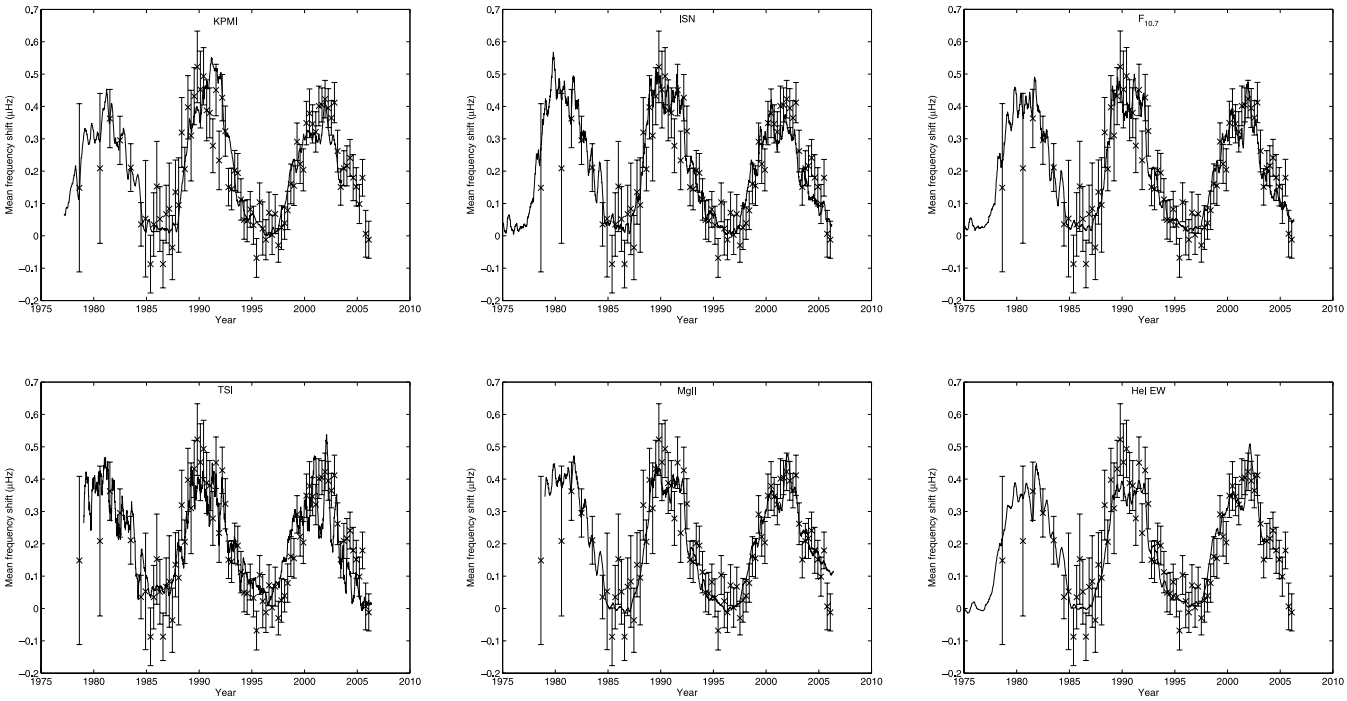

FIG. 11.—Mean cross-correlation fitted frequency shifts (*crosses*) with 108 day smoothed activity levels scaled according to a linear regression over all three cycles (*solid line*). The activity proxy is indicated in the title of each plot. The first two sets are not shown because the error scaling for these very low fill data produces large error bars that are not indicative of the uncertainties in these cases.

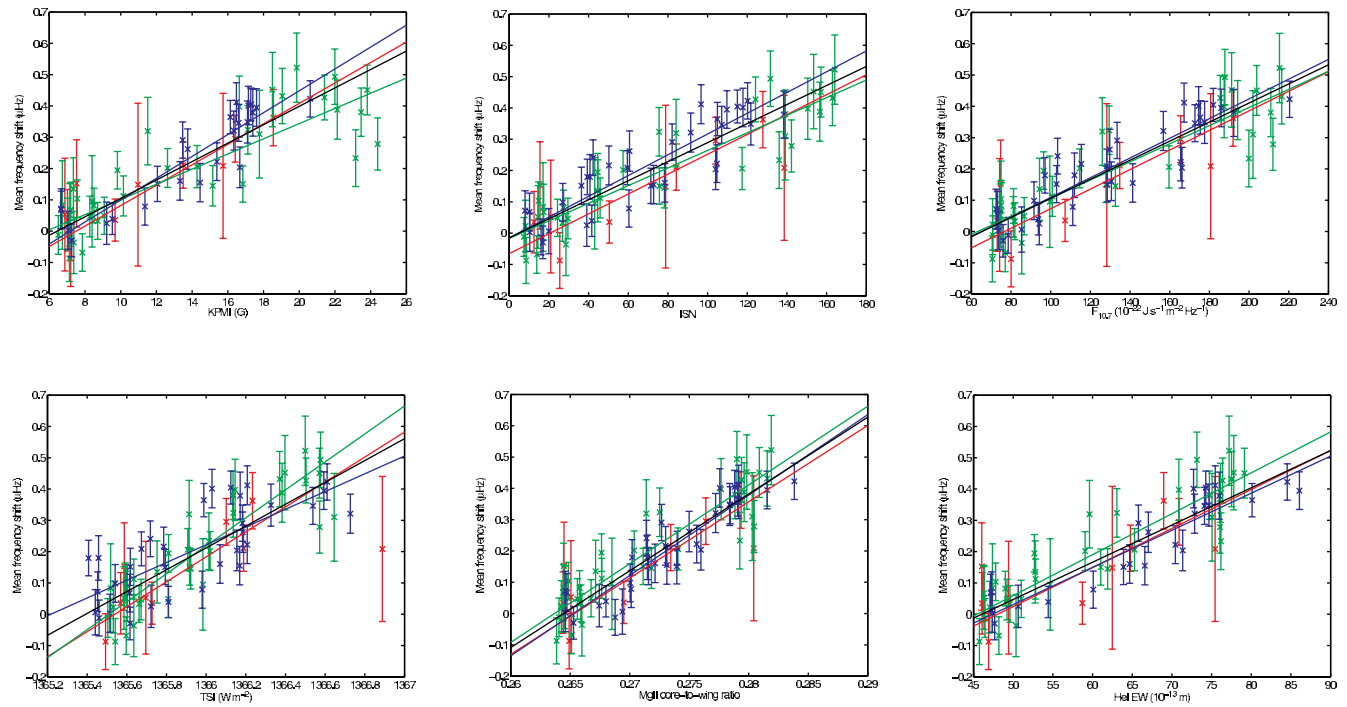

FIG. 12.—Correlations of the average activity proxies corresponding to the epochs used to calculate the mean frequency shifts using cross-correlation fitting. The activity proxies are given in the abscissae of the plots. Linear regressions are indicated for cycle 21 (*red*), cycle 22 (*green*), cycle 23 (*blue*), and each complete data set (*black*).

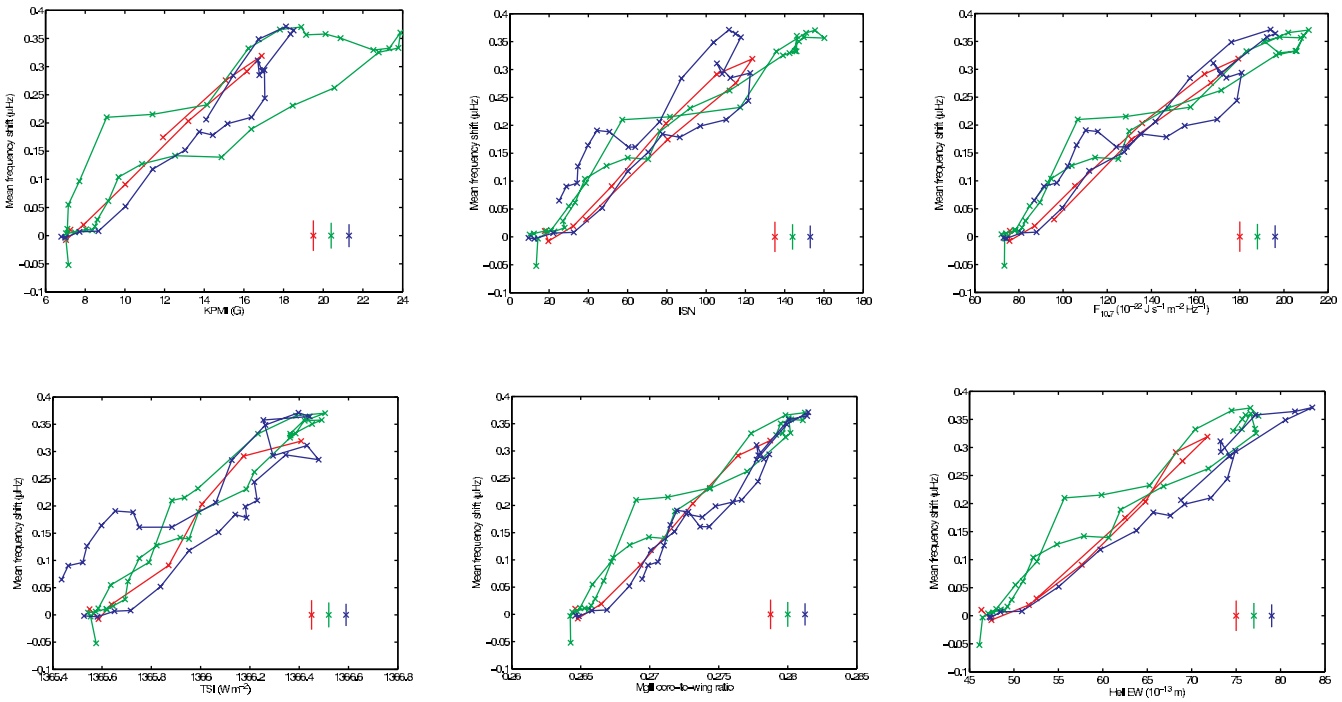

FIG. 13.—Cyclic curves for each activity proxy over successive cycles generated from mode-pair fitted frequency shifts: cycle 21 (*red*), cycle 22 (*green*), and cycle 23 (*blue*). Each activity proxy is given in the abscissae of the plots. The error bars are indicative of the mean error in frequency shift for the corresponding cycle.

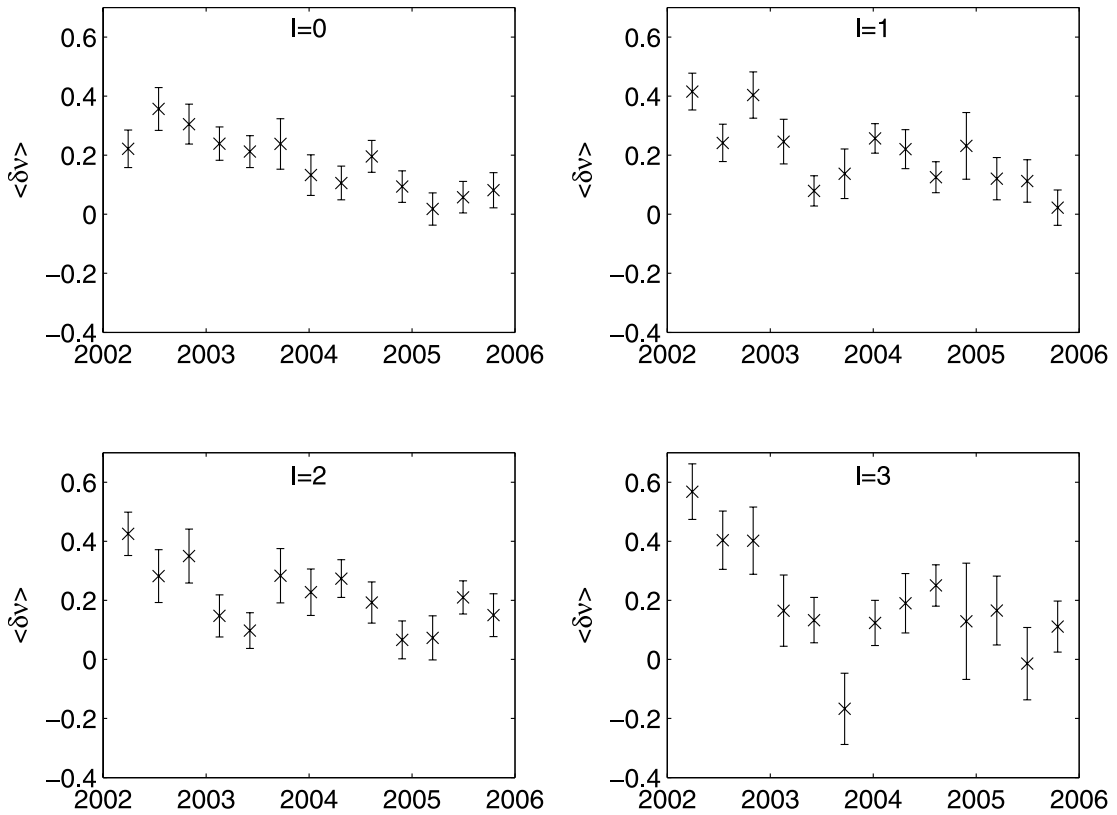

FIG. 14.—Individual degree ($l = 0, 1, 2, 3$) p-mode frequency shifts (in microhertz) over the declining phase of cycle 23.

Both frequency shift extraction methods indicate a sharp decrease in mean frequency shift during the early declining phase of cycle 23 (2003–2004), followed by a small increase in mean shifts. This behavior is not observed in any of the activity proxies. Figure 14 shows the individual degree components of the frequency shifts. The feature is notably absent in the $l = 0$ p-mode frequency shifts and appears to increase in magnitude with increasing l.

5. SUMMARY

We extracted, and analyzed, mean p-mode frequency shifts of low-degree modes covering three solar cycles. The data came from observations made by the Birmingham Solar Oscillations Network (BiSON). The 11 yr cycles covered were cycles 21, 22, and 23. Two methods of analysis were used to extract the frequency shifts: one method used results on fitted frequencies of individual modes, which were averaged to give mean frequency shifts; the other method made use of cross-correlations of power frequency spectra made from data shifted in time.

We compared the behavior of mean frequency shifts of low-l modes with several well-known proxies of global solar activity. The comparisons were made over three solar cycles. The proxies we used were the NSO Kitt Peak Magnetic Index (KPMI), the International Sunspot Number (ISN), the Ottawa/Penticton 10.7 cm radio flux, the total solar irradiance (TSI), the Mg II H and K core-to-wing ratio, and the NSO He I equivalent width.

In this section we summarize the main findings of the paper. We also make some important remarks on use of the cross-correlation technique.

5.1. *Main Findings*

1. We sought to answer the question, which of the activity proxies follows the frequency shifts most closely? We looked not only for similarities but also for differences in behavior, between cycles. We assumed a linear dependence of the mean frequency shifts on each proxy and extracted best-fitting gradients (shift per unit change in activity) on each of the cycles. The proxies that provided the most consistent description of the frequency shifts over all three cycles—in the sense that they had the most consistent gradients from one cycle to another—were the He I equivalent width and Mg II core-to-wing data. The 10.7 cm flux, which is often chosen as a "standard" proxy for comparison with the global-mode frequency shifts, also performed well. The ISN did a good job in cycles 21 and 22 but fell out of agreement with the p-mode data in cycle 23. The KPMI and TSI were the poorest performing proxies.

a) It is interesting to note that the proxies that perform best are those that have good sensitivity to weak-component magnetic flux, which is more widely distributed in latitude than the strong-component flux in the active regions. The Mg II core-to-wing ratio, which is strongly correlated with other chromospheric proxies such as the Ca II H and K data, is often used as a proxy for the bright faculae and network. Harvey & Livingston (1994) showed that the largest contribution to the He I absorption—which serves to provide an excellent "picture" of the chromosphere—comes from the quiet Sun. Even the 10.7 cm radio flux has a contribution from thermal free-free (bremsstrahlung) emission from weak, widely distributed activity (Tapping & DeTracey 1990). In contrast the KPMI fails to capture a sizeable fraction of the weak-component flux (as demonstrated by Harvey 1994).

b) Our analysis showed that despite a reduction in sunspots, and therefore in strong emerging flux, in cycle 23, the frequency shifts continued to rise once the increase of the ISN had flattened out. While the ISN had provided a good proxy for the frequency shifts in the magnetically stronger cycle 22, it clearly fails in the magnetically weaker, and simpler (de Toma et al. 2004), cycle 23.

c) The implication of the cycle-to-cycle comparison is that the modes appear have a significant response to the weak-component flux.

2. We also calculated the minimum-to-maximum swing of each proxy over cycle 22 *divided by* the minimum-to-maximum swing over cycle 23. These ratios provided measures of the relative strengths of cycles 22 and 23, as seen in each proxy. The largest ratios were seen in the KPMI and ISN proxies, reflecting the fact that cycle 23 is a comparatively weak cycle in the stronger, active region flux. Both proxies had ratios that were higher than the ratio calculated for the mean frequency shifts. The KPMI ratio was higher by $\sim 3\ \sigma$ (normalization by mean frequency shift uncertainty), and the ISN ratio was higher by $\sim 2\ \sigma$. The ratios calculated for the 10.7 cm flux and the TSI were closest to the mean frequency-shift ratio (within $1\ \sigma$), while those for the Mg II and He I data were lower, by about $1\ \sigma$.

a) We therefore find a close correspondence of the frequency-shift ratio to the ratios of proxies that have sensitivity to widely distributed, weak-component flux. We regard the good agreement with the TSI ratio as fortuitous. If the TSI variations are regarded as being caused by the relative balance of bright network and faculae and dark spots, they are not a measure of the change in the *total* magnetic flux, to which the modes presumably respond.

b) It is worth adding that calculation of the cycle 22 to cycle 23 ratios of modes of different angular degrees, l, uncovered tantalizing suggestions of an interesting trend. The $l = 0$ ratio had a value close to unity and was therefore in good agreement with the ratios of the Mg II and He I data. The ratio was larger in the $l = 1$ modes and higher still in the $l = 2$ modes, where a value was returned that was similar to the "high" ratios of the KPMI and ISN. However, the difference between the $l = 2$ and 0 ratios was significant only at the $\sim 1.5\ \sigma$ level. The ordering of the ratios of the different l does meet with expectation. In spite of differences between cycles 22 and 23, the implication is that the $l = 0$ frequency shifts showed remarkable consistency from one cycle to the next and that the proxies with good sensitivity to widely distributed flux are in close agreement with the $l = 0$ shifts, as one might expect.

3. Cycle 22 experiences a faster rise acoustically than is suggested by, in particular, the KPMI proxy. There is also evidence of some disparity in the He I equivalent width and Mg II core-to-wing ratio, although at a smaller level than the KPMI correlations, during the rising phase of cycle 22 (Fig. 13, *green*). As the KPMI has the highest relative sensitivity to strong flux, this suggests that the p-mode frequencies were responding to increasing weak flux during the rising part of the cycle that is not detected by the Kitt Peak magnetograms. The fact that the He I and Mg II show some disparity does however demonstrate that no proxy gives a perfect match to the p-mode shifts.

4. During the declining phase of cycle 23, there was a notable deviation of the frequency shifts from the description suggested by many of the activity proxies. The deviation is most apparent for the TSI and accounts for the poor linear correlation of the TSI and mean frequency shifts during cycle 23. The feature is also prominent in the ISN and, to a lesser extent, the 10.7 cm flux and Mg II correlations. The KPMI and He I equivalent width data were not collected during this period. Figure 14 shows that the deviation is notably absent in frequency shifts made from just the $l = 0$ modes. The deviation is present in data made from modes

with $l \geq 1$ and gets stronger going from $l = 1$ to 3. For higher degree modes, the eigenfunctions are restricted to lower latitudes. The inference on this result is that *p*-modes were responding to the reduced low-latitude activity during the declining phase of cycle 23.

5. The Mg II data have been shown to provide a robust, consistent description of the *p*-mode shifts over three solar cycles. These Mg II data are highly correlated with Ca II H and K data. Ca II data have been collected on many stars for use as a proxy of stellar activity and to uncover evidence for, and then track, stellar activity cycles. Our work here suggests that, provided allowance is made for the effect of the rotation axis offered by stars, stellar Ca II data may be used to give an indication of likely acoustic variability (see Chaplin et al. 2006a).

5.2. *Comments on Analysis Methods: Peak-fitting versus Spectrum Cross-Correlations*

1. Points in the cross-correlation function of two power spectra are strongly correlated. The strong correlation causes underestimation of the uncertainty in the frequency offset of the maximum of the function—the estimated mean frequency shift—if no allowance is made for the correlation, and uncertainties are estimated in accordance with the usual assumption of the input data being statistically independent. We have estimated the true uncertainties on the frequency shifts given by cross-correlation fitting through use of 50 simulated solarFLAG time series. In the case of 108 day spectra, fitted over a range in frequency lag of ± 20 μHz, the true errors are actually a factor of ~ 7.5 larger than those given by the usual independent, least-squares fit approach. Once a suitable scaling is applied to correct for this overestimation, the cross-correlation method is shown to have a precision that is comparable to that given by averaging shifts of individual mode frequencies (where the frequencies come from results on fits to modes in pairs).

2. Closer scrutiny of the results demonstrated that the true precision in the cross-correlation method is actually slightly inferior to that in mode-pair fitting results in cases in which a high duty cycle is available. However, for short duration, low-fill data, the cross-correlation method may produce results of similar or slightly higher precision. This is because it can be difficult to converge on reliable peak fits in such data. Regardless, fitting the cross-correlation function to obtain frequency shifts provides a useful comparison with the shifts obtained by mode fitting. Furthermore, there is often an element of selection associated with use of the mode-pair fitting results, i.e., the issue of defining objective fit-rejection criteria can be a tricky one, and often visual inspection of the fits by an experienced eye does a better job of weeding out poor results. Use of the cross-correlation function to some extent removes this selection effect and any influence it may have on the mean frequency shifts.

3. The functional form of the cross-correlation function used previously by Jiménez-Reyes et al. (2001) is unsuitable for high-fill data, where the cross-correlation sidebands are dominated by $l = 0, 2$ and $l = 1, 3$ overlapping pairs, rather than the diurnal sideband structure (shown in Fig. 2). The more complicated function we have used in this paper (eq. [1]), which takes account of the mode-pair sidebands, is found to work well on both high-fill and low-fill data.

We are indebted to J. Allison, I. Barnes, B. Jackson, and S. Hale for their technical and analysis support in Birmingham and to former colleagues, in particular C. P. McLeod, H. Williams, J. Litherland, and R. Lines. The authors would like to acknowledge the significant contributions to helioseismology made by their recently deceased colleague, G. R. Isaak. We also thank P. Whitelock and P. Fourie at SAAO, the Carnegie Institution of Washington, the Australia Telescope National Facility (CSIRO), E. J. Rhodes (Mt. Wilson, California), and members (past and present) of the IAC, Tenerife. The 10.7 cm radio flux data and sunspot numbers were obtained from the Solar Geophysical Data Center. The NSO/Kitt Peak data are produced cooperatively by NSF/NOAO, NASA/GSFC, and NOAA/SEL. BiSON is funded by the UK Particle Physics and Astronomy Research Council (PPARC). G. A. V. acknowledges the support of PPARC. We acknowledge the solar Fitting at Low Angular degree Group (solarFLAG) for use of their artificial helioseismic data.

REFERENCES

Andretta, V., & Jones, H. P. 1997, ApJ, 489, 375
Antia, H. M., Basu, S., Hill, F., Howe, R., Komm, R. W., & Schou, J. 2001, MNRAS, 327, 1029
Brookes, J. R., Isaak, G. R., & van der Raay, H. B. 1978, MNRAS, 185, 1
Chaplin, W. J., Elsworth, Y., Isaak, G. R., Lines, R., McLeod, C. P., Miller, B. A., & New, R. 1998, MNRAS, 300, 1077
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. 1999, MNRAS, 308, 424
———. 2002, MNRAS, 330, 731
———. 2004, MNRAS, 352, 1102
Chaplin, W. J., Houdek, G., Elsworth, Y., & New, R. 2006a, MNRAS, submitted
Chaplin, W. J., et al. 2006b, MNRAS, 369, 985
Christensen-Dalsgaard, J., & Berthomieu, G. 1991, in Solar Interior and Atmosphere, ed. A. N. Cox, W. C. Livingston, & M. Matthews (Tucson: Univ. Arizona Press), 401
de Toma, G., White, O. R., Chapman, G. A., Walton, S. R., Preminger, D. G., & Cookson, A. M. 2004, ApJ, 609, 1140
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
Elsworth, Y., Howe, R., Isaak, G. R., McLeod, C. P., Miller, B. A., New, R., & Wheeler, S. J. 1995, A&AS, 113, 379
Elsworth, Y., Howe, R., Isaak, G. R., McLeod, C. P., & New, R. 1990, Nature, 345, 322
Fröhlich, C. 2000, Space Sci. Rev., 94, 15
Harvey, J. W., & Livingston, W. C. 1994, in IAU Symp. 154, Infrared Solar Physics, ed. D. M. Rabin, J. T. Jefferies, & C. Lindsey (Dordrecht: Kluwer), 59
Harvey, K. L. 1994, in The Sun as a Variable Star: Solar and Stellar Irradiance Variations, ed. J. M. Pap et al. (Cambridge: Cambridge Univ. Press), 217
Howe, R., Komm, R. W., & Hill, F. 2002, ApJ, 580, 1172
Jiménez-Reyes, S. J., Corbard, T., Pallé, P. L., Roca Cortés, T., & Tomczyk, S. 2001, A&A, 379, 622
Jiménez-Reyes, S. J., Régulo, C., Pallé, P. L., & Roca Cortés, T. 1998, A&A, 329, 1119
Jones, H. P., Duvall, T. L., Jr., Harvey, J. W., Mahaffey, C. T., Schwitters, J. D., & Simmons, J. E. 1992, Sol. Phys., 139, 211
Nigam, R., & Kosovichev, A. G. 1998, ApJ, 505, L51
Régulo, C., Jiménez, A., Pallé, P. L., Perez Hernández, F., & Roca Cortés, T. 1994, ApJ, 434, 384
Tapping, K. F., & DeTracey, B. 1990, Sol. Phys., 127, 321
Viereck, R., Puga, L., McMullin, D., Judge, D., Weber, M., & Tobiska, W. K. 2001, Geophys. Res. Lett., 28, 1343
Woodard, M. F., & Noyes, R. W. 1985, Nature, 318, 449