Chapter 3

The Analysis of the Fourier Tachometer Data

In this chapter I will describe the Fourier Tachometer and parts of the analysis of the data obtained with it. This chapter is the only one in the thesis that is not simply an introduction to the papers included in the appendices. Even for this chapter, some of the ideas used are described in the included papers, in particular the principle behind the analysis of the time-series is described in Schou (1993a, cf. App. 5). Bachmann, Schou, and Brown (1993, cf. App. 6) contains a description of the analysis of the 1989 dataset, with some details not included here (and vice-versa). Schou and Brown (1993, cf. App. 4) contains a description of the fake data program used to test the analysis programs.

I will start with a very short description of the Fourier Tachometer instrument and some of the details of when and how it was operated. Then I will again briefly describe the part of the reduction up to the generation of the time-series and the subsequent analysis code. I have not been involved much in the part of the analysis up to the generation of the time-series, but as it affects the subsequent analysis, I have included this short section. Technical details of how the different parts of the code have been implemented are left out. Descriptions of many of these parts can be found in Haynes (1992).

Section 3.3 is a description of some of the more-or-less processed data from the various parts of the reduction described in the previous section. I will give a description of the general appearance of the various types of images and describe some of their statistical properties.

The next two sections (3.4 and 3.5) contain details about the reduction of the so-called aj89 and aj87 datasets. As the aj89 dataset was the first to be reduced, and the one most effort has been put into, it will be described first. Also, the details common to the two datasets will only be described in connection with the aj89 dataset.

In section 3.6, some of the tests performed to look for sources of possible systematic errors will be described. This search was in part motivated by the so-called 'bump' discovered in a_1 for the aj89 dataset (and later in other datasets), and described in the next section. While the search was motivated by the 'bump', these tests have given a good feeling for the magnitude of systematic errors in the final results one can expect from various incorrect assumptions concerning the raw data.

As this bump is perhaps the most surprising feature found in the analysis, a whole section (3.7) has been devoted to it and its possible explanations. While no cause has been identified, several possibilities have been excluded.

Finally in section 3.8, I will briefly compare the results from the different datasets reduced. As some descriptions of the time variability of various parameters can be found in Bachmann, Schou and Brown (1993, cf. App. 6), and is similar to that found by others (eg., Libbrecht and Woodard 1990), I will not go into details in this section. When we have analyzed the remaining datasets, we plan to publish a more complete comparison of all the Fourier Tachometer datasets.

3.1 The Fourier Tachometer

This is only a short description of the Fourier Tachometer (strictly speaking, the Fourier Tachometer MkII), not a complete description. Some parts will be described in more detail than others, particularly those that are important for the reduction of the data. More detailed description of some parts and operating procedures may by found in Brown (1980), Evans (1980), Brown (1984) and Morrow (1988). Unfortunately, the instrument, the operating procedures etc. have not been described in detail in any publications.

This section is by no means a presentation of original ideas, but mostly a summary of other works, as relevant to the parts of the reduction I have been concerned with. It also serves as an introduction to some of the things about the instrument that are discussed later.

The Fourier Tachometer was build over the period 1980-1984 and operated jointly by the High Altitude Observatory (HAO) and the National Solar Observatories (NSO) from 1984-1990. Until October 1985 the instrument was at Sacramento Peak Observatory, Sunspot, New Mexico. For the remainder of the period it observed from NSO's headquarters in Tucson, Arizona. In addition to this change several things have changed over the years, most of these *should*, however, not have affected the operation significantly, although some of them might be the cause for some of the differences between the years observed. After the last run the instrument was permanently disassembled, and it is thus not possible, or very difficult, to find the causes of many problems discovered in the later analysis.

It should probably also be mentioned that I have never seen the instrument, only some parts, let alone operated it, and that I am therefore not well qualified to comment on the various technical details.

3.1.1 Description of the Instrument

The instrument has been described in Brown (1980, 1984). Figure 3.1 shows the general idea behind the instrument. Basically, the instrument consists of a light feed (not shown), a collimating lens (not shown), a prefilter, a Michelson interferometer with a path length modulator in one arm, a (re)imaging lens, and a detector array. The light was provided by a heliostat tracking the Sun and a 6 inch refractor. The prefilter selects a fairly narrow band around the desired solar absorption line.

The Michelson interferometer modulates the prefilter passband by its transmission profile which is given approximately by $0.5*(1+\cos(k\lambda+\phi))$, where $k = 1, 2, ...$ and λ is the wavelength. By varying the phase ϕ in steps of $\pi/2$ (by changing the pathlength in one leg), both the real and imaginary part of the Fourier component at wavenumber k can be found, modulus a zero-point offset. Given that the phase shift of this Fourier component is proportional to the shift of any features, such as the target absorption line, in the passband, it is possible to find the Doppler shift of the line and hence, the solar surface velocity. In practice the intensity is integrated over four quadrants, which is enough to get both the values of the Fourier component and the constant offset of the intensity. The summation of the subimages (taken every 1/60s) was done

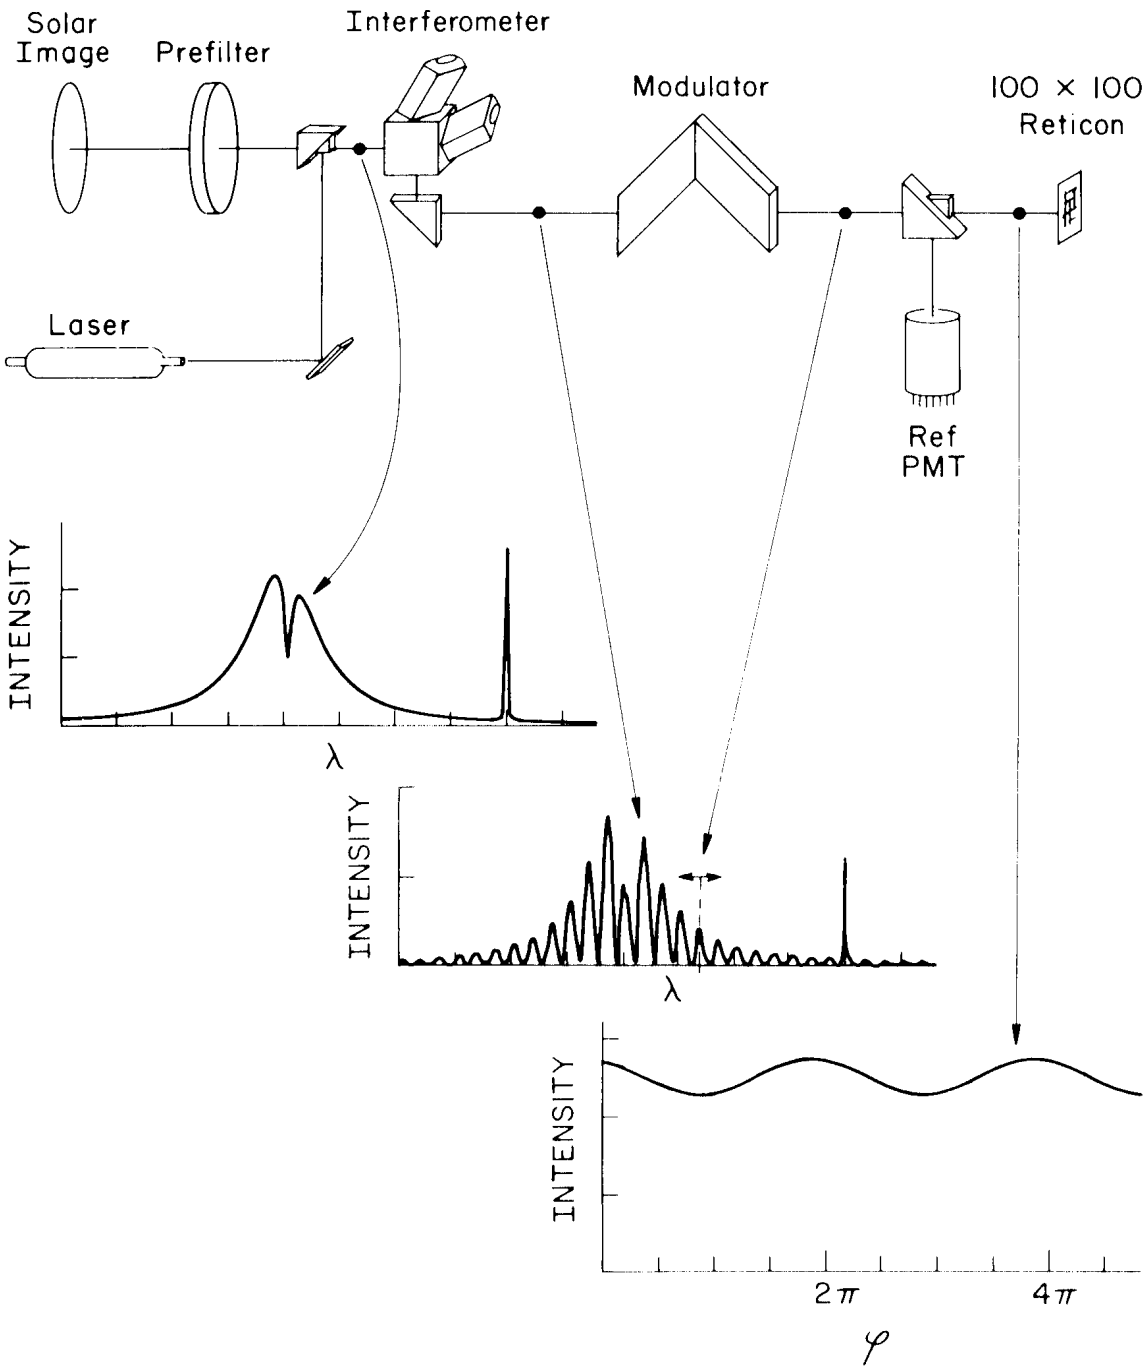

Figure 3.1: Schematic of the Fourier Tachometer.

in real time by a dedicated computer. Note that since the equivalent width of the spectral line is small compared to the width of the passband for the prefilter (which is fairly wide in order to make it uniform), the modulation of the signal is very small, of the order 1%. This is unfortunate, as the (photon) noise depends on the signal, not the modulation. If the prefilter does not have a uniform passband (which will tend to be the case if it is very narrow), there is a risk that drifts will introduce systematic errors in the measured velocities.

The detector used is a photo diode (Reticon) array with 100×100 pixels, the solar images typically had a radius of 44-48 pixels. The detector and the associated electronics had many 'interesting' features, some of which affect the data reduction and will be mentioned later. One feature was that in the original use of the detector, a bias was applied to each line read out, such as to make the integrated intensity over the line constant. This was disabled in the Fourier Tachometer, but unfortunately some noise was introduced this way, in the form of a random offset added to each row. Also, the even and odd columns were brought out to different amplifiers, which, as shown later, also has led to some problems.

Given the design of the lightfeed, the orientation of the solar image is not fixed with respect to the optics during the day. Instead of correcting for this rotation in the later analysis, the detector is rotated by an amount found from an ephemeris, such as to keep the orientation fixed on the detector. Ideally, the instrument would be adjusted such that the solar image would have north up, but for some runs this was not done accurately.

3.1.2 Mode of Operation

Although the instrument was capable of observing different spectral lines and with different image scales, most observations (and all those used here) used the Fe I 6678Å line and the full disk observing mode. Except for setup, calibration, etc., the instrument would take intensity and velocity images every 60s or 75s for as large a fraction of the day as possible. Every hour this sequence would be interrupted to take an 'objective image', in which an image of the objective is formed on the detector rather than an image of the Sun. Occasionally, other forms of calibration images were also taken. The original purpose of taking the objective images were to determine the zero point velocity introduced by the optics. Originally, these images were then interpolated in time and subtracted from the velocity images. In the current reduction code this is not done for reasons explained later.

3.1.3 Observing Periods

The observing runs made are summarized in table 1. Note the large variation in length of the observing runs, cadence, and the aliases used for the runs. Figure 3.2 shows the combined time coverage of all the runs.

Observing periods

Name	Year	Start	End	Cadence	Aliases
FULL84	1984	October 28	November 14	60s	
BILLOB1	1986	May 17	June 28	60s	
BILLOB2	1986	September 10	October 30	60s	
FULL87	1987	April 11	July 5	60s	Spring 87, aj87
POLE87	1987	November 8	December 4	75s	nd87
FULL89	1989	April 3	June 30	75s	Spring 89, aj89
FULL89F	1989	August 23	December 13	75s	
FULL90	1990	April 2	July 1	75s	
FULL90F	1990	August 17	December 19	60s	

Table 1: List of observing periods. Dates are inclusive and local time. The POLE87 dataset was coordinated with observations from the South Pole.

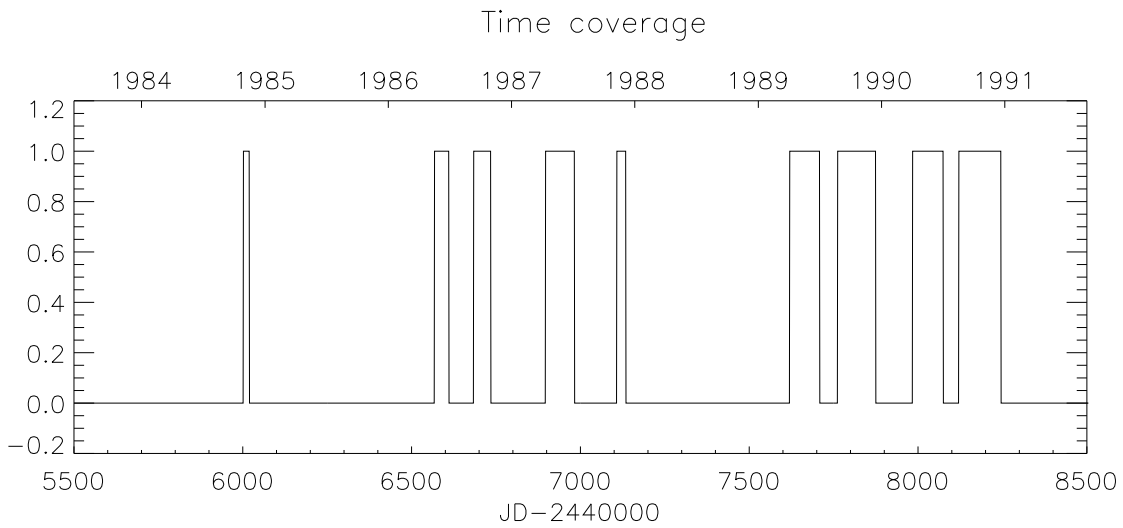

Figure 3.2: The time coverage for all the full disk observing runs. The lower x-axis shows the Julian day number, the upper axis shows the beginning of each year. Single missing days inside runs have not been excluded.

3.2 Reduction Procedure

Again this is only a brief description of the reduction process up to the generation of the time-series of Y_l^m coefficients, and Fourier transforms thereof. Some parts have been described in more detail in Schou (1993a, cf. App. 5), Schou, and Brown (1993, cf. App. 4), Bachmann, Schou and Brown (1993, cf. App. 6), and Haynes (1992).

The analysis of the time series/Fourier transforms is described in Schou (1993b, cf. App. 5), but will be summarized at the end of this section.

The reduction consists of several major substeps, each of which is described below. Some of these substeps, in particular the RDC code, actually consist of a number of separate programs, but this will not be addressed here.

Many of the programs have been around for a while and thus contain various 'historical' features not used any more. They were also written at a time when the available computer resources were very different from now and are thus somewhat inefficient and would probably benefit from some major reworking.

As with the instrument, I have not run many of the different parts of the reduction code myself, and while I have looked at substantial fractions of the codes, I do not know all the details of the various codes or why some things were implemented the way they were.

3.2.1 The RDC Code

This part of the reduction code processes the raw images. The first step is to classify each image as normal velocity/intensity images, objective images or various other types not used here. The raw velocity images are then changed to images of velocity in standard units.

There are a few bad pixels on the detector. At the moment nothing is done about those as their contribution to the overall noise is fairly small. This would probably be the place to interpolate over these pixels if it were decided that they were a problem.

The intensity images are then used to find the center and radius of the solar image. Both the intensity and velocity images are then centered and the intensity images are left unchanged as they are not used in the subsequent analysis. The interpolation used to shift the images obviously introduces a certain amount of blurring, degrading the effective point spread function (PSF). As discussed later it is actually possible to avoid this interpolation, but for the purpose of being able to look at the images as a movie without them drifting around during the day, this has not been attempted.

There has been some concern about the accuracy of the algorithm used to find the center and radius of the images. In particular, it appears that for some years, the motion of sunspots during the day is not consistent with the motion from day to day caused by the solar rotation. In other words, it would appear that the supposedly centered images are drifting through the day. A known problem is that the images often drift far enough to hit the edges of the detector (as discussed in section 3.7) which, although the centering routine is supposed to get around this, could lead to problems.

Recently, it was also discovered that the solar images were, at least for some datasets, somewhat oblate. An algorithm to detect and correct this, together with the centering, was implemented, but none of the datasets used in the rest of the thesis have been analyzed with this newer version of the code.

A model of the solar surface velocity caused by the solar rotation and the limb shift is then subtracted from the velocity images in order to reduce the large velocity gradient across the

image. As the P-angle rotation (see later) is no longer done before this step and the model does not take the P-angle into account, this removal is far from perfect, but as the term is almost constant, it should be of little concern.

In the previous version of the code, that was used for the reduction of some datasets, the objective images were subtracted from the velocity images at this point. In order to do this, each of the objective images was first rotated back to the position at the beginning of the day. The reason for having to rotate these images is that the detector is rotated with respect to the optics, to compensate for the rotation of the solar image during the day. As the effects that the objective images are supposed to remove are fixed with respect to the optics, the images must be rotated.

Each velocity image was then also rotated back to the start of the day and an interpolation between the two adjacent objective images to the time of the current image was performed. This interpolated objective image was then subtracted from the velocity image, and the velocity image was rotated back to its original position (where it supposedly has a fixed effective P-angle).

Obviously, these two rotations of the velocity images degraded the PSF, due to the interpolations performed. A way around this is to rotate the objective images to the position of the velocity images, instead of the other way around. This reduces the number of rotations of the velocity images by two. It is important to note that the objective images are generally fairly smooth, and thus presumably less affected by interpolations (rotations) than the velocity images, which contain a lot of small scale structure.

Instead of doing this modified objective removal, it was decided to not subtract any objective images at all. That the objective image removal is essentially without effect is fairly obvious when one considers the fact that the images are almost constant and only taken every hour. Anything removed changing on a timescale of an hour is only going to change the signal at frequencies of the order a few times 1 hour^{-1} $\approx 278\mu Hz$ and lower, which is far below the lowest frequency p-mode observed.

A low order polynomial of the form

$$p(x, y) = a + bx + cy + dx^2 + ey^2 + fxy$$

where x and y are the coordinates relative to the center of the Sun, is fitted to the velocity image and subtracted in order to remove large scale gradients in the image. The fit is only done out to a fractional solar radius of roughly 0.8 (that is within the area of the images used in the later parts of the analysis). While this removes modes with low l and m, which might be the cause of problems with these modes discussed later, it does remove some of the spurious velocity gradients introduced by imperfections in the Michelson interferometer and some of the residual solar velocity (like the limb shift) not removed when the model of the solar surface velocity was subtracted.

In this part of the code the outermost 10% (by radius) of the velocity images are also apodized away, as they can be extremely noisy.

3.2.2 The YLM Code

This part of the reduction code projects the velocity images onto masks designed to isolate individual (l, m)s. Parts of this code have been described in some detail in Haynes (1992). Some of the general considerations for this part of the reduction and the physical background have also been described in Schou and Brown (1993, cf. App. 4).

As the solar image generally does not have the solar rotation axis up and the B-angle is generally nonzero, the first step of this code is to remap the image coordinates to a uniform longitude/sine latitude grid. Also an apodization as a function of fractional radius r on the solar image of the form

$$Ap(r) = \begin{cases} 1, & \text{if } r \le r_1 \\ 0.5 * (1 + \cos(\pi \frac{|r| - r_1}{r_2 - r_1})), & \text{if } r_1 \le |r| \le r_2 \\ 0, & \text{if } r \ge r_2 . \end{cases}$$

is applied to the images (see figure 3.3), where r is the fractional radius, $r_1 \approx 0.8$ and $r_2 \approx 0.9$. The reason for applying an apodization is that the signal is dropping very fast close to the solar limb, due to the velocity projection factor, while the noise is increasing, and that using the outermost parts of the solar image are thus likely to contribute more noise than useful signal. The exact limits 0.8 and 0.9 were determined fairly arbitrarily.

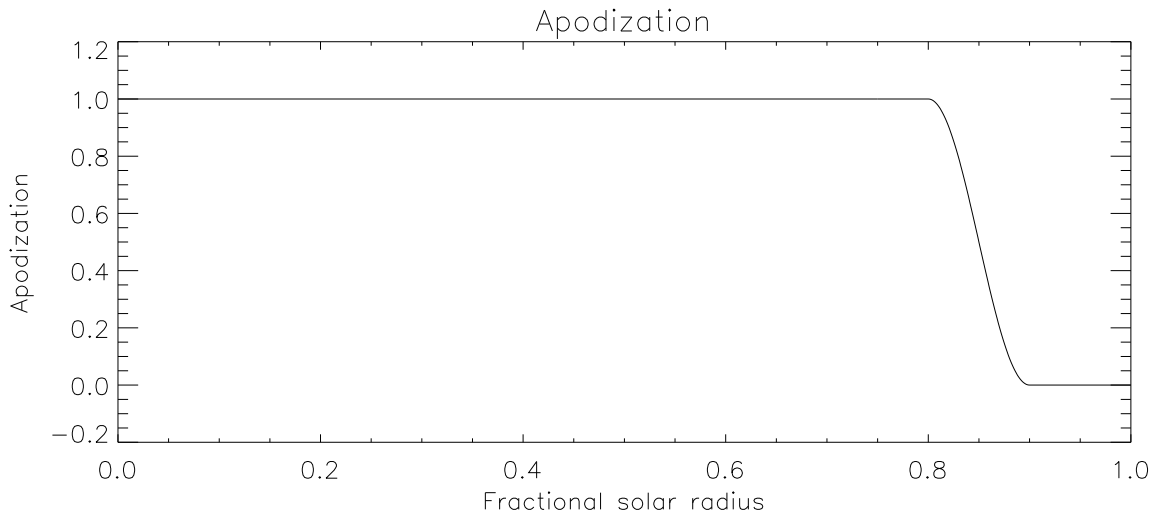

Figure 3.3: The apodization function as a function of the fractional solar radius.

The radius used to do the remapping is the radius from the RDC code, plus a constant determined such as to make the spatial sidelobes of the modes with $\Delta l = \pm 1$ in the power spectra from the following parts of the reduction code have the same magnitude for selected ls. This is obviously not quite satisfactory, but it is at least a way to estimate the correction that needs to be applied to the radius determination from the RDC code.

The fact that the Sun is tilted due to the (variable) B-angle means that the velocity projection factor and the apodization changes slowly with time as a function of longitude and latitude, and is somewhat different from what one would expect from B=0. This leads to a different (and time varying) weighting of the northern and southern hemispheres, especially at high latitudes. While it is possible to correct this effect in this part of the code, nothing is done at the moment.

Notice that the P-angle needed for this part of the reduction is not necessarily equal to the P-angle obtained from an ephemeris. As a matter of fact, it was attempted to align the camera such that the effective P-angle (the angle between the solar rotation axis and up on the detector) was 0. Unfortunately, this was not always the case and the effective P-angle has to be determined from the observations. This can be done in different ways.

One way to determine the P-angle is to follow one or more sunspots as the Sun rotates. For some runs close to the solar minimum it is very difficult to find *any* sunspots to follow. Also, the sunspots change shape and move slightly, leading to some uncertainty. However, it appears that it is possible to determine the P-angle with an accuracy of about $\approx 0.5^\circ$ using this method.

Another way is to look at the direction of the gradient of the velocity caused by the solar surface rotation. While this method can, in principle, give a very precise value for the P-angle, it is susceptible to a number of systematic errors. In particular, it is necessary to remove the spurious velocity introduced by the optics. This can be done by subtracting the objective images, but even after doing this the P-angle is not constant during a day. The P-angle determined this way for one day has been shown in figure 3.4. As is clear from this figure, this method has serious problems. This method should, however, be very useful for catching occasional large slips. Such slips, where a whole day has a different P-angle, have been observed on at least one occasion.

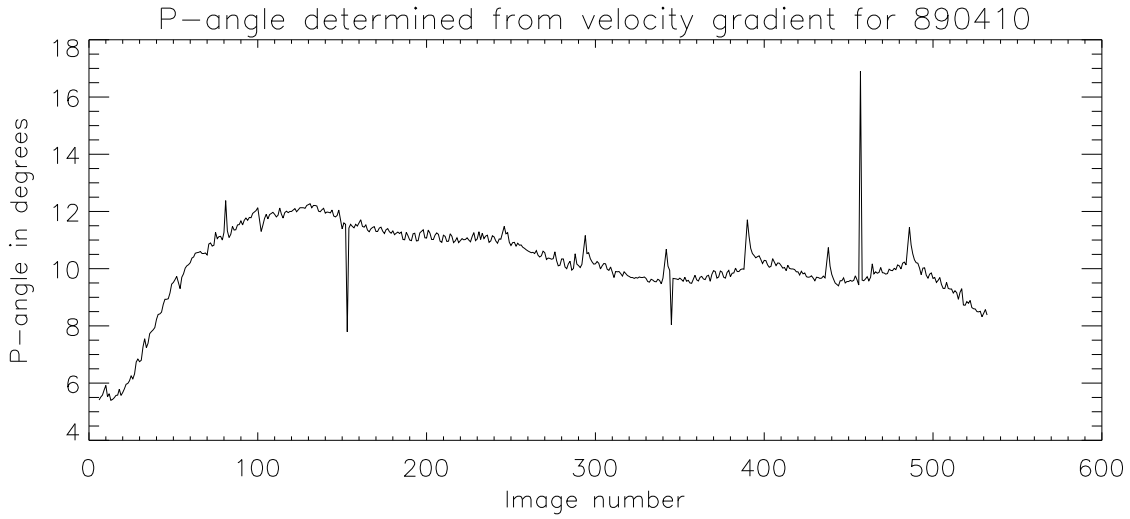

Figure 3.4: P-angle as a function of time of day for the day 890410.

At the moment, we have used P-angles determined from sunspots, as we have had problems in determining the true accuracy of the velocity gradient method and what the systematic

problems are due to.

At each latitude in the interpolated image, the pixel values as a function of longitude are then Fourier transformed to isolate the individual ms. At each m the Fourier transform coefficients are multiplied by P_l^ms as a function of latitude to give the final coefficients for each l. The coefficients for each image are then written out. Notice that the coefficients are written out image by image. At the times when objective images were taken, the program interpolates to fill the gap. This gap filling is likely to introduce small temporal sidelobes with a spacing of 1 hour^{-1}, but we have not observed those, presumably because of the small amplitude.

At the moment, only the coefficients for the even ms are calculated and written out, but the capability exists in the program for calculating the coefficients for both the even and the odd m values, which, as discussed later, could prove useful.

A possible improvement would be not to center the images in the RDC code and instead combine the centering with the interpolation to the latitude/longitude grid. This should improve the PSF by a significant amount, as discussed in a later section. The total computation time should also go down somewhat. The drawback of this modification is that the images output by the RDC code will no longer be centered, making them more difficult to look at. A partial solution making the images more easy too look at would be to only do the shift to the nearest integer number of pixels in the RDC code and leave the subpixel interpolation (and the oblateness correction) to this part of the code.

Another possibility is to apply the masks directly to the uncentered images and recompute the masks instead. This option is, unfortunately, extremely computationally intensive, but should improve the effective PSF.

3.2.3 The Strext Code

As the output of the YLM code is ordered by image number instead of l and m the data have to be transposed. The transposition is done in this part of the code. As getting all the coefficients for a given (l, m) involves reading parts of all the output files from the YLM code, this is very I/O intensive. This is probably the most inefficient part of the code at the moment. This piece of the code also fills the larger gaps in the time-series with zeros.

An example of a short section of a time-series is shown in figure 3.5. Notice the clear signature of the 5-min oscillations; the rapid variation of the amplitude is due to the beating of the different modes going into the time-series.

3.2.4 The MNU Code

This program produces power spectra from the time-series files. I have not used the power spectra for my analysis, but they are used for other analysis methods we have used.

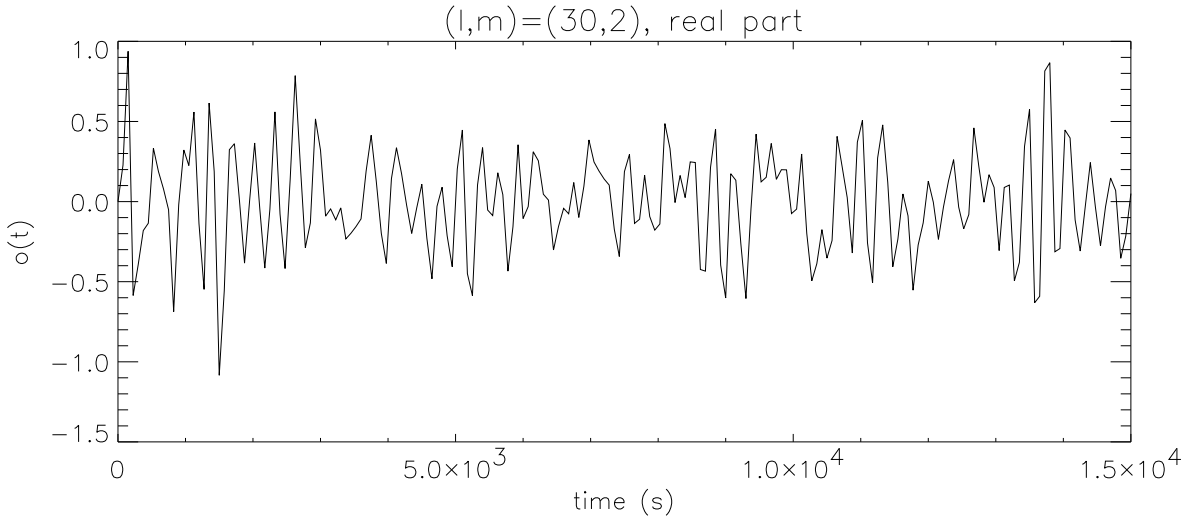

Figure 3.5: A short stretch of a time-series with $(l, m) = (30, 2)$ from the day 890403.

3.2.5 The ttof Code

This piece of the code takes the time-series from the strext program, Fourier transforms them, and writes them out. The Fourier transforms are not restricted to lengths of 2^n, unlike the MNU code. No bad point removal is performed.

3.2.6 The Analysis Code

The principle behind the main analysis code is described in Schou (1993a, cf. App. 5).

Briefly the program fits one (l, n) at a time, taking into account both the leakage among the different ms and the the leakage from other ls and ns. Contrary to other programs used for the fitting, this program uses the raw Fourier transforms, not the power spectra, in order to be able to take into account the correlations among the different ms caused by the leakage. In addition to the leakage, such things as the correlations in the noise between different ms and the temporal sidelobes caused by the time gaps are taken into account.

As mentioned, the program fits one (l, n) at a time, using a set of initial estimates for the various parameters. In the fit the parameters for modes with l or n different from the target mode are held fixed at their initial values (except for the amplitudes, which can vary, cf. Schou 1993a (App. 5)). For the target mode all parameters are varied. For most of the results presented here, the set of parameters used for each (l, n) was the mode frequency, the power in the mode, the mode linewidth, the level of the background power, and a number of a-coefficients (typically 6).

Only modes within a certain Δl of the target modes are taken into account, for all the results presented here, $\Delta l = 6$. For each of these ls, all ns and ms are taken into account. For the target (l, n), all ms out to a distance depending on the linewidth are taken into account. The fit is done using all frequencies within a certain distance from the expected center frequency.

The size of this frequency window is determined by the linewidth, and is, except for the lowest and highest frequency modes, five times the initial estimate for the mode linewidth.

As the fitted parameters for the target mode depend on the assumed parameters for the other modes (which are held at their initial values), cf. section 3.6, the fit needs to be iterated until the mode parameters have converged. In each iteration the parameters from the previous iteration are used as initial estimates, filling in the missing modes by interpolating/extrapolating the parameters for the fitted modes. Unfortunately, it turns out that the feedbacks can be rather large, making it necessary to do a large number of iterations or use various tricks to speed up the rate of convergence.

In addition to the fitted parameters, the program also provides an estimate of the covariance matrix for the parameters and (at the moment), two types of diagnostics. The covariance matrix, and thereby the estimated errors on the parameters, is an internal estimate, based on the second derivatives of the merit function, as described in Schou (1993a, cf. App. 5).

The diagnostics have also been described in Schou (1993a, cf. App. 5). At the moment two types are implemented. The first diagnostic is essentially a measure of the deviation of the crosstalks between neighboring ms from those given by the estimated crosstalks. This measure is constructed by comparing the correlation between all pairs of points in the Fourier transforms with the same frequency and a Δm of 2, 4, . . . with the correlation expected from the model covariance matrix. Generally, only the measures for $\Delta m = 2$ and $\Delta m = 4$ are used. The second diagnostic is a measure of the total power as a function of m compared to the expected power as a function of m. In order to simplify the measure, the ratio of the observed to the estimated power is expanded in Legendre polynomials. Again, only the lowest order coefficients are generally used. For an exact definition, see Schou (1993a, cf. App. 5).

3.3 Description of the Fourier Tachometer Images

In this section the raw or slightly processed data will be described. In addition to showing both normal and atypical images, examples of some of the more subtle features will be shown. Most of these unusual features were found when trying to explain various problems in later stages of the analysis, but have been collected in this section, as many of them turned out to be irrelevant to the problems motivating the search.

Some of the bizarre effects have only been seen once or affects only one series of observations, while others exist in several datasets. Many datasets have not yet been looked at in detail, so other effects could still be present in those datasets.

3.3.1 Normal Images

Typical intensity, velocity and objective images are shown in figure 3.6. The velocity images are images output from the RDC code, and have thus had removed an approximation to the velocity introduced by the surface solar rotation (see the section on the RDC code). Some of the

velocity and intensity images have been centered while others have not; when important, this has been noted. The outer two pixels of the velocity images have been cut away in the RDC code.

Some of the effects mentioned in the following sections can easily be seen in the figures, while some are best seen directly on a monitor. Sometimes it is necessary to see a whole series of images as a movie, which is unfortunately not possible here.

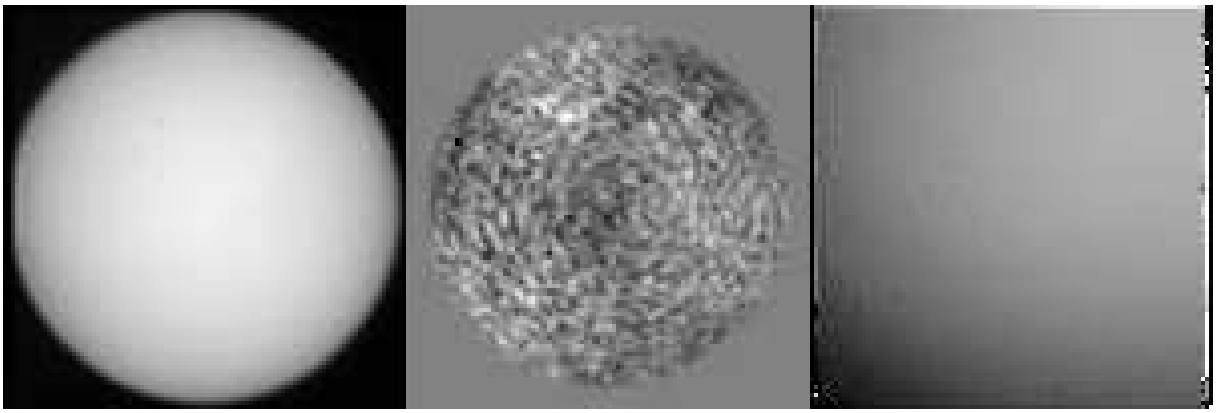

Figure 3.6: Examples of intensity, velocity, and objective images. The images are from the day 871108.

The most obvious defect on the intensity images (at least on good-quality plots, unlike that shown in figure 3.6) is probably the ghost which can be seen more clearly in figure 3.9. The ghost can be seen to rotate on the solar image as the camera rotates with respect to the optics. It was caused by a reflection off the pick-up mirror to the guider.

Sunspots are clearly visible on some of the intensity images, and the disturbing effect on the velocity images is apparent. Some of the 'sunspots' are actually bad pixels, which are especially obvious on movies of the intensity pictures.

Another defect in the intensity images is the cut on the left side, which can be seen in figures 3.6 and 3.8. This is due to a light sensitive shift register at the side of the detector, which had to be blocked off, so that the solar image would not fall on it. It was not possible to block this register without also blocking off some of the detector. Of course, it would have been possible to design the optics such that the solar image would have been small enough not to drift into this region. Occasionally the images also drift enough in the vertical direction to hit the top or bottom of the detector. There is some concern that this might affect the centering of the images. This problem will be discussed in another section.

An obvious flaw in some of the velocity images is the presence of roughly circular disturbances, shown in figure 3.7. These disturbances are caused by the A/D converter, which was not quite linear around major bit transitions. As the subimages taken every 1s are summed on the dedicated computer, which was part of the instrument, it has proved difficult to remove this effect after the fact. If the individual subimages were available, one could in principle have removed the problem, at least if the nonlinearities are time-invariant, by measuring and compen-

sating for the nonlinearity. Since the intensity decreases toward the solar limb and the intensity increases in the morning and decreases in the evening, these rings will typically move outwards in the morning and inwards in the evening. If the images are viewed as a movie, these effects are quite obvious, especially for the fall 1990 dataset. In the middle of the day when the intensity is roughly constant, the rings do not move. Also the main ring from the most significant bit transition was deliberately placed just outside the image when the Sun is at full intensity.

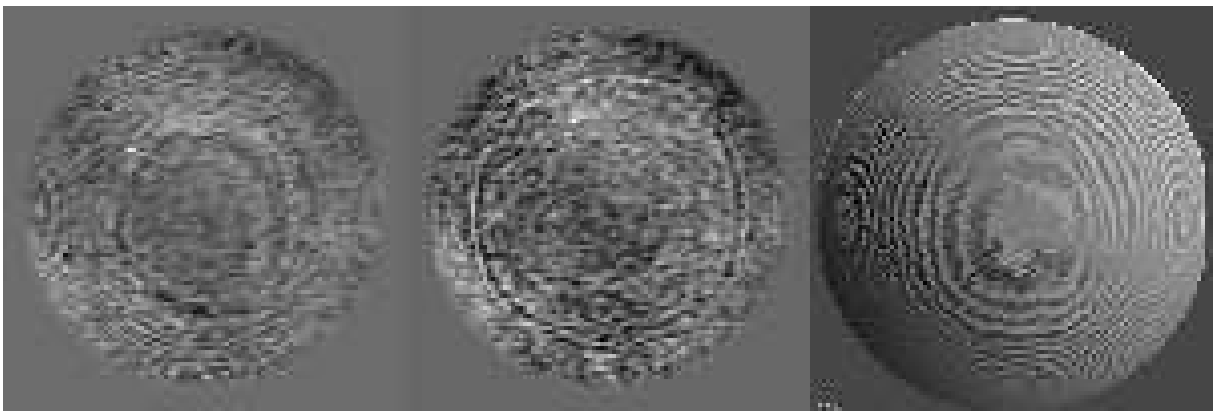

Figure 3.7: The left two images are two RDC velocity images from the day 871108, the rightmost image is a raw velocity image from the day 901207.

The effect shown on the rightmost image of figure 3.7 is probably also due to the problems with the A/D converter, but it is clear that there is a special problem with the later part of the FULL90F dataset, which appears to make it useless. A preliminary analysis seems to show that one of the bits in the A/D converter started to behave erratically.

The objective images are generally fairly featureless, but not quite uniform. An example is shown in figure 3.6. As explained later, they are not currently used, but were used for the aj89 dataset.

3.3.2 Statistical Description of the Images

Figure 3.8 shows 'low' and 'high' frequency intensity power, velocity power and covariance between the intensity and the velocity for approximately five hours of data from November 8, 1987. The low frequency plot includes frequencies $\nu < 833\mu Hz$, while the high frequency plots are for $\nu > 833\mu Hz$. The intensity power appears to have roughly the same appearance at low and high frequencies. Also the power looks very similar to the raw intensity images, except that the ghost has rotated during the run, which has smeared it out. Most of the intensity variance is probably due to the varying intensity of the Sun during the day.

If the intensities of the images are scaled such as to make total intensity constant for all the images, as shown in figure 3.9, the power no longer looks like the intensity images, but is instead dominated by the few bad pixels and the moving ghost.

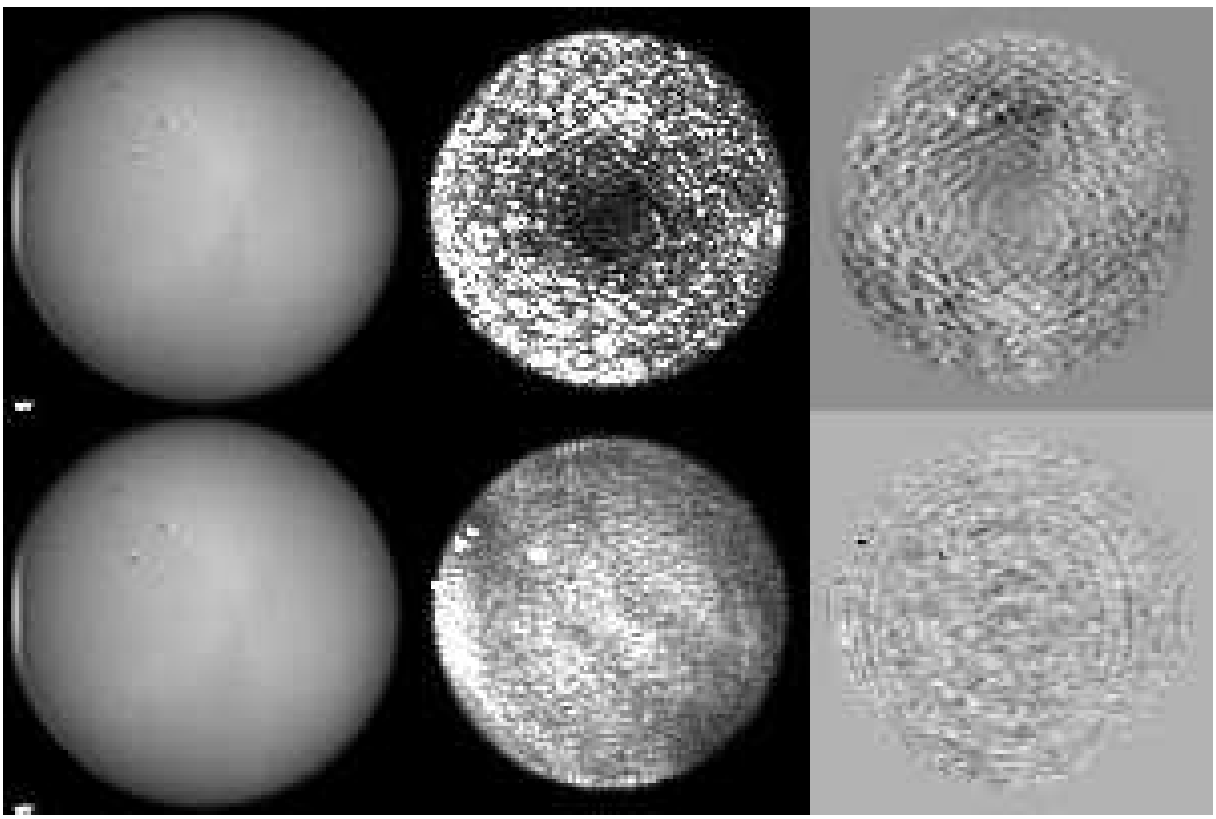

Figure 3.8: Images of the intensity power, the velocity power and the covariance between the intensity and the velocity, for five hours of data from November 8, 1987. The upper panel shows the low frequency part, while the lower panel shows the high frequency part.

The velocity power images are very different depending on the frequency range. The high frequency image is dominated by the p-modes and thus has most power towards the center, due to the velocity projection factor. The low frequency image is clearly dominated by another source, presumably supergranulation.

If a small square, in this case of size 40X40 pixels, is taken out of each image, a 3-D Fourier transform is performed, and the norm is taken and squared, it is possible to learn a number of different things.

First of all one can produce a $k-\omega$ diagram by averaging over the direction of the waves, that is by averaging together all points with the same wavenumber $k = \sqrt{k_x^2 + k_y^2}$. An example of such a diagram is shown in figure 3.10. Notice that as the radius of the solar images are $R \approx 48$ pixels, $l = 2\pi Rk \approx 300k$, with k in pixels^{-1} (as has been used here). As only the center of the image was used (in order not to have problems with foreshortening, and the line of sight velocity projection factor), individual ls are not isolated. In addition, only a few hours of data were used and the frequency resolution is thus very low. Also notice that the solar rotation has not been taken out, which leads to a small smearing, as does the neglected foreshortening.

It is also possible to collapse the 3-D power spectrum in frequency. (That is average over frequency.) An example of this is shown in figure 3.11. Only the frequencies inside of the

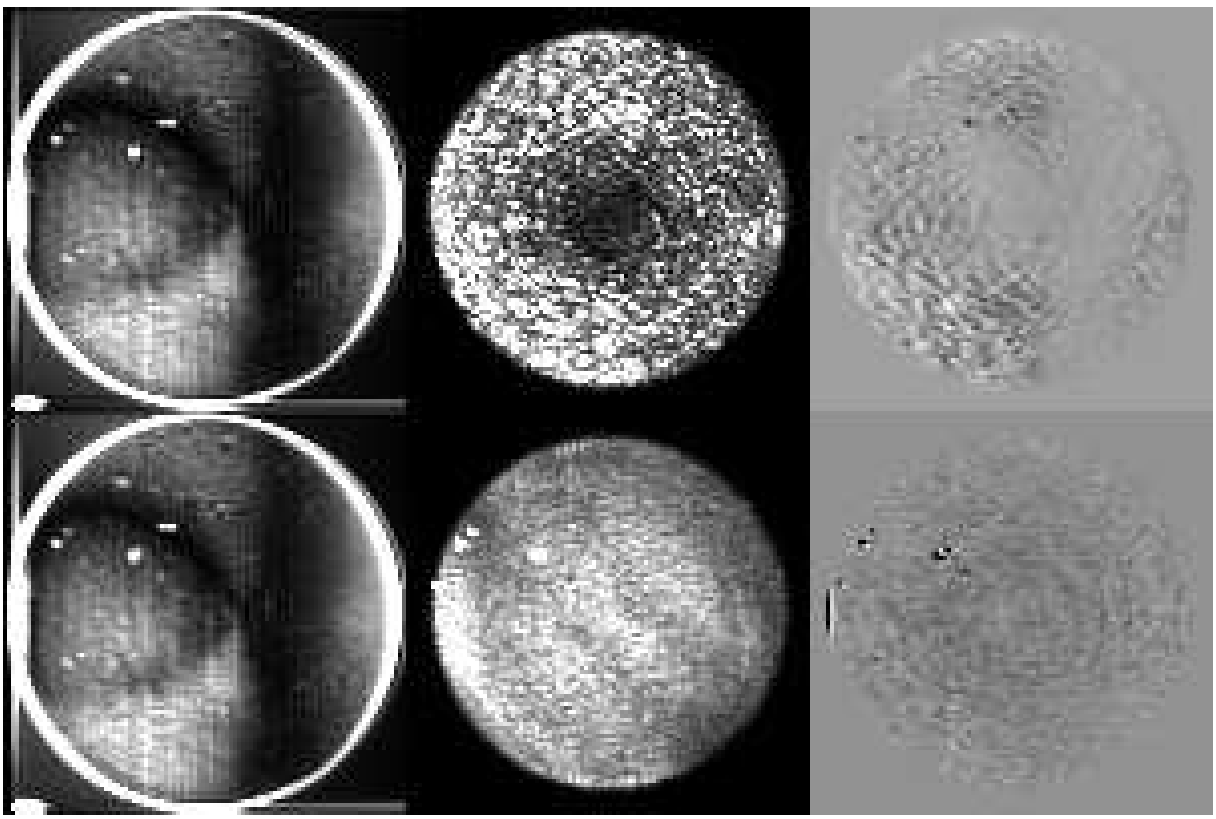

Figure 3.9: As figure 3.8, but after the intensity images were normalized to have the same total intensity.

p-mode band were used (in this case $2000\mu Hz < \nu < 4000\mu Hz$), cf., the discussion of the images of the velocity power. If it is *assumed* that the power is independent of wavenumber, this shows the combined Modulation Transfer Function (MTF) (within a multiplicative factor) of the atmosphere, the optics, the detector, and any modifications done to the images before this stage. The MTF is the square root of the ratio of power at a given wavenumber after the light has passed through the optics to the power before. The Fourier transform of the MTF is the Point Spread Function (PSF), which describes how a point on an ideal image of the Sun is spread out due to optical problems etc. Notice that the effects of the interpolations performed after this stage of the analysis have not been taken into account. For later use the MTF has been fitted to a Gaussian of the form (actually the fitting is done in the log of the two functions, using an unweighted least squares fit):

$$f(x,y) = e^{-\frac{x^2}{2\sigma_x^2} - \frac{y^2}{2\sigma_y^2}},$$

where σ_x and σ_y are measures of the widths in the x and y directions, which, if it assumed that the MTF is a Gaussian, can be converted to widths of the PSF. While the MTF is not always perfectly represented by a two-dimensional Gaussian (see for instance figure 3.13), it is generally not too far off, and it is a very convenient form to use, as the Fourier transform of a Gaussian is a Gaussian.

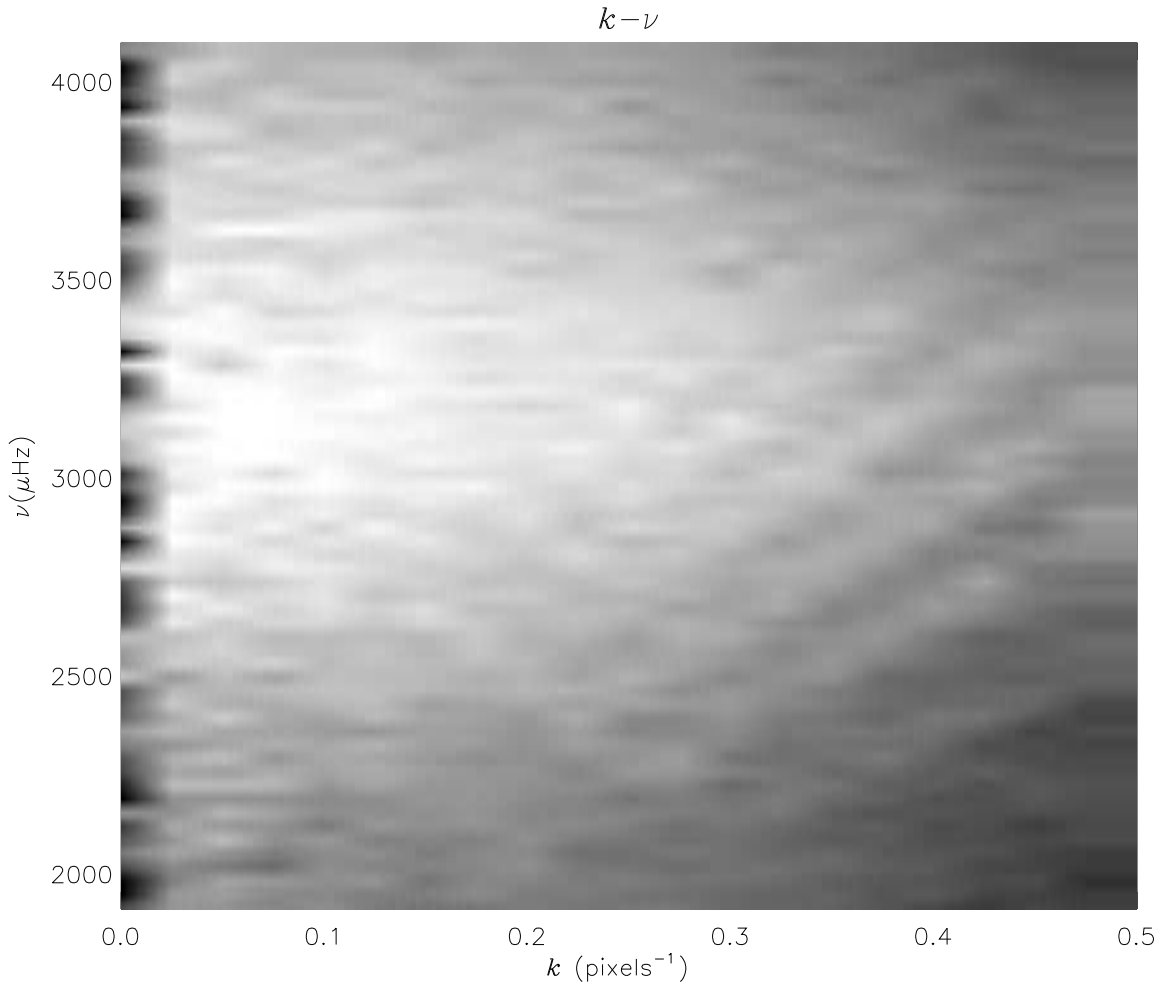

Figure 3.10: $k-\omega$ diagram for a small patch of size 40X40 pixels on the center of the Sun. 6.4 hours of data from the day 890408 were used. The true resolution in k is 0.025pixels^{-1}, the frequency resolution $\approx 35\mu Hz$.

It is evident that the MTF/PSF is significantly worse horizontally than vertically for this (and most other datasets). This is presumably due to a known horizontal charge leakage problem in the detector. The rest of the PSF is probably due to a combination of seeing (presumably a small effect, as the pixels are approximately 20 arcseconds across, which is much larger than the seeing), misfocusing/various optical defects, movement of the solar image during the exposures, the finite pixel size, and the interpolation used when centering the images. The PSF determined this way is not constant in time. There are large differences between different runs and minor differences between different days of a given run.

The differences between runs are presumably due to the changes made to the instrument and light feed, in particular the focusing was redone for each run. Possibly the degradation of the optics and the detector may be a problem too.

The differences between different days from the same run could be due to change in seeing etc. A likely effect is also the fact that the direction of the image drift could change from day to

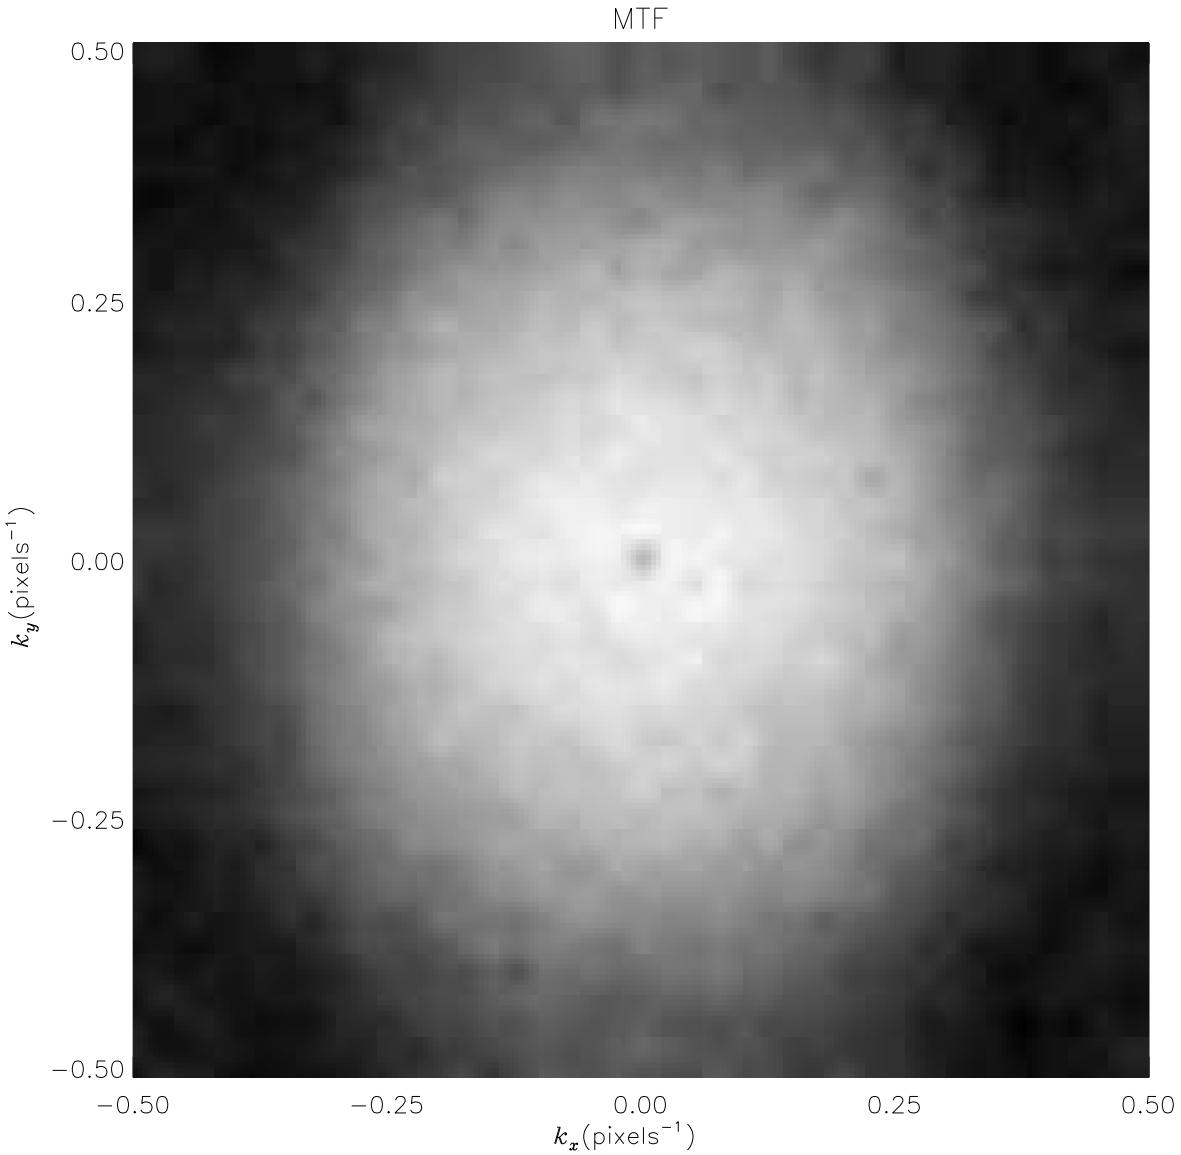

Figure 3.11: The power in the p-mode band as a function of k_x and k_y. 6.4 hours of data from the day 890408 were used. Only the power in the range $2000\mu Hz \leq \nu \leq 4000\mu Hz$ was used. The 'hole' in the center is due to the the removal of a polynomial fit to the images in the RDC code.

day leading to differences in the contribution of the centering interpolation to the PSF. If the images are only centered by a small fraction of a pixel in, say, the x-direction, the interpolation will only degrade the PSF slightly in this direction. As discussed elsewhere, this short-time variation can lead to problems in the subsequent analysis.

3.3.3 The Image Rotation Problem

As mentioned earlier the original RDC code rotated the velocity images back and forth, when subtracting the objective images. It was decided that this was both unnecessary and in fact

degraded the images. An example of the MTF before and after these rotations were taken out is shown in figure 3.12 and 3.13. The improvement is quite substantial. Before and after $k-\omega$ diagrams are shown in figure 3.14. Again it is evident that the signal to noise ratio at high k (l) has improved substantially. The width of the PSF (for the day 890408 used for figure 3.12) changed from $\sigma_x = 0.819$ to $\sigma_x = 0.644$ pixels, while the change in the y-direction was from $\sigma_y = 0.708$ to $\sigma_y = 0.395$ pixels.

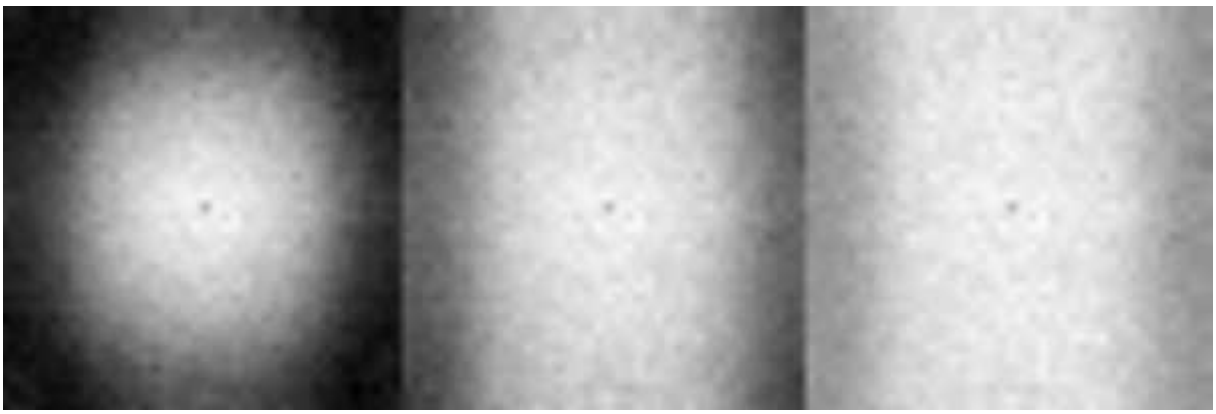

Figure 3.12: As figure 3.11. The left image shows the logarithm of the power as a function of k_x and k_y when the images were both rotated and centered, the middle image when they were only centered and the right image when the images were neither centered nor rotated. 6.4 hours from the day 890408 were used. The left plot was used for figure 3.11.

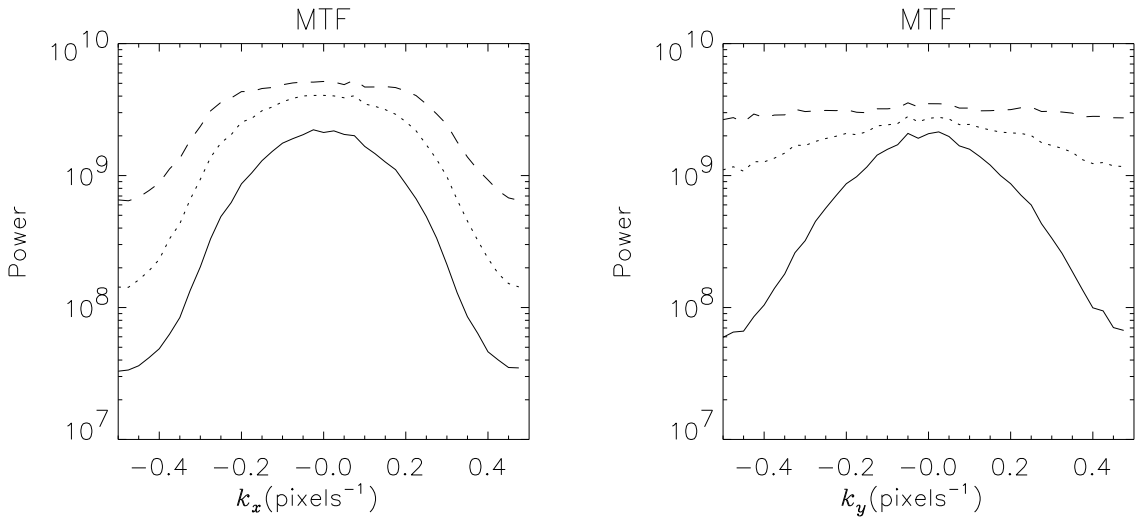

Figure 3.13: The MTF's from figure 3.12 collapsed onto k_x and k_y. Solid lines are centered and rotated, dotted centered but not rotated, and dashed neither centered nor rotated.

3.3.4 The 1989 Noise Problem

When the spring 1989 run was first analyzed, it was realised that the noise was very high and that the real and imaginary parts of the time-series correlated strongly, which they should not

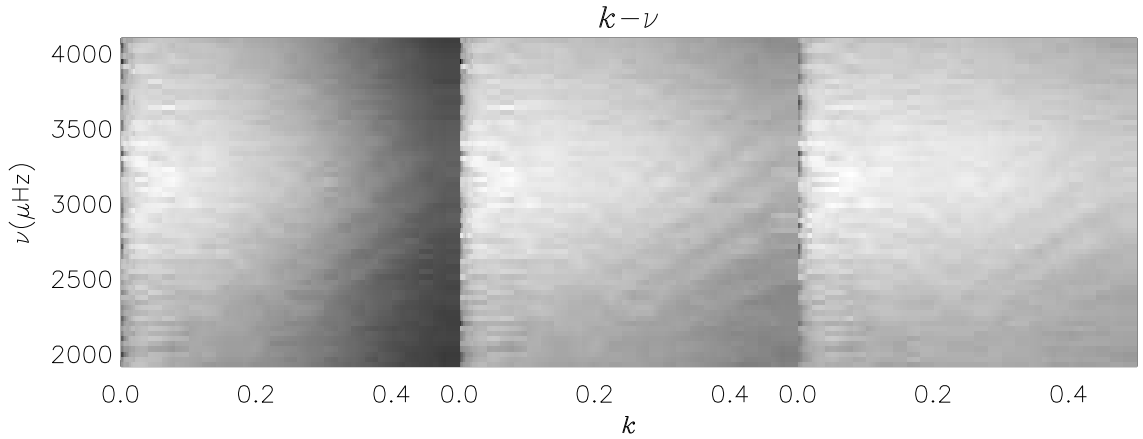

Figure 3.14: $k-\omega$ diagrams corresponding to the images in figure 3.12. The grayscale is the same on the 3 subplots. The center plot was used for figure 3.10.

do according to the parity rules (see eg., Schou and Brown 1993, cf. App. 4). An analysis of the velocity images revealed that there was a very strong noise source on the left side of the images. An example of the velocity power distribution for this run is shown in figure 3.15. If the velocity power image is collapsed in the vertical direction, as is also shown in figure 3.15, it is obvious that the noise is dominated by this problem. In view of this it was decided to make a symmetrical vertical cut at $x = \pm 0.6R$ in order to reduce the noise. Making this cut obviously reduces the effective area observed and hence makes the crosstalks larger, which needs to be taken into account, but as the reduction in noise is quite substantial, it should be an advantage to cut the noise away.

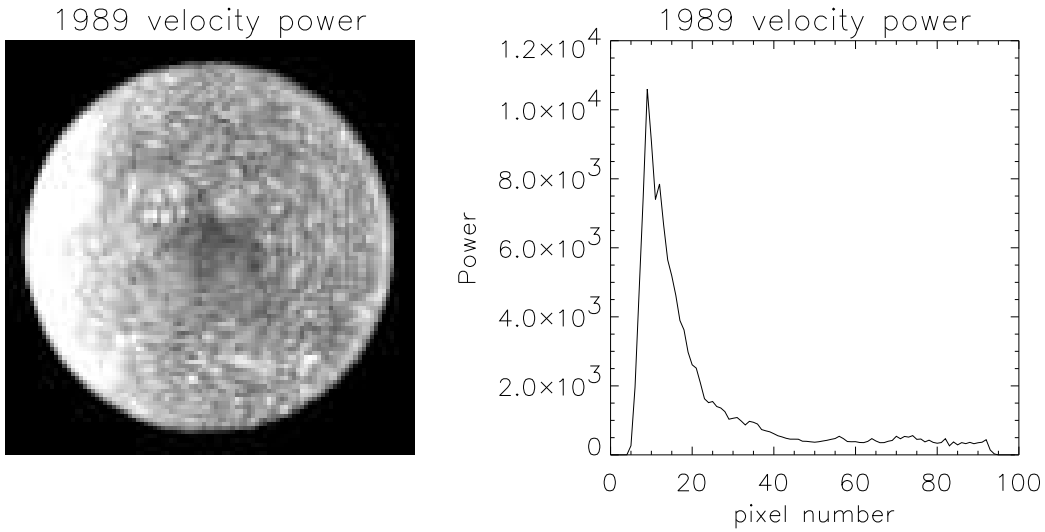

Figure 3.15: The *logarithm* of the low frequency noise for the day 890408 as a grayscale plot and the result of summing each column. The data for the left hand plot in figures 3.12 and 3.14 were used.

The cause of this noise is not known. It appears that the objective images have a slight but varying gradient in the same area as the increased noise. As a matter of fact, the reason for not subtracting the objective images is that they contributed strongly to the low frequency noise. As the objective images are only likely to affect low frequencies, removing the subtraction probably has little effect on the noise in the p-mode band.

The noise seems not to rotate during the day as would be expected from a problem in the optics. The most likely source seems to be the detector (including the electronics) or some component rotating with the detector. It could for example, be related to the light blocking device, which was removed and put in again before each run. It is, however, not clear how the blocking can create noise problems.

3.3.5 The Centering Problem

As mentioned earlier, any interpolation in the images will degrade the PSF. In particular, the fact that the centering and the interpolation to the latitude/longitude grid are done separately is somewhat unfortunate. If the centering is not performed (which is of course slightly problematic, but not too bad as the images are generally drifting slowly) and one calculates the PSF, one obtains $\sigma_x = 0.504$ and $\sigma_y = 0.201$, for the case previously considered, which is again a substantial improvement. The MTF for this case was also shown in figures 3.12 and 3.13. It is possible to combine the centering and the interpolation, and this is certainly a place where the current reduction could be improved.

The PSF determined without the centering is probably a good measure of the part coming from the seeing, optics, and detector. The substantial difference between the x and y directions mentioned earlier is most likely due to the charge leakage in the detector. The charge leakage, which is not likely to be well represented by a Gaussian PSF, could also be the cause for the non-Gaussian MTF in figure 3.13. It is interesting, however that the FULL90F run has a wider PSF in the y-direction than in the x-direction.

3.3.6 The Even/Odd Problem

On the low frequency velocity power images a striking feature is the difference between the even and odd columns. This is particularly clear on the noncentered images, an example of which is figure 3.16. Figure 3.16 also shows the total power in each column and it is obvious that every other column is very noisy compared to the columns in between. Unfortunately, the noise is not correlated between the pixels in a given column, which would have enabled one to remove most of it by subtracting a column average from each column (which would be a problem for sectoral modes) or estimating the offset from the unused part of the image.

The cause of this behavior is related to the fact that the even and odd columns are brought to different amplifiers, one of which was apparently much more noisy than the other. It is interesting that the noise seems to be larger where the intensity is low.

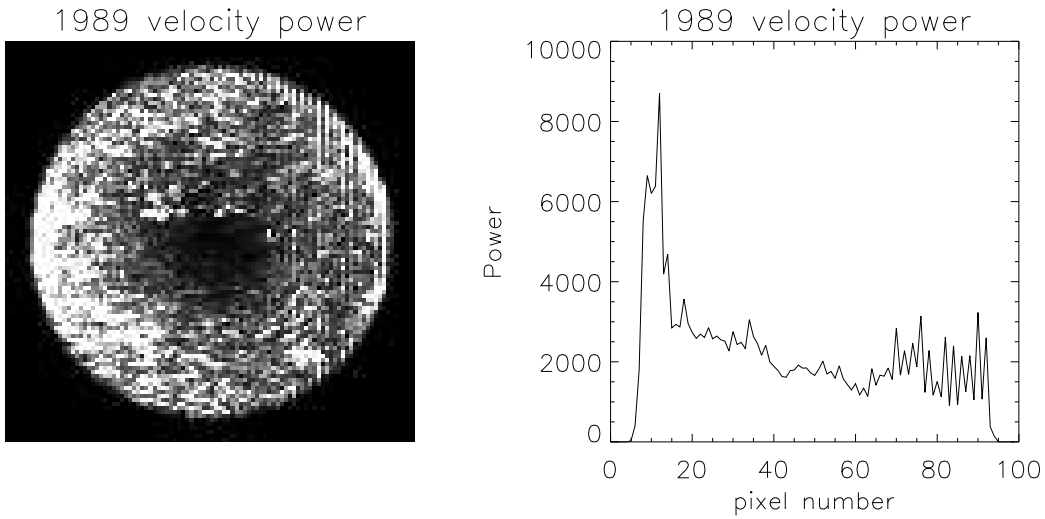

Figure 3.16: As figure 3.15, but using the power instead of the logarithm and for the case used for the right-hand plots of figures 3.12 and 3.14, that is, using uncentered images.

3.4 Reduction of the aj89 Dataset

Of the observing runs reduced until now, most effort has been put into the reduction of the data from aj89/spring 89/FULL89 run. While this is not the longest or best run, it was for historical reasons picked as the first major run to analyze.

The basic statistics for this run were given in the section on the instrument.

In this section some of the details of the reduction of this run will be presented. Given that this set was the first reduced, some of the problems are unlikely to affect other datasets. Also, as this run has been given most attention, the descriptions will be more detailed than for other runs.

3.4.1 Initial Reduction

The effective P-angle for this run was determined to be 9.24° from observations of sunspots. The accuracy is believed to be of the order 0.5°.

The PSF was determined to have $\sigma_x = 0.016R$ and $\sigma_y = 0.0144R$ using an early version of the PSF finding program and the day 890408, which should be a fairly good day. The PSF has not been corrected for the interpolations done in the YLM code for the purpose of calculating the crosstalks. The significance of this will be discussed later.

The RDC code used for this run was the old version with objective image removal and rotations of the velocity images. Given that most of the analysis had been done before these problems were discovered, it was decided not to redo the reduction until a later time in case other problems would turn up. As mentioned earlier, a noise problem towards the one side of the solar images was found and the sides of the solar images were cut off. The cutoff was more or less arbitrarily put at 0.6R as a compromise between decreasing the noise and removing too much of the signal.

The removal of bad images was mainly done manually. As some of the data were later reanalyzed for the tests described elsewhere, a small number of images included in the main run have later been found to be of marginal quality. Since the number of bad images found later is small, nothing has been done to remove those images at a later time, which might otherwise have decreased the noise slightly. The number of hours observed each day is shown in figure 3.17.

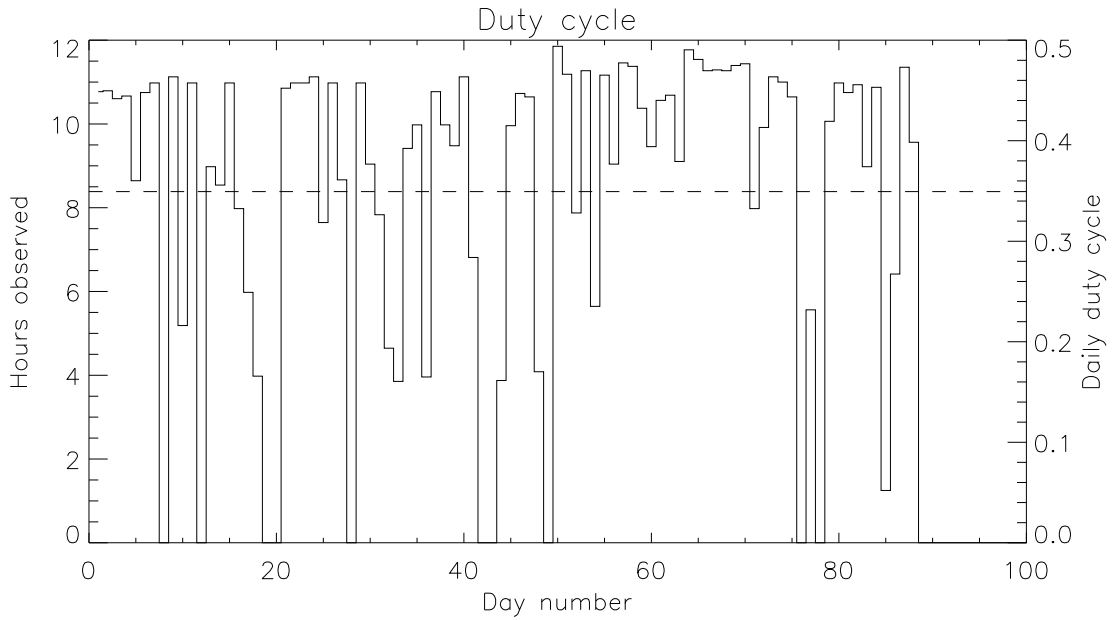

Figure 3.17: The number of hours observed each day (left-hand axis)/daily duty cycle (right-hand axis). The average duty cycle has been indicated by a dashed line.

The YLM code is the standard version, with the already mentioned cut and standard apodization. No changes were made in the RTS code for this run.

As the spatial Nyquist frequency at the center of the images is around $l = R\pi \approx 48\pi \approx 151$, where R is the radius of the solar image in pixels, it was decided only to fit up to $l = 120$. Also the imperfect PSF means that the power decreases rapidly above this point. Notice that if the PSF were good (that is narrow) some spatial aliasing should occur. Due to the foreshortening this will not be in the form of nice ridges going 'backwards', but more probably as very fuzzy ridges, which could potentially be a problem at ls close to the Nyquist frequency.

3.4.2 Time-Series Analysis

As the time difference between the first and last Y_l^m coefficient was 100720 timesteps of 75s, the Fourier transforms were done with a length of $100800 = 2^6 3^2 5^2 7$, which is the first higher number that is a product of small primes. Given the sampling interval, the Nyquist frequency is $0.5 \times (75s)^{-1} \approx 6667 \mu Hz$. The total length of the Fourier transform is $100800 \times 75s = 7560000s$, resulting in a frequency spacing between points in the discrete transform of $\approx 0.132 \mu Hz$. The

effective duty cycle (the ratio of the number of good images to the length of the Fourier transform) was 34.91%. The power spectrum of the window function is shown in the section on the tests performed (figure 3.31).

As already mentioned, the crosstalks were calculated using the values estimated from the day 890408, without correcting for additional degradation caused by the interpolation, but taking into account the apodization and cut.

The noise covariance matrix was determined from the time-series as a combination of low and high frequency noise.

For the fitting, all modes with $|\Delta l| \leq 6$ were used. This should be more than adequate for all modes. The frequency limits were set to $1500\mu Hz \leq \nu \leq 4000\mu Hz$. The lower limit should include all modes with a sufficient S/N, while the high limit corresponds roughly to when the linewidths become comparable to the spacing of the temporal sidelobes ($11.57\mu Hz$) and to the size of the window used for the fit. Instead of trying to fit for the modes above $4000\mu Hz$ (which is perfectly possible), we decided not to try, as the results would likely be unreliable.

For $l \geq 9$ a_1 through a_6 were fitted for. The type of a-coefficients used are similar to those of Ritzwoller and Lavely (1991), but modified to be asymptotically equal to the standard coefficients in the limit of high l (see Schou, Christensen-Dalsgaard, and Thompson 1993, cf. App. 3, for an exact definition). For $l < 9$ a smaller number of a-coefficients were used, such as to make the fits more stable. It has not been possible to fit $l \leq 4$, presumably due to the increased noise and the subtraction of the polynomial fit to velocity images, which removes any low spatial frequency components.

The fitting window was set to five times the FWHM of the mode, but kept within the range 40-100 frequency bins, corresponding to $5.29\mu Hz$-$13.23\mu Hz$. The high limit to the window size was set such as to restrict the number of ms to the range where the estimated covariance matrix used in the analysis stays reasonably nonsingular.

The level of other modes was set to follow that of the target mode for the final iterations, which makes the fits considerably more stable.

As good values for the various parameters were not known in advance, a fit to results published by Libbrecht (1988) was used. The mode amplitude presents a special problem, as the inaccurate PSF leads to a decrease in the amplitudes for high ls. Libbrechts values were then used as initial estimates, and an attempt was made to fit as many modes as possible. The new parameters determined this way were then smoothed and used as new initial estimates, extrapolating to higher ls. Repeating this process, it was possible to eventually fit modes with $5 \leq l \leq 120$.

After this the process was iterated until all parameters converged. To generate new initial estimates, the fitted parameters from the previous run were weeded and occasionally parameters from previous iterations were added to fill holes. These parameters were then gap filled using adjacent ls and smoothed slightly in l.

The mode frequencies and the a_1s turned out to be particularly difficult to converge for the procedure used, as the interference from modes with $|\Delta l| = 1$ is large in the area of the 'bump'.

This, combined with the fact that a_1 deviated so much from the constant value assumed in the initial estimates, made a fairly large number of iterations necessary. In some cases the iteration was speeded up by extrapolating the fitted parameters, as a function of iteration number, using the behavior observed for previous iterations.

3.4.3 Results

An l-ν diagram of the modes fitted in the last iteration is shown in figure 3.18. The rejection criteria used were the following:

1. Failure to converge before exceeding the limits on the number of various iteration steps
2. Non positive definite estimated covariance matrix
3. Test(0) outside the interval [0.99,1.01] (see below)
4. $\sigma(\log(P)) > 0.5$, where P is the mode power
5. $\log_e(P_{\rm other})$ outside the interval [-1,1]

Note that as the level $P_{\rm other}$ (see Schou, 1993a, cf. App. 5) of modes other than the target mode was set to follow that of the target mode, 5. is essentially equal to requiring that the mode power does not change by a large factor during the iteration. Test(0) used in 3. is essentially the ratio of the power in the real and model spectra, and is the first diagnostic discussed in Schou (1993a, cf. App. 5).

After applying these criteria, a total of 1218 modes remained.

These criteria are obviously not enough to ensure that the parameters found are those corresponding to the relevant minimum of the merit function, or that the solution is even close to a minimum. Looking at plots of various parameters, in particular a_1, shows that some fits have failed. For most modes the parameters changed very little between the previous iteration and the last, suggesting that the previous iteration was already almost converged. An obvious further rejection criteria is thus that some parameter changed by a large amount between the two last iterations. To be specific, all modes where ν, A, w or an a-coefficient changed by more than 0.5σ between the last two iterations, were rejected. It is improtant to note that this criteria was only applied after the final iteration, not after the previous iterations.

These rejection criteria leave a number of isolated modes at high frequency and high l. As it is unlikely that these modes are any better than the surrounding rejected modes, they were also rejected. Finally, a few modes with implausible parameters were removed. The resulting set of modes is shown in figure 3.18.

The resulting set of 1100 modes contains very few modes with implausible parameters, but a few are left, most of which have high frequencies. The results for various parameters are shown in figures 3.19 to 3.21.

A more proper way to find bad fits would be to check the sensitivity of the output parameters to perturbations of the input parameters. Unfortunately, this is a computationally very expensive thing to do, as it requires the whole time-series analysis procedure to be repeated. Also, it is difficult to determine the proper magnitude of the perturbation - with a sufficiently

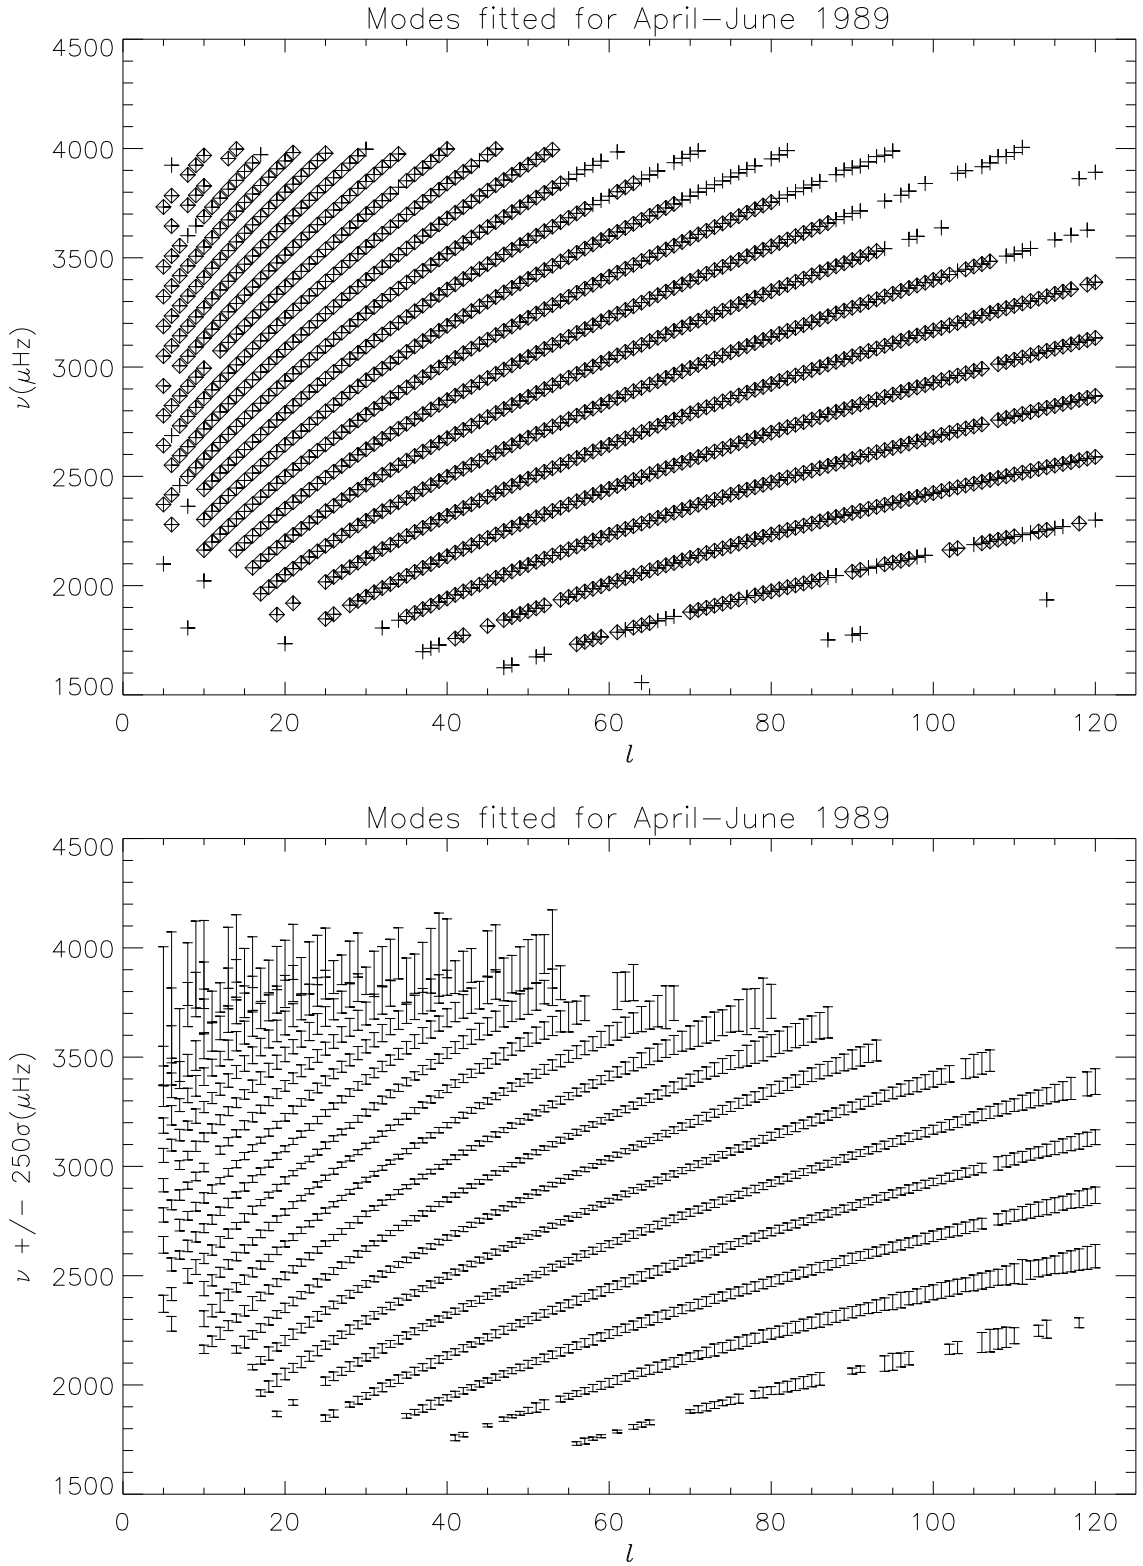

Figure 3.18: The modes fitted in the aj89 dataset. + indicates the original set of modes, diamonds the weeded set of modes. 250σ error bars were used on the bottom plot for clarity.

large perturbation, even the best method will fail. Some tests of the sensitivity to perturbations are discussed in the section on the tests performed.

While the modes at low and medium l and high frequency were formally fitted, there are good reasons to be careful when using them. First of all, the linewidths are comparable to the spacing between the temporal sidelobes ($11.57\mu Hz$), making it difficult to separate them from the sidelobes and to get reliable estimated errors. The linewidths are also comparable to the size of the window used for the fitting, leading to vastly increased random errors as the information outside the window is ignored. In any case the errors are generally so large on the various parameters for these modes that they carry very little weight, when, eg., they are used in inversions.

3.4.4 Discussion

As is obvious from figure 3.19, both the amplitudes and linewidths depend on both ν and l.

The dependence of the amplitude on l is probably not real, and is somewhat unexpected given that the PSF was effectively determined such as to make the power independent of k (and hence l). The explanation is presumably that the PSF does not include the contribution from the interpolation performed in the YLM code. This point will be discussed in more detail later.

The dependence of A on ν is close to that found by others (eg., Jefferies et al. 1991).

The dependence of the linewidth on the frequency is also qualitatively close to that found by Jefferies and others. The absolute level and the dependence on l is are somewhat different, however. In particular, it appears that the linewidths determined here are generally larger than those determined by Jefferies et al. (1991).

The a-coefficients are also very interesting, in particular a_1 shows considerable structure for $\nu/L < 50\mu Hz$. This structure is what has become known as the 'bump', and a whole section (3.7) has been devoted to it, as has a major fraction of section 3.6. The problem with this bump is that the position of it corresponds to the place where the temporal sidelobes of the neighboring ls interfere with the target mode. a_3 on the other hand shows no sign of a bump. a_5 shows some structure around the bump, but of a considerably lower magnitude. The even a-coefficients are discussed in more detail in section 3.8.

The dependence of the background noise level on l is interesting. In particular, there is a strong odd/even effect. Given that only even ms were used in the analysis, all the masks are even functions of the latitude for even ls and odd functions for odd ls. If the noise close to the equator is not uniform, this will lead to an odd/even effect in the noise, and is probably the cause for this curious phenomena. Except for this effect, the noise dependence on l appears to be dominated by the decrease with l due to the bad PSF, and possibly other effects. Also, there is a sharp increase at low l, probably due to some instrumental effect, changes in atmospheric transparency, or similar effects.

The noise dependence on frequency, with the maximum around the peak of the p-mode power, would seem to indicate that a major contribution to the noise in the center of the p-mode band is unresolved modes. Given the imperfect PSF and the resulting not totally negligible

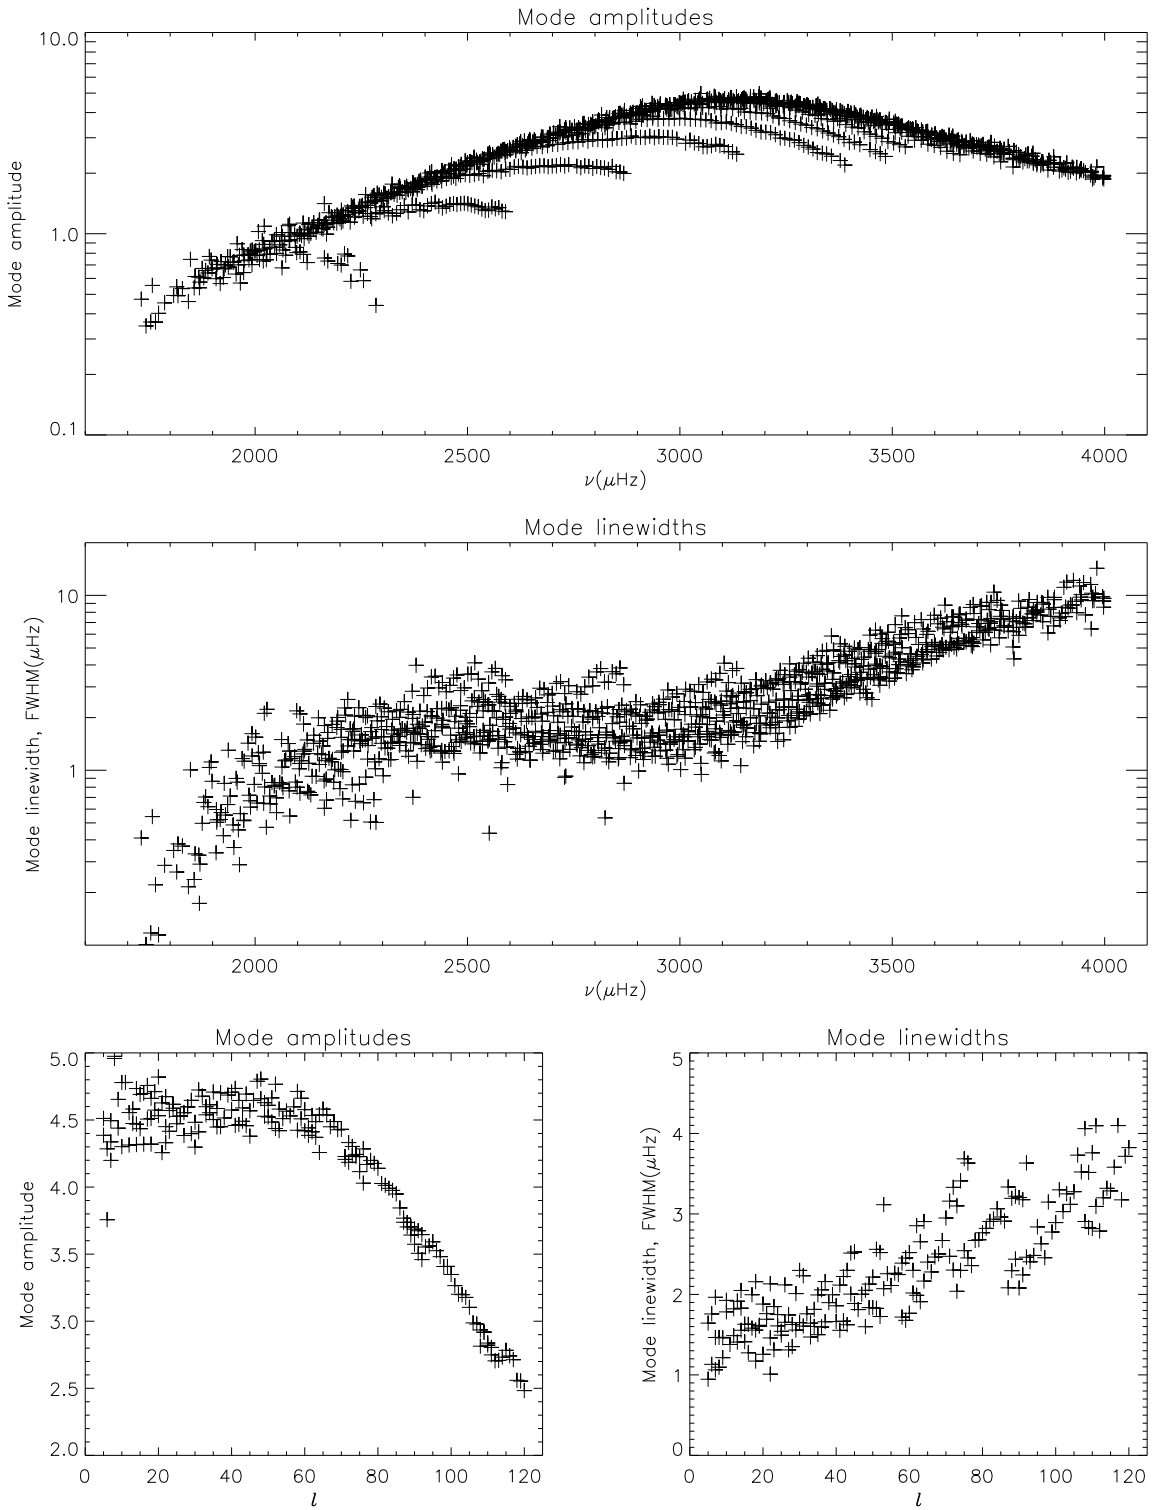

Figure 3.19: Mode amplitudes and linewidths as a function of frequency and l. For the plot of amplitude versus l only modes with $3000\mu Hz \leq \nu \leq 3300\mu Hz$ were used to reduce the effect of the dependency on frequency. For the linewidth plot only modes with $2500\mu Hz \leq \nu \leq 3000\mu Hz$ were used.

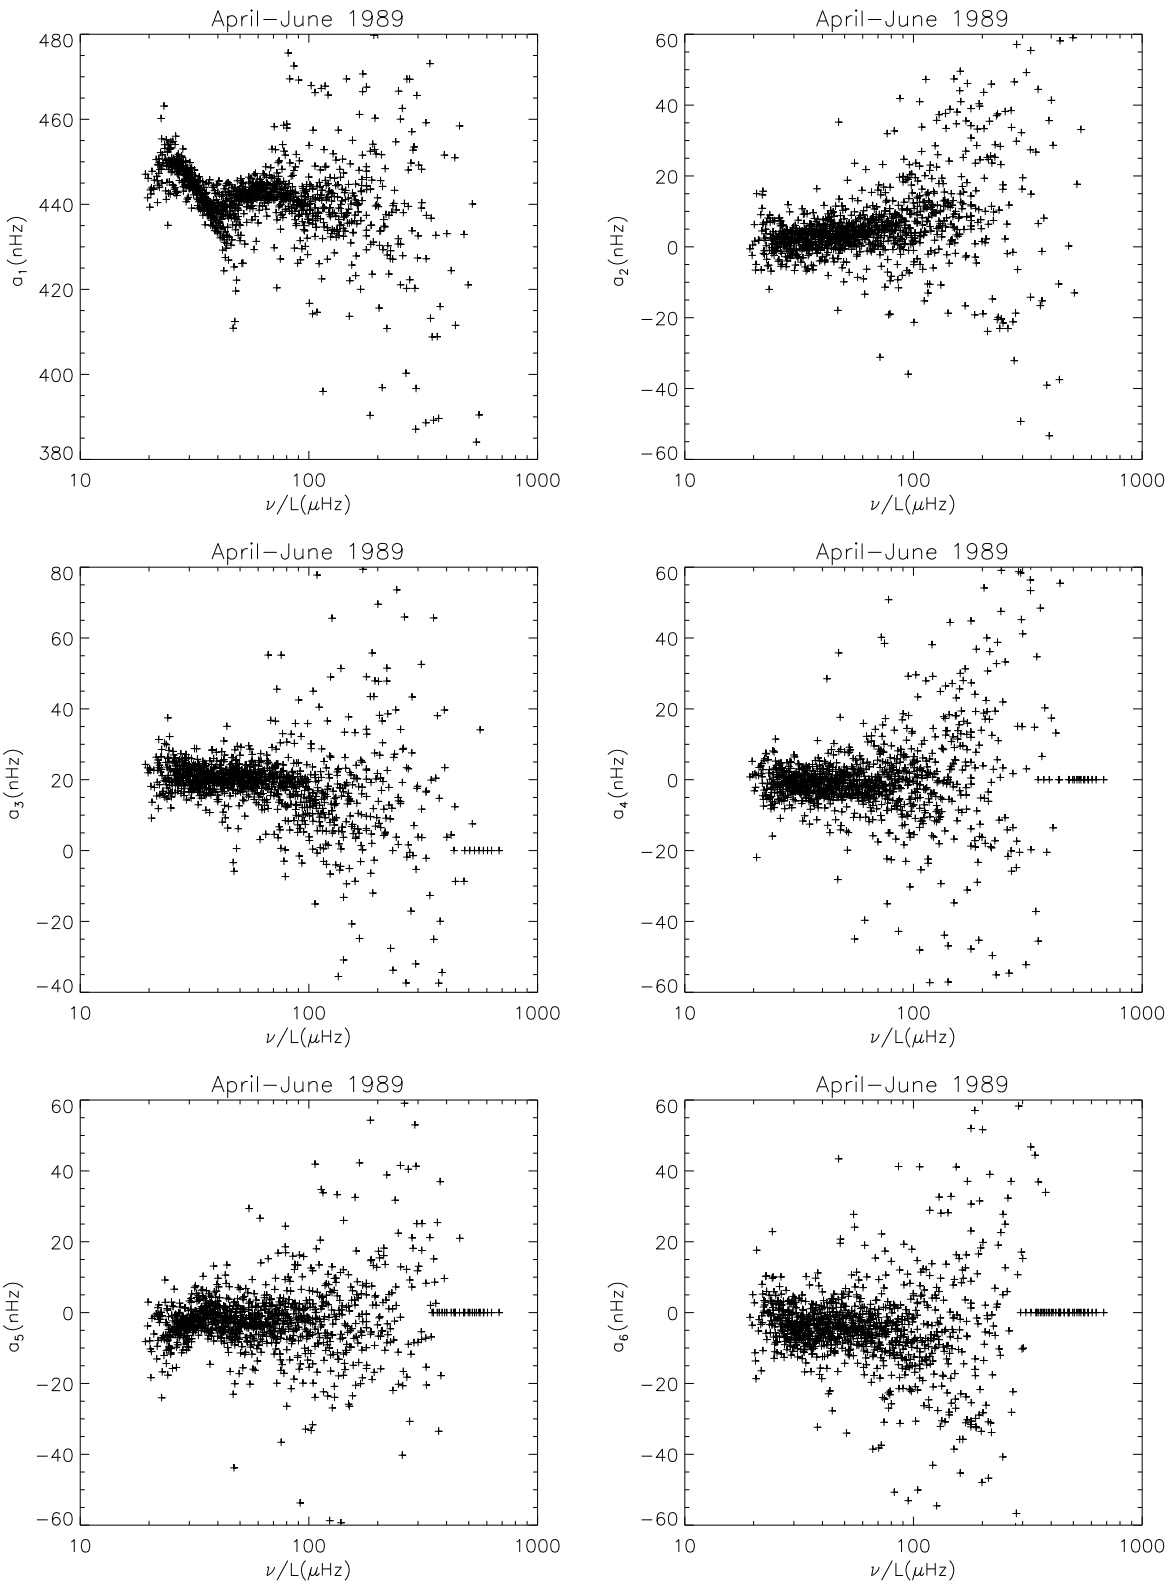

Figure 3.20: a-coefficients as a function of ν/L. The modes with $a_i \equiv 0$ for high ν/L are modes with $l \leq 8$, where not all a-coefficients were fitted for.

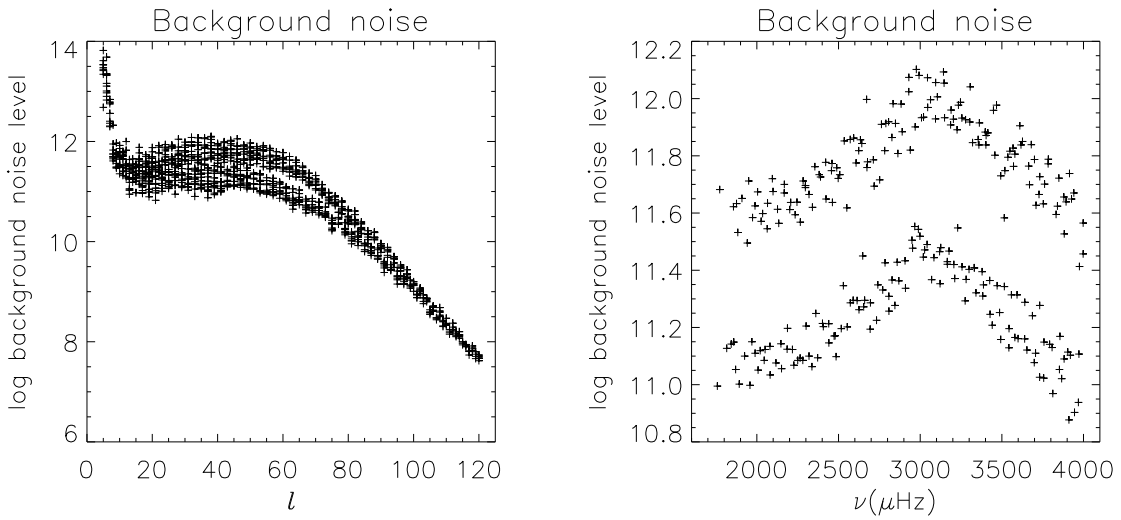

Figure 3.21: The background noise as a function of l (left panel) and ν (right panel). For the plot versus ν only $30 \leq l \leq 50$ was used to reduce the effect of the l dependency.

crosstalk from modes with $|\Delta l| > 6$, this is not surprising. Also there may be a contribution from the modes taken into account (those with $|\Delta l| \leq 6$) as the temporal sidelobes are only taken into account out to a finite distance in frequency from the central peak. Incorrect sidelobes, due to varying PSF or other effects, are also potential problems.

The expected correlations between a few of the parameters, that is

$$Cor_{ij} = \frac{Cov_{ij}}{\sqrt{Cov_{ii}Cov_{jj}}},$$

where Cov is the covariance estimated from the fit (see Schou 1993a, cf. App. 5), have been plotted in 3.22.

As can be seen from figure 3.22, the correlations are generally fairly moderate (less than, say, 0.5). However, the power (amplitude) and the noise level are for some modes strongly negatively correlated, meaning that their sum is well determined. This correlation is basically due to the fact that the total power in the fitting window, which is generally well determined, is the sum of the contribution for the mode and the contribution from the noise. That the correlations are such strong functions of the frequency is probably due to the dependence of the mode power (and hence the S/N) on the frequency.

It turns out that if the power of the modes other than the target mode (see Schou 1993a, cf. App. 5) are fitted for, the correlations between this power and that of the mode power become extremely strong and negative (worse than -0.95), where the temporal sidelobes from the neighboring l cross the target mode (at $\nu/L \approx 31\mu Hz$). This is due to the fact that the sidelobes and the target mode look alike when the fitting window is small enough that the temporal sidelobes of the target mode (and the main peaks of the neighboring ls) are not included. If the fitting window is large enough, the two contributions are easy to distinguish, as the ratios of the powers in the main and side lobes are different.

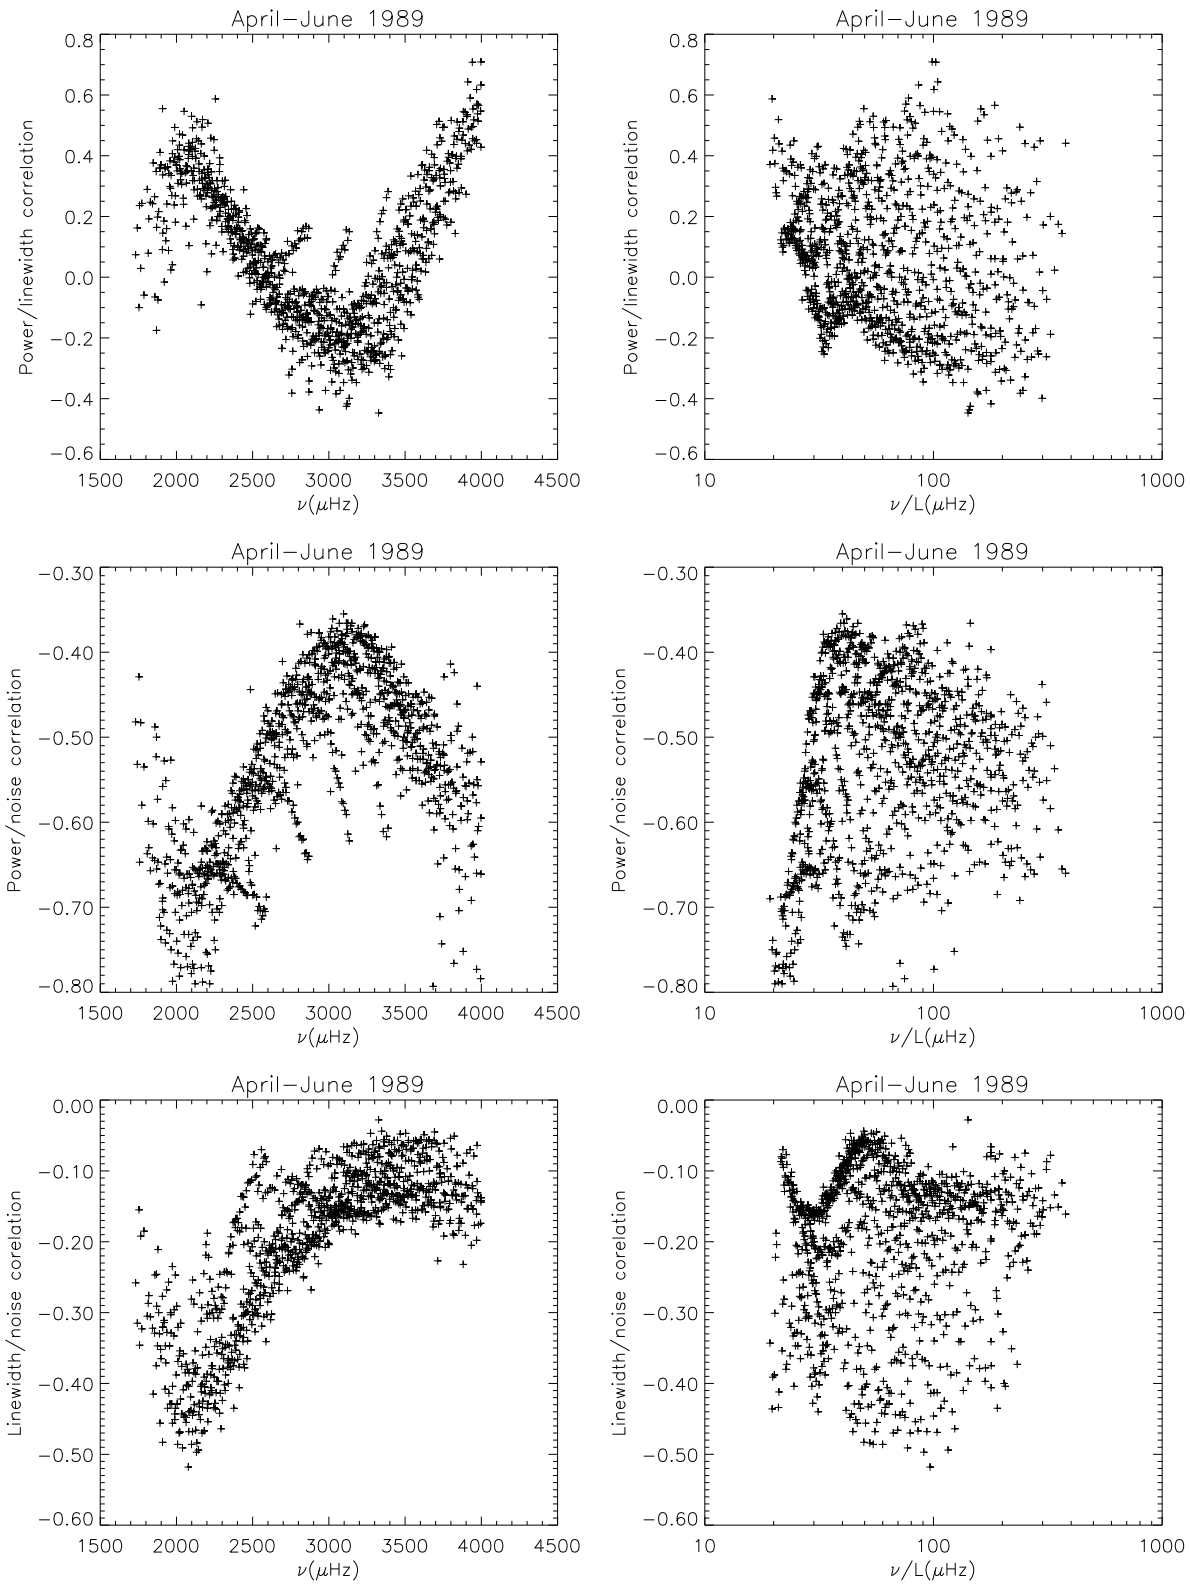

Figure 3.22: The estimated correlations between various parameters as a function of ν and ν/L.

The values of a few of the diagnostics are plotted in figure 3.23. As discussed in Schou (1993a, cf. App. 5), the first of the diagnostics is generally fairly difficult to interpret. The second diagnostic, on the other hand, clearly indicates that the PSF has been incorrectly estimated. As discussed before in connection with the amplitude dependence on l (note the striking similarity between the plots of the amplitude as a function of l in 3.19 and the second diagnostic in 3.23), the likely cause for this is the neglect of the degradation of the PSF caused by the interpolation in the YLM code. It is also interesting to note the similarity to the plot of the background noise in figure 3.21.

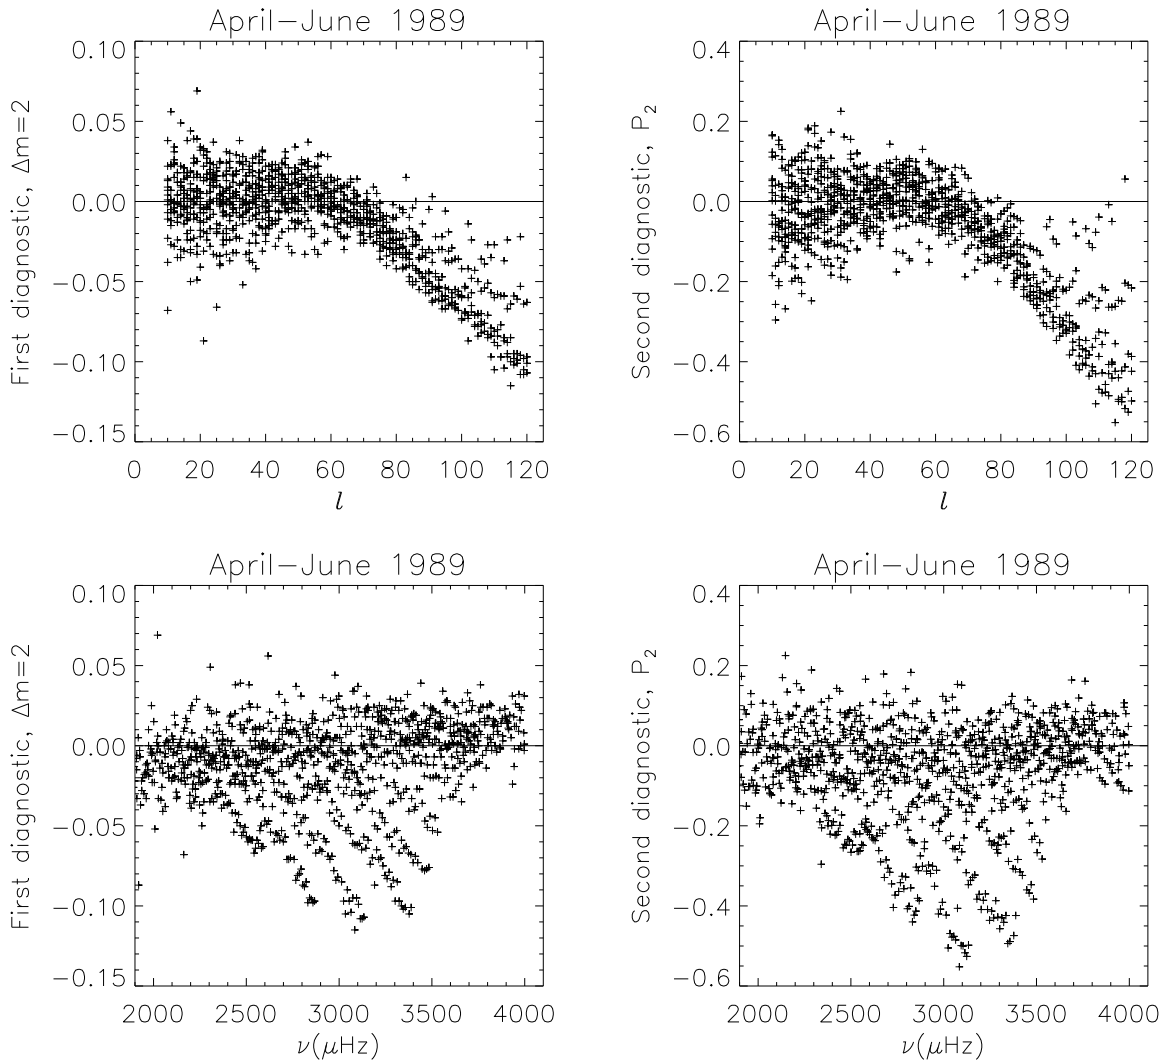

Figure 3.23: The values of two of the diagnostics. Only modes with $10 \leq l \leq 120$ are plotted.

The frequency dependence of both the diagnostics is fairly easy to explain. The problems with the PSF only affect the contribution to the covariance matrix from the modes whereas the estimate of the noise covariance matrix is fairly good (as it is derived from the actual time-

series). Hence, the error will be an increasing function of the ratio of the mode power to the noise, which, as the noise is fairly constant, mainly depends on the mode power.

It is clear that several improvements could be made in the analysis, both in the earlier steps and in the final part of the analysis that I have been most concerned with.

In the initial part of the analysis several minor and a few major improvements could be made. First of all, the noise could presumably be reduced by interpolating over the bad pixels and being more careful with the rejection of bad images. The improvement of the final results from these two improvements is likely to be rather small, as the contributions to the total noise are small. Also it is potentially possible to remove some of the contributions from active regions.

A major improvement in the analysis of this dataset would be to reanalyze it without the image rotations, as has been discussed earlier. This should improve the PSF substantially, which would improve the S/N considerably, especially at high l, as shown in the section on the description of the images.

It would also be possible to combine the centering and the interpolation on the uniform net in latitude and longitude, again improving the PSF. While this should be fairly straightforward to implement, it does mean that the images output by the RDC code would move around during the day. It would also create problems for the part of the apodization done in the RDC code.

In the final analysis code several things could be improved. As already mentioned, the PSF is clearly worse than assumed and probably non-Gaussian. Improving this should lead to smaller random errors, more confidence in some of the results, more converged modes, and improved diagnostics. It is not quite clear how to go about doing this in a more practical sense. In particular, it is difficult to find the correct PSF. Using the measured one and trying to model the additional degradation caused by the interpolation is certainly an option, but judging from comparisons of centered/noncentered images and rotated/nonrotated images (both effects should behave similarly), it has become clear that the degradation is difficult to model properly. Doing a crude estimation is likely to improve things considerably, but is likely to be far from perfect.

Also, note that the PSF is determined from the RDC images which are rotated by the effective P-angle (9.24° in this case). If it is assumed that the PSF on the RDC images are well modeled by a two-dimensional Gaussian (which is in itself unlikely), the PSF used in the calculation of the crosstalks should be rotated by the effective P-angle, which is not done.

From a more fundamental point of view, it would clearly be more satisfactory to estimate the PSF from more direct measurements. In particular, this would make it possible to estimate the true variation of power with l. This is unfortunately very difficult. A possible way to find the PSF is to look at the edges of the solar intensity images and using a model of the limb darkening. Possible problems with this approach include the spatial and temporal variation of the PSF and problems with knowing the limb darkening for the spectral range used, including the observed line. Given that the continuum observations cover a substantial spectral range, and the velocity observations are essentially only sensitive to the spectral line, it is likely that the limb darkening is different for the two types of observations.

There are several ways to create a spatially and temporally varying PSF. Spatial variations can be caused by optical distortions, improper alignments, detector problems etc., temporal variations by seeing, temperature changes affecting the optics, etc. It is also important to note that the degradation caused by the interpolations will vary depending on the magnitude of the centering shifts and the P and B angles. If the images happen to be close to the perfect centering or an integer number of pixels off in some direction, the centering will not degrade the PSF in that direction. Looking at the centering corrections, it is not uncommon for the image to stay at the same position in some direction for a substantial fraction of a day. Also note that the interpolations will contribute differently to the PSF at different points on the image.

One might expect several problems from spatial and temporal variations in the PSF. The spatial variation means that the PSF determined from the center of the disk, as done here, will not be valid in other parts of the image. The temporal variation will effectively change the window function as a function of l and m, leading to different temporal sidelobes.

A more fundamental question is whether one should include the odd ms in the analysis. The argument for not including them is that since we only observe a fraction of the solar surface, the information in the odd ms is redundant. While this may be correct, it is not clear that the information is useless. Most of the redundancy is in that the odd ms from one l appears in the even ms for $l \pm 1$. In these ls the information is not used and is actually unwanted. Including the odd ms is thus likely to improve the errors by a factor of roughly $\sqrt{2}$. The computational expense should be about twice as large. Tests with artificial data seem to indicate that the improvement in accuracy is indeed $\sqrt{2}$.

Another possible improvement is to change the masks used to isolate the individual (l, m)s. While it is generally impossible to perfectly separate the modes, it should be possible to remove a large fraction of the interfering crosstalk from neighboring modes. What makes this possible is that only a small number of other modes will have frequencies (or temporal sidelobes) close to that of the target mode. Doing this is likely to increase the background noise around the modes, but as the crosstalks can be very large, it might nonetheless be worth the effort. The possibility of removing interfering modes by changing the masks is currently being considered for the low l instrument built by Steve Tomczyk, but we have not made any tests yet. Given the nature of this modification, it should be easier to test on low l modes.

It is also possible to modify the temporal sidelobes to a certain extent. In particular, it should be possible to remove much of the power far from the central peak by tapering the window function at the beginning and end of the observing periods. Given a duty cycle of only 35%, it is, however, practically impossible to do anything about the first set of temporal sidelobes.

3.5 Reduction of the aj87 Dataset

The reduction of the aj87 dataset did not significantly differ from the reduction of the aj89 dataset. Some minor details obviously had to be changed, most of which will be discussed below. The basic statistics have been shown earlier.

3.5.1 Initial Reduction

For this run the effective P-angle was determined to be 0° to within the measurable accuracy.

The RDC code used for this run was the new version without the rotations and without removal of objective images. As the noise was found to be fairly uniformly distributed, that is, without the problem found in the aj89 dataset, only the standard apodization was used and no cut was applied. The PSF was determined to have $\sigma_x = 0.016R$ and $\sigma_y = 0.0152R$, from an average of the days 870422, 870523, and 870703. Notice that while the image rotations were not done in this run, the PSF is actually *worse* than in the April-June 1989 run. Again, the effects of the interpolation done in the YLM code were not taken into account in the calculation of the crosstalks.

As the solar radius is roughly the same as in the aj89 run, it was again decided only to reduce up to $l = 120$.

3.5.2 Time-Series Analysis

The length of this time-series was 123097 60s time steps, and the Fourier transforms were done with a length of $125000 = 2^3 5^6$. The effective duty-cycle was 24.23%, which is rather low compared to eg., the aj89 run. The resulting Nyquist frequency was $0.5 * (60s)^{-1} \approx 8333\mu Hz$ and the frequency resolution $(125000 * 60s)^{-1} \approx 0.133\mu Hz$. The daily duty cycle has be shown in figure 3.24.

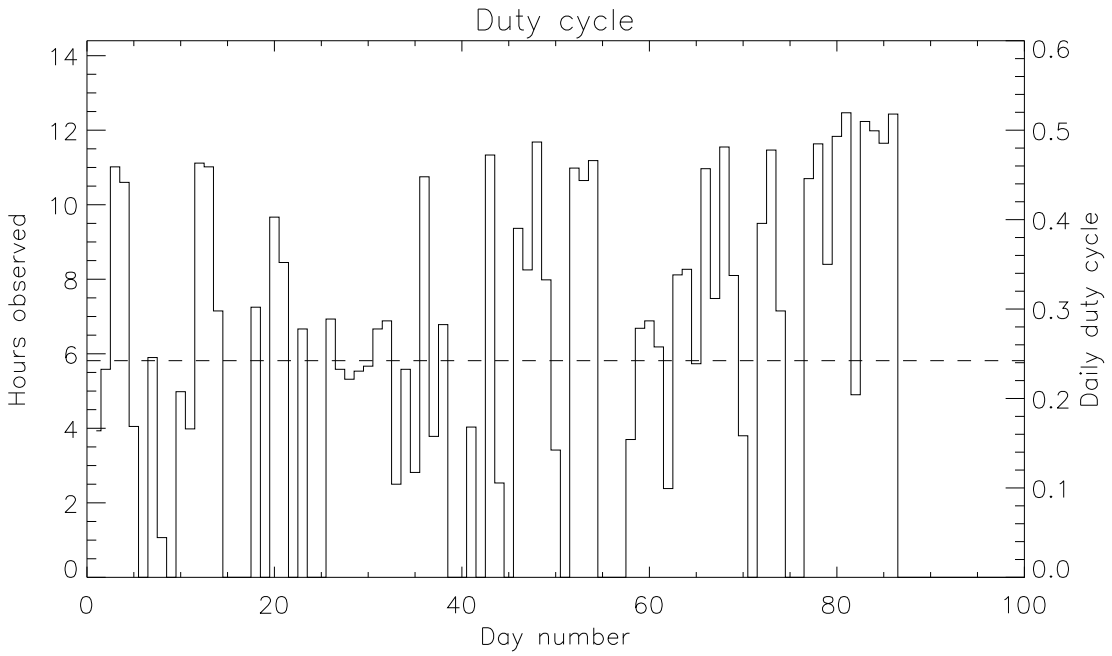

Figure 3.24: The number of hours observed each day (left-hand axis)/daily duty cycle (right-hand axis). The average duty cycle has been indicated by a dashed line.

The noise covariance matrix was determined as an average of parts of the low and high frequency noise.

The fitting was done as in the aj89 dataset, except that the frequency range attempted was $1350\mu Hz \leq \nu \leq 4000\mu Hz$. The lower minimum frequency was used as it appeared possible to fit lower frequency modes as discussed later. Also the fitting window was kept within the range of 40 to 100 frequency bins, corresponding to $5.33\mu Hz - 13.33\mu Hz$.

As initial parameters the values from the aj89 run were used, except that the a-coefficients were reset to the 'standard' values ($440nHz$, $0nHz$, $22nHz$, $0nHz$, $-5nHz$ and $0nHz$). The frequencies used were an average of runs from Libbrecht 1986 and 1988. Again, the fitting was iterated until the various parameters converged. (See the discussion later for a possible problem.)

3.5.3 Discussion

The rejection criteria used were the same as those used for the aj89 dataset. The modeset before (1214 modes) and after (926 modes) the weeding is shown in figure 3.25. The results for various parameters are shown in figures 3.26 and 3.27.

Generally, most parameters behave similarly to the aj89 dataset. There are however, several interesting features in this dataset, not appearing in the aj89 dataset. The 'physically interesting' differences are discussed in a separate section; in this section only the 'pathological' differences will be discussed.

One difference is the different set of modes successfully fit. The two main differences are that it was possible to fit modes with lower frequencies and that there appears to be problems with the high l's. As can be seen from figure 3.25, most of the low frequency modes were rejected, generally because the change from the previous iteration was too large. As the various parameters were not as completely converged as in the aj89 dataset, it is likely that the cause for this is that the fit simply had not converged. It is clear that this should be investigated, possibly by performing more iterations, but this has not been done yet.

Given that the noise problem in the aj89 dataset is absent here (leading to both decreased noise and no chop), the signal to noise ratio would be expected to be better. This would seem to explain why it was possible to fit lower frequency modes. On the other hand, it is interesting that although the new RDC code was used, the PSF is actually slightly *worse*, preventing one from getting the increased quality at high l. It is not clear why the original PSF is so much worse compared to 1989, but one possible cause is that the optics may have been less perfectly aligned. Also, the duty cycle is lower, which should lead to a degraded S/N.

It is, however, not clear why there appear to be problems with the high l, high frequency modes since the PSF is only slightly worse than in aj89. Also the lack of the chop should lead to less m crosstalk. Apart from the difficulties with fitting the modes at all, there appears to be a problem with determining the linewidths, as discussed later, which could be related to this. Also notice that even in the aj89 dataset, there appeared to be problems with many of these modes.

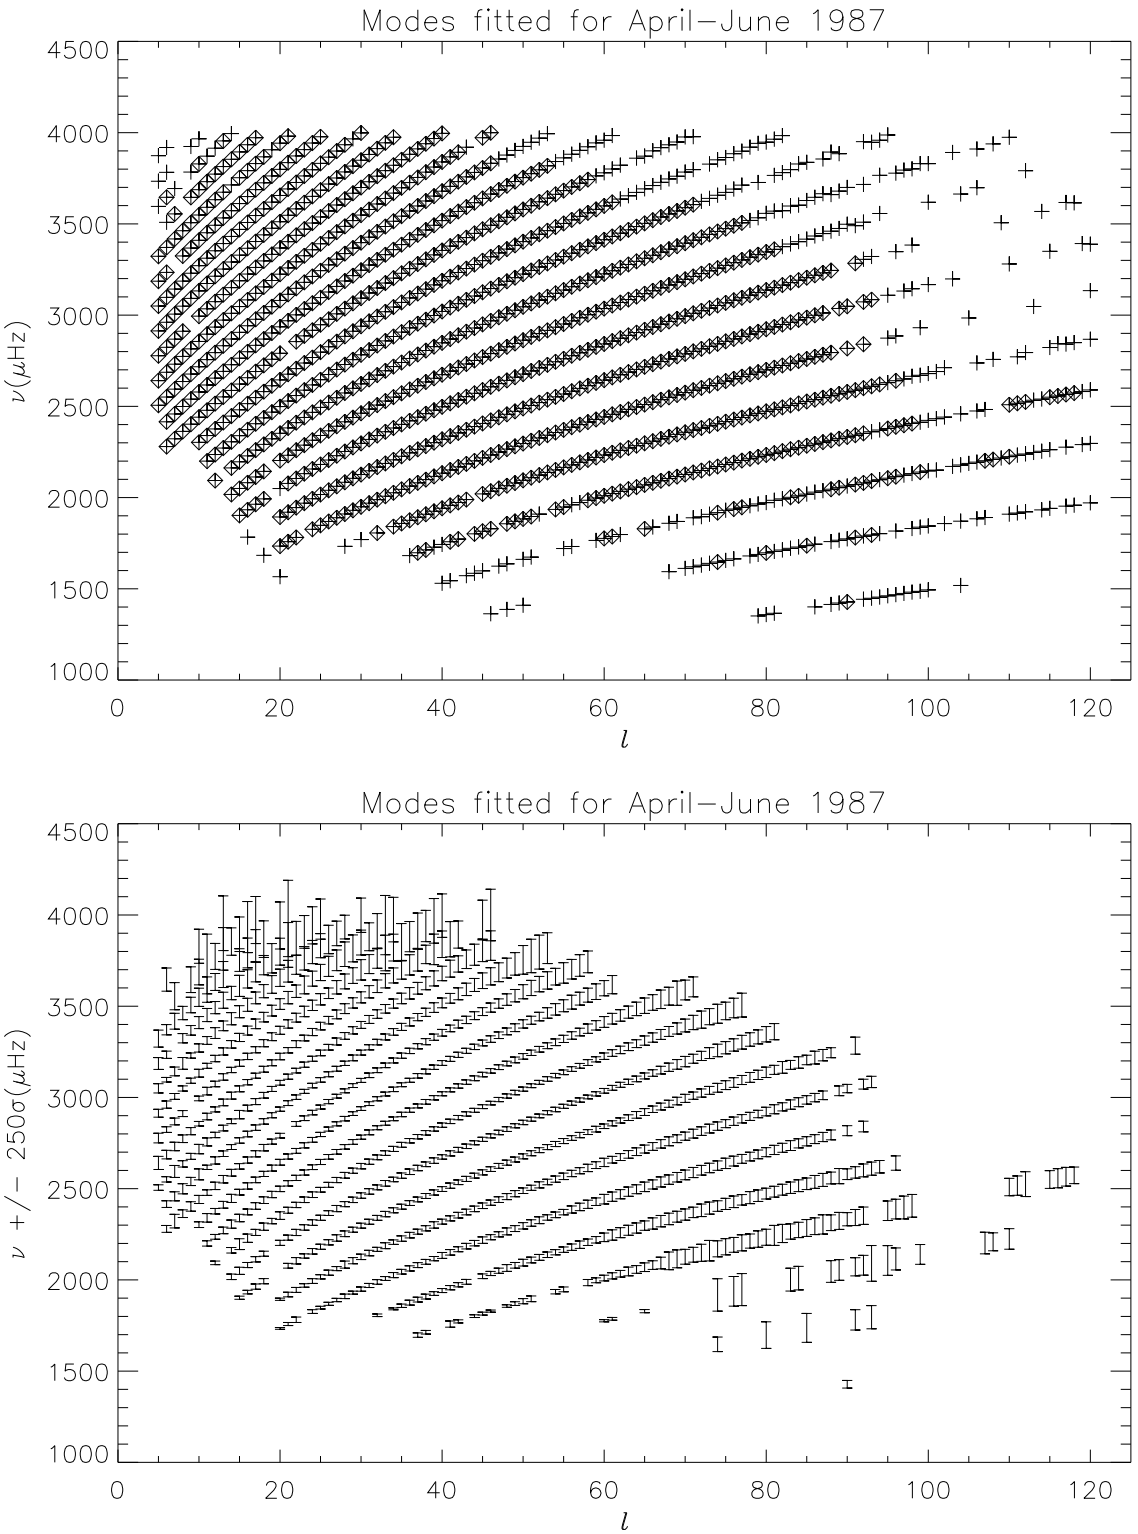

Figure 3.25: The modes fitted in the aj89 dataset. + indicates the original set of modes, diamonds the weeded set of modes. On the bottom plot 250σ error bars were used to better show the size of errors.

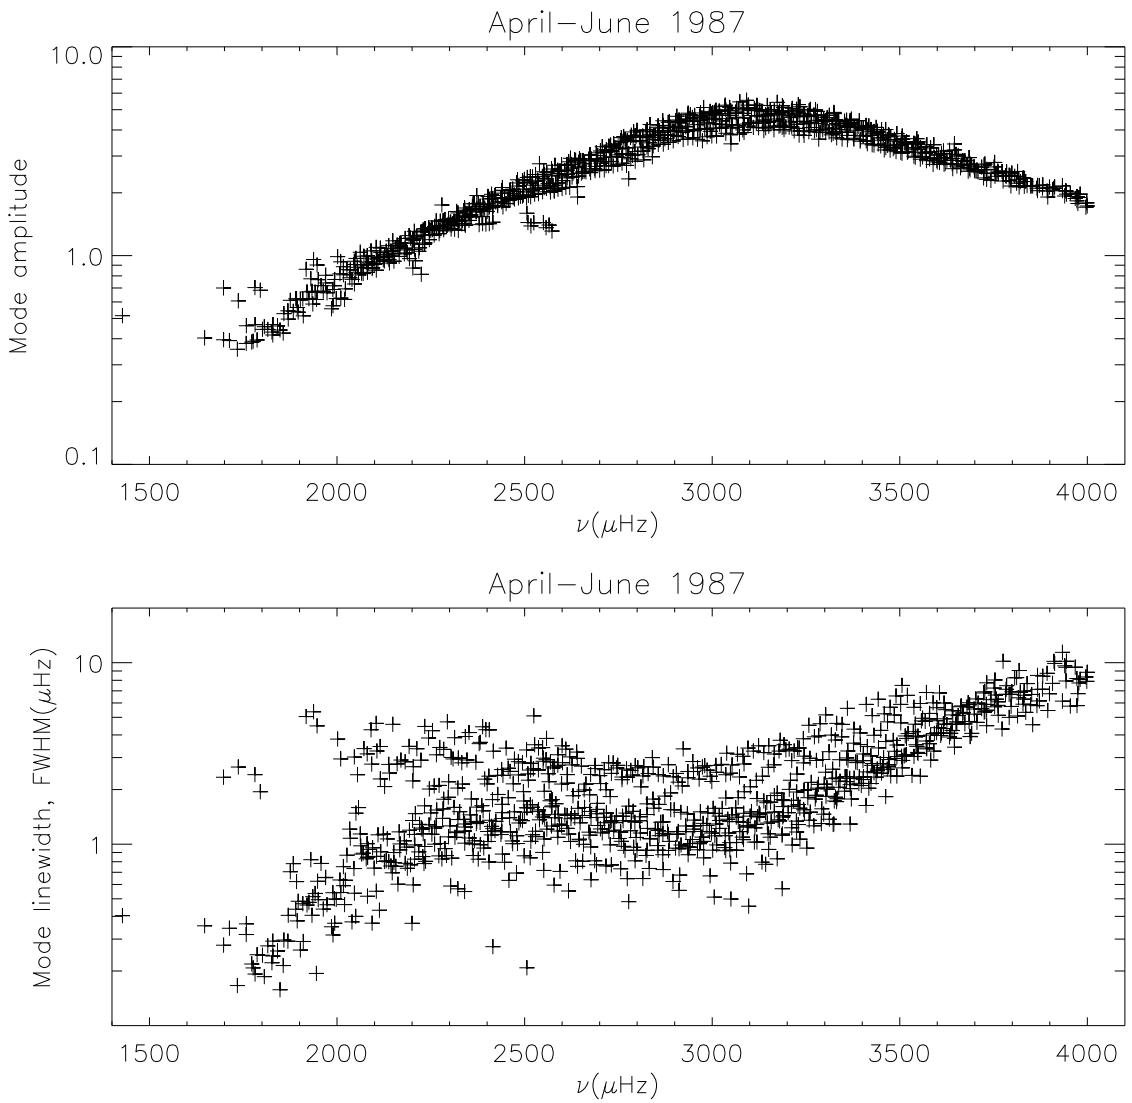

Figure 3.26: Mode amplitudes and linewidths as a function of frequency.

As just mentioned, there appears to be a problem with the linewidths in this dataset. Basically, the problem is that some modes appear to have very large linewidth, while others with similar ls and frequencies do not. This problem is illustrated in figure 3.28. It is not clear if there is a connection between the linewidth problem and the problem of fitting the modes in the first place, but as mentioned earlier, the problem with the high l, low frequency modes (which are the ones behaving differently from 1989) seems to be due to the fit not being converged.

As can be seen from figure 3.28, it appears that, at least for the low ns where the ridges are complete, the linewidths suddenly increase as a function of frequency (which is equivalent to l on each ridge), stays large for a while and drops down to the 'normal' level again. It is interesting that it seems like the low frequency (low linewidth) modes are generally affected more than the high frequency modes at the same l. This appears to be the case even if one looks at the absolute jump in the linewidths, instead of the relative jump. As the problem is only clearly

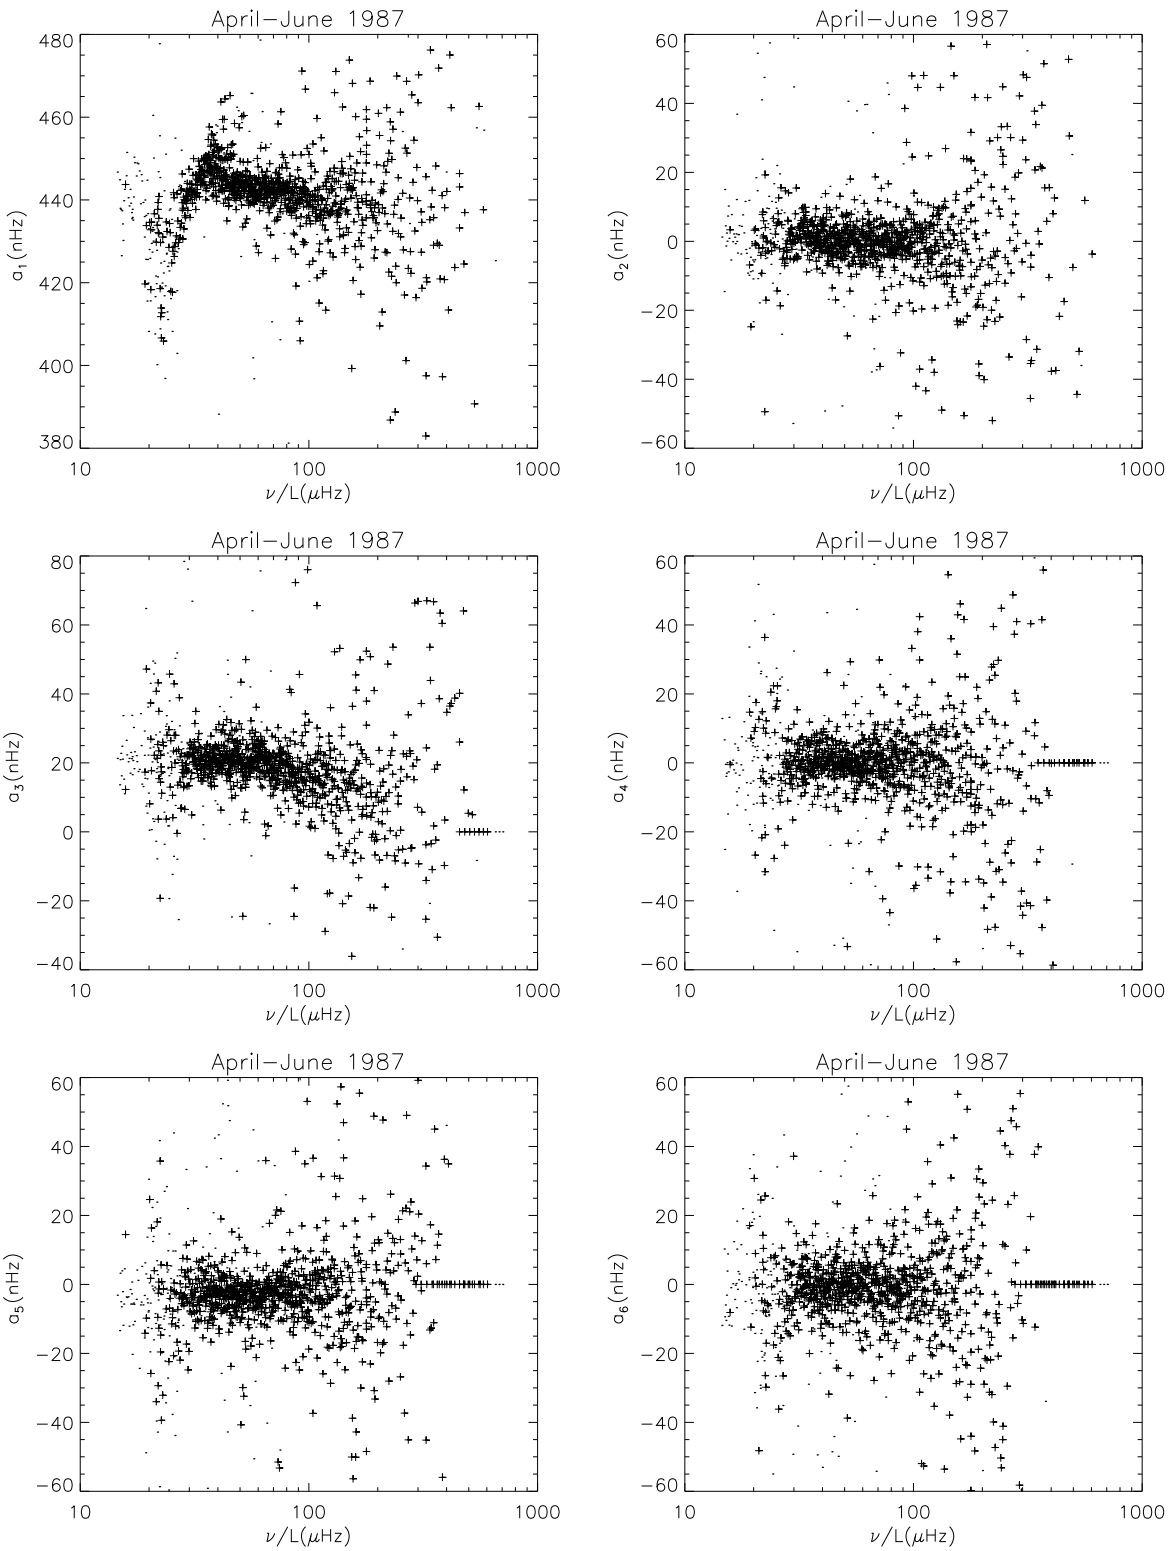

Figure 3.27: a-coefficients as a function of ν/L. The weeded modeset has been indicated by +, the unweeded by dots. The modes with $a_i \equiv 0$ for high ν/L are modes with $l \leq 8$, where not all a-coefficients were fitted for.

visible for the lowest few ns, it is difficult to see what it appears to depend on. It looks like the perturbation is mostly a function of l, instead of say the frequency (which generally determines the linewidths) or ν/L (which determines the spacing between ls).

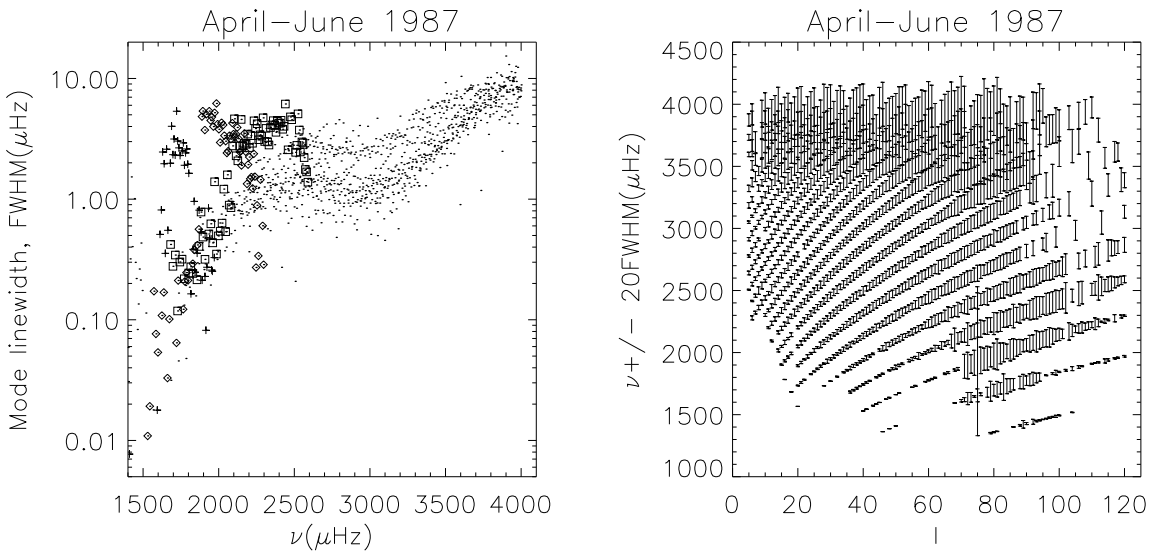

Figure 3.28: The linewidths for the aj87 run. The different n values have been shown with different symbols on the left hand plot, n=2 are indicated by +, n=3 by diamonds and n=4 by squares. All other modes are indicated by a dot. The original (unweeded) set of modes was used for these figures, to better illustrate the problem.

A possible explanation for the linewidth problem could be that the crosstalks are for some reason wrong. If the original ridges can be modeled by a Gaussian of width a and the smearing caused by the incorrect crosstalks by a convolution with a Gaussian of width b, the resulting ridge is a Gaussian of width $\sqrt{a^2+b^2}$, and hence the absolute increase in the linewidth is the largest for the lowest frequencies at a given l. Given that the broadening is of the order a few μHz, it appears most likely that the l crosstalks are the main problem (the relevant ms are spaced by $\approx 0.8\mu Hz$). If this is the explanation, the effect should depend on a combination of the linewidth (which depends mainly on the frequency), the frequency spacing between ls (which in turn depends on ν/L), and l, which is likely to also affect the crosstalks. That the linewidth increases suddenly along a ridge could be because the program only gets confused when some parameter such as the linewidth or the frequency spacing between ls exceeds a certain threshold.

Another major difference is of course the values of the a-coefficients, which are discussed in section 3.7 and 3.8. In particular, the bump in a_1 is very different from the aj89 run.

Examples of the various diagnostics are shown in figure 3.29. As can be seen, they are similar to those from 1989 in that the diagnostics are generally good up to about $l = 60$ after which they develop problems. As in the aj89 case, the likely cause is the incorrectly estimated PSF due to the neglect of the interpolation performed in the YLM code.

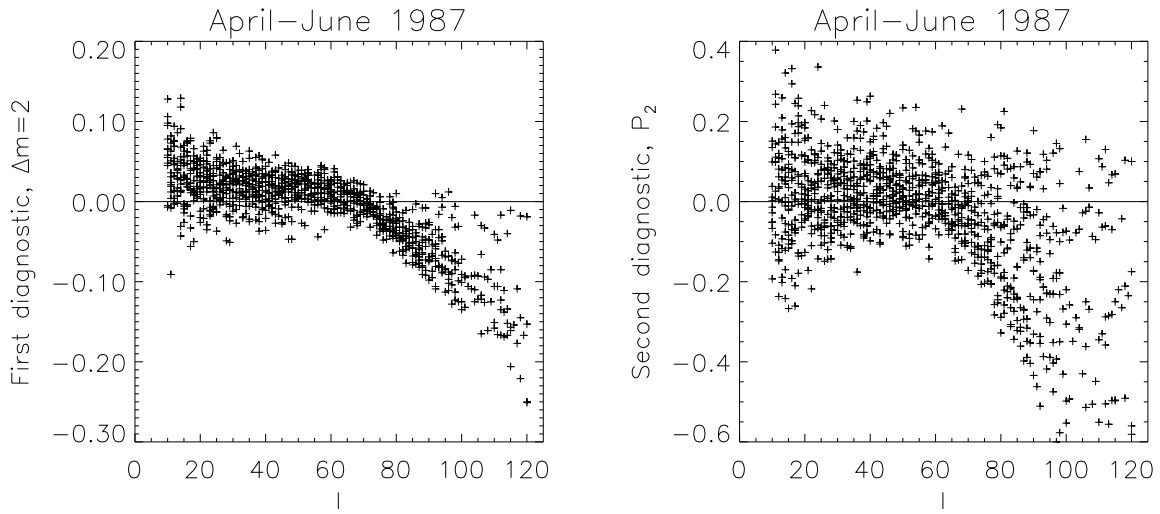

Figure 3.29: The values of two of the diagnostics. The unweeded set of modes was used and only modes with $10 \leq l \leq 120$ were included.

3.6 Tests for Possible Systematic Errors

When the 'bump' in the a-coefficients was found (see eg., the section on the aj89 dataset), it was decided to check whether any of the uncertainties in the underlying assumptions could be the cause. In particular, the fact that the bump is close to the place where temporal sidelobes from neighboring ls are expected to be a problem, has been a point of some concern. Also the effects of various distortions of the solar image, incorrect P-angle, and the cut done to reduce the noise was investigated.

In the following, some of the results of these investigations are presented. For the purpose of reducing the computational burden, only 35 days from the beginning of the aj89 dataset were generally reduced with the various modifications. Given the magnitude of the bump, this should be more than adequate to determine if any of the modifications are the cause of the bump. Many of the comparisons were only done in the range $10 \leq l \leq 110$.

Even if the bump was absent, it would still be a valuable exercise to see how robust the results are to the known uncertainties in the underlying assumptions. Obviously, some implicit underlying assumption may have been overlooked, but an idea of the magnitude of the expected effects has been obtained.

Interestingly, the distance between neighboring ls (that is $\partial\nu/\partial l$ at constant n) is mainly a function of ν/L (which follows from various asymptotic expressions, such as (2.12)), which is also the parameter one would expect a real feature in the solar rotation to show up as a function of, since the turning point depth r_t of a mode is related to ν/L through $r_t/c(r_t) = \nu/L$, where c is the sound speed (cf. eq. (2.8)). The relationship between ν/L and $\partial\nu/\partial l$ is shown in figure 3.30.

As the main concern is the bump, the focus will be on the change in a_1 introduced by the various modifications. The change in other variables will also be shown when they are

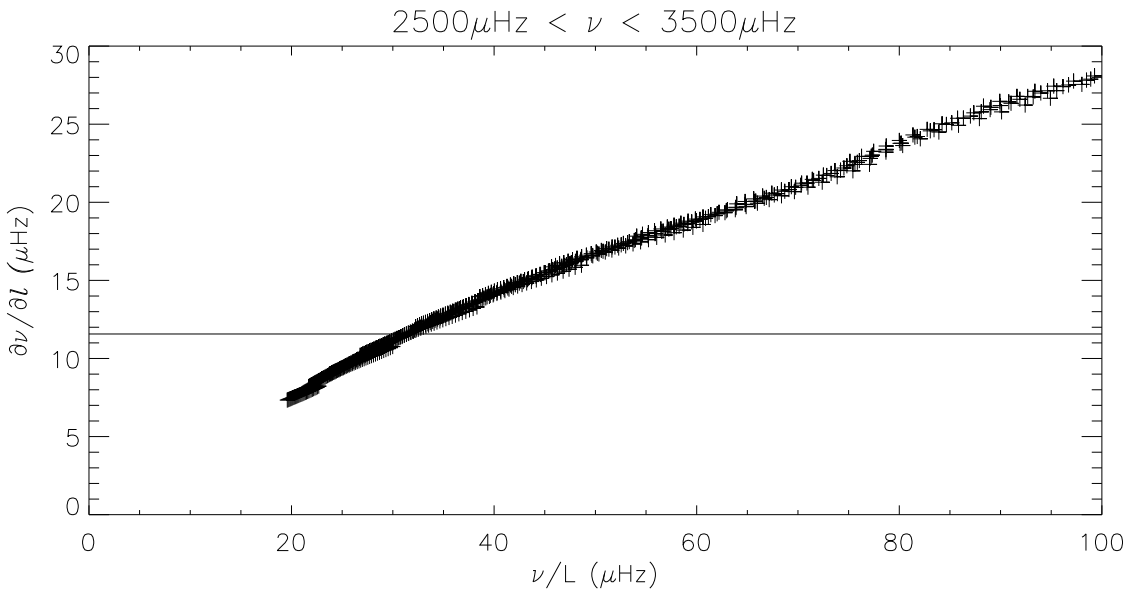

Figure 3.30: $\partial\nu/\partial l$ at fixed n as a function of ν/L. Only modes with $1 \leq l \leq 139$ and $2500\mu Hz \leq \nu \leq 3500\mu Hz$ have been included. Frequencies are from a model by Christensen-Dalsgaard (1991, personal communications). $\nu/L = 11.57\mu Hz$ corresponding to 1 day^{-1} has been indicated by a horizontal line on the plot. Notice that 1 day^{-1} corresponds to $\nu/L \approx 31\mu Hz$.

significant. It is important to remember that the first temporal sidelobes of $\Delta l = \pm 1$ cross the mode of interest when $\nu/L \approx 31\mu Hz$, and that the second set of temporal sidelobes from $\Delta l = \pm 1$ cross when $\nu/L \approx 77\mu Hz$. The temporal sidelobes are placed symmetrically when $\nu/L \approx 53\mu Hz$.

As the sidelobe spacing is mainly a function of ν/L, many of the changes have been plotted as a function of ν/L. Some parameter changes have been plotted as a function of l or ν, especially when they are unrelated to the temporal sidelobes or the bump. As it would require an extreme number of plots to plot everything against everything, the things plotted vary from test to test. Generally, the quantities changing significantly are plotted against the one or two parameters they depend most clearly on. Sometimes a quantity is not plotted even if it does change by a small amount.

Generally an explanation will be given for the observed changes. As some of the changes are fairly unimportant, the explanations will often be brief, and even possibly incorrect. Also sometimes no explanation will be offered.

Most of the modifications were tried before the mode parameters had been iterated to convergence, but results are always compared to a reference run with the same parameters, except when otherwise noted. The output has generally not been weeded and some diverged points are included. Also note that the mode sets successfully fit in each of the tests were not the same and that the modes plotted, which are those common to the reference, are thus generally not the same either.

3.6.1 Time Gap Modification

Since one of the major concerns is that of the 1 day^{-1} sidelobes of neighboring ls interfering with the target mode, this was one of the first effects tested. Since it was impossible to improve the time gap structure, it was instead decided to degrade it. This was done by ignoring the first and last two hours of observations each observing day, thereby reducing the duty-cycle from 34.94% to 17.81%. The sidelobe spectrum before and after the duty cycle reduction has been shown in figure 3.31.

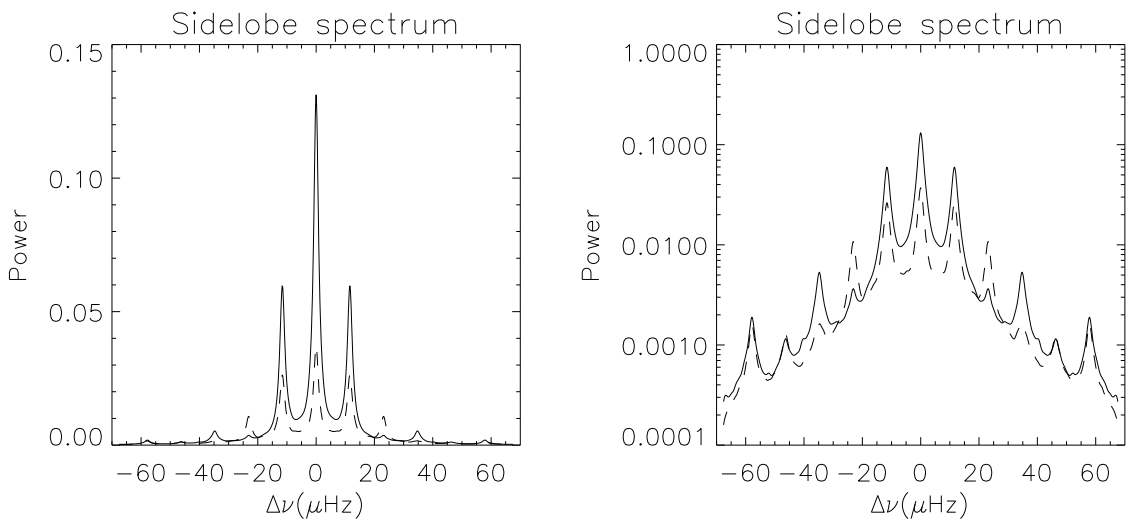

Figure 3.31: Sidelobe spectrum before (solid line) and after (dashed line) the duty cycle reduction. The spectrum has been convolved with a Lorenzian with a HWHM of $1\mu Hz$. The power is normalized such that a perfect window function would have given a peak power of 1.

As can been seen from figure 3.31, the sidelobe spectrum is substantially modified. In the convolved spectrum the ratio of the heights of the first sidelobe to the main lobe changed from 0.45 to 0.71. It is also interesting to note that the second pair of sidelobes goes from being very small to quite noticeable and the third the opposite way. Note that for a cycle with 12 hours on, 12 hours off, 12 hours on . . . (which is almost the case during good periods), only the odd numbered sidelobes (and the main lobe) will appear.

Note that as the data quality is generally better in the middle of the day, it is generally the lower quality data that get rejected this way. An unfortunate side effect of the duty-cycle reduction is that the signal-to-noise ratio is degraded, leading to the loss of some of the low frequency modes.

This test was done with the whole time-series, after the various parameters were almost converged, and with the full set of modes.

The changes in the amplitudes and linewidths are fairly small (figure 3.32), but significant, and would seem to indicate that the sidelobes are not perfectly taken care of by the analysis. This could be due either to a problem with the analysis or some modulation of the sensitivity other

than the day/night cycle. In particular, it is not unlikely that the PSF varies systematically during the day. Although the Fourier Tachometer itself should be insensitive to the varying intensity level during the day, the images are generally degraded early and late in the day, which might lead to a decreased sensitivity. Also the problem with the rings caused by the A/D converter (described elsewhere) is generally worse in the morning and evening, due to the low intensity levels.

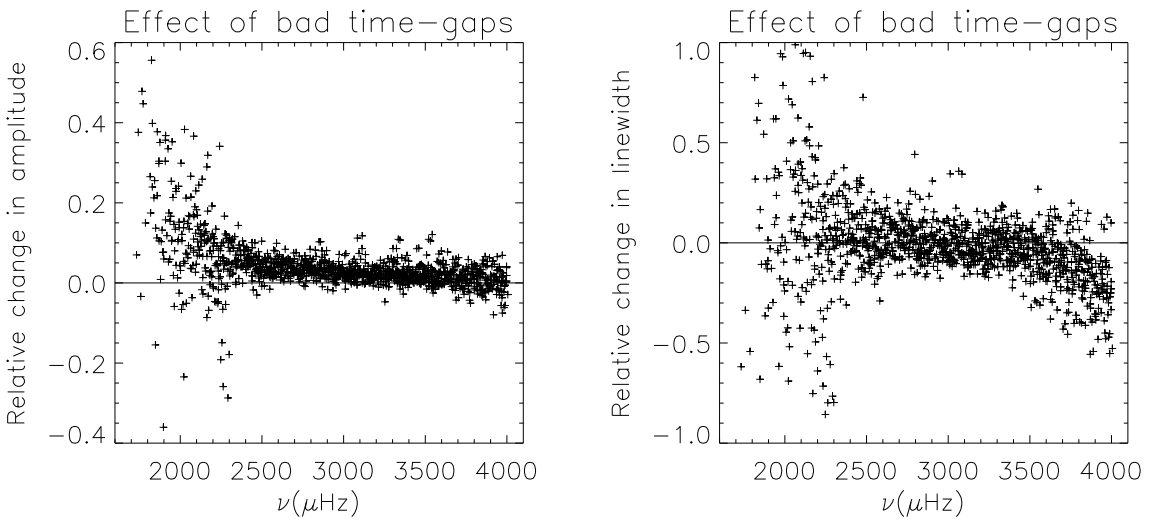

Figure 3.32: The change in the amplitudes and linewidths caused by the degraded time gaps.

The change in the background noise level shown in figure 3.33 is interesting. The general change in the level of the noise is a combination of the decreased duty cycle and the way this is handled in the analysis and is thus not significant. The bump as a function of frequency is most likely caused by the increased spreading of power due to the power in the window function far from the central peaks (compare with the figure of the original noise in the section on the aj89 run). Note that even for the ls taken into account, only the power in the sidelobe structure out to a certain distance from the main lobe is taken into account. The dependence on l, on the other hand, is not easily explainable.

The change in a_1 resulting from this modification is shown in figure 3.34. There is a feature around the high end of the bump/wiggle, but nothing able to explain the wiggle. The changes in the other a-coefficients are generally very small.

Another way of testing whether the bump is an artifact due to a combination of temporal and spatial sidelobes, is to change the spatial sidelobes. This is done in some of the following tests.

3.6.2 Dividing the Time-Series

Another way to test for various problems caused by the temporal sidelobes is to divide the time-series into chunks and analyze them separately. For this purpose the original time-series

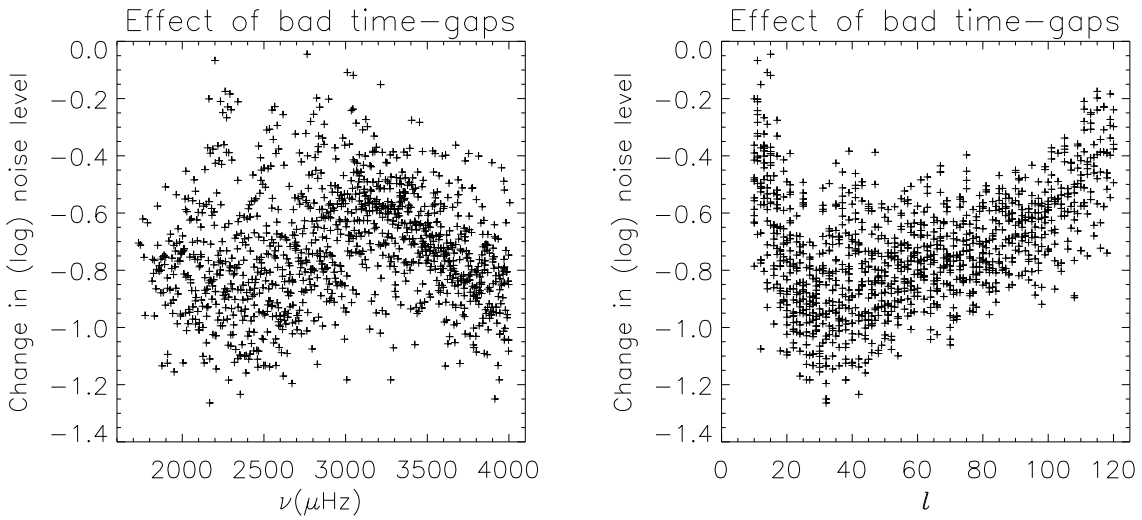

Figure 3.33: The change in the background noise level caused by the change in the time-gap structure.

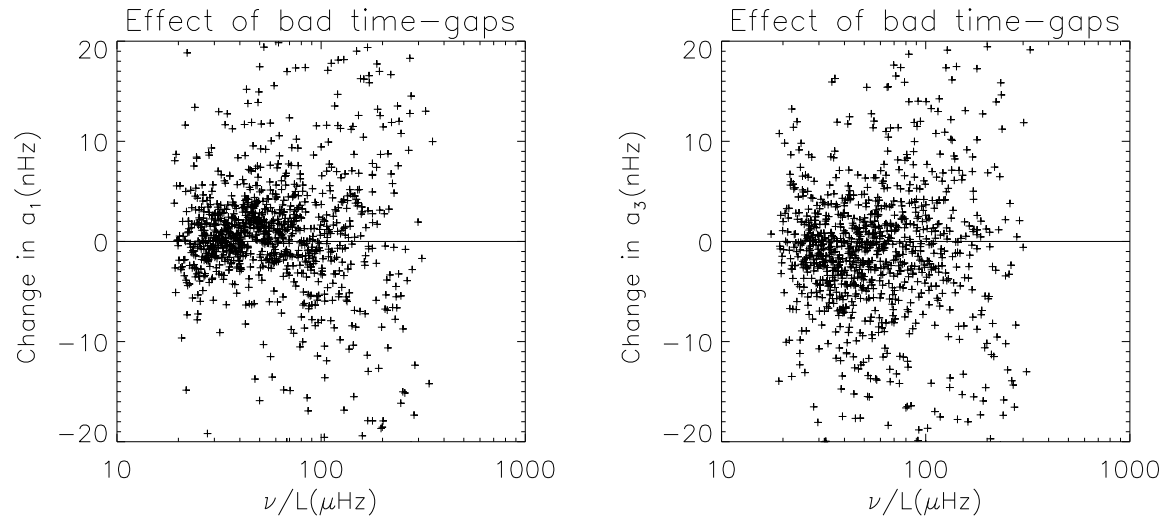

Figure 3.34: The change in a_1 and a_3 as a result of the degradation of the sidelobe spectrum.

was divided into three chunks with the following starting times, lengths, and duty cycles:

Start	Length	Duty Cycle
0	30375	34.48%
31000	34560	33.06%
66200	34560	38.55%

Recall that the original time-series had a length of 100800 time steps and a duty cycle of 34.94%. Note that the three time series are not of the same length and that there are small gaps between them (all when there were nighttime gaps in the original series). Each of these time-series was

then analyzed separately and the results averaged. The original time-series was also analyzed using the same initial parameter estimates, except for the background noise level, which had to be changed to reflect the different lengths.

The differences between the results of the analysis of the original time-series and the average of the chunks have been shown in figure 3.35.

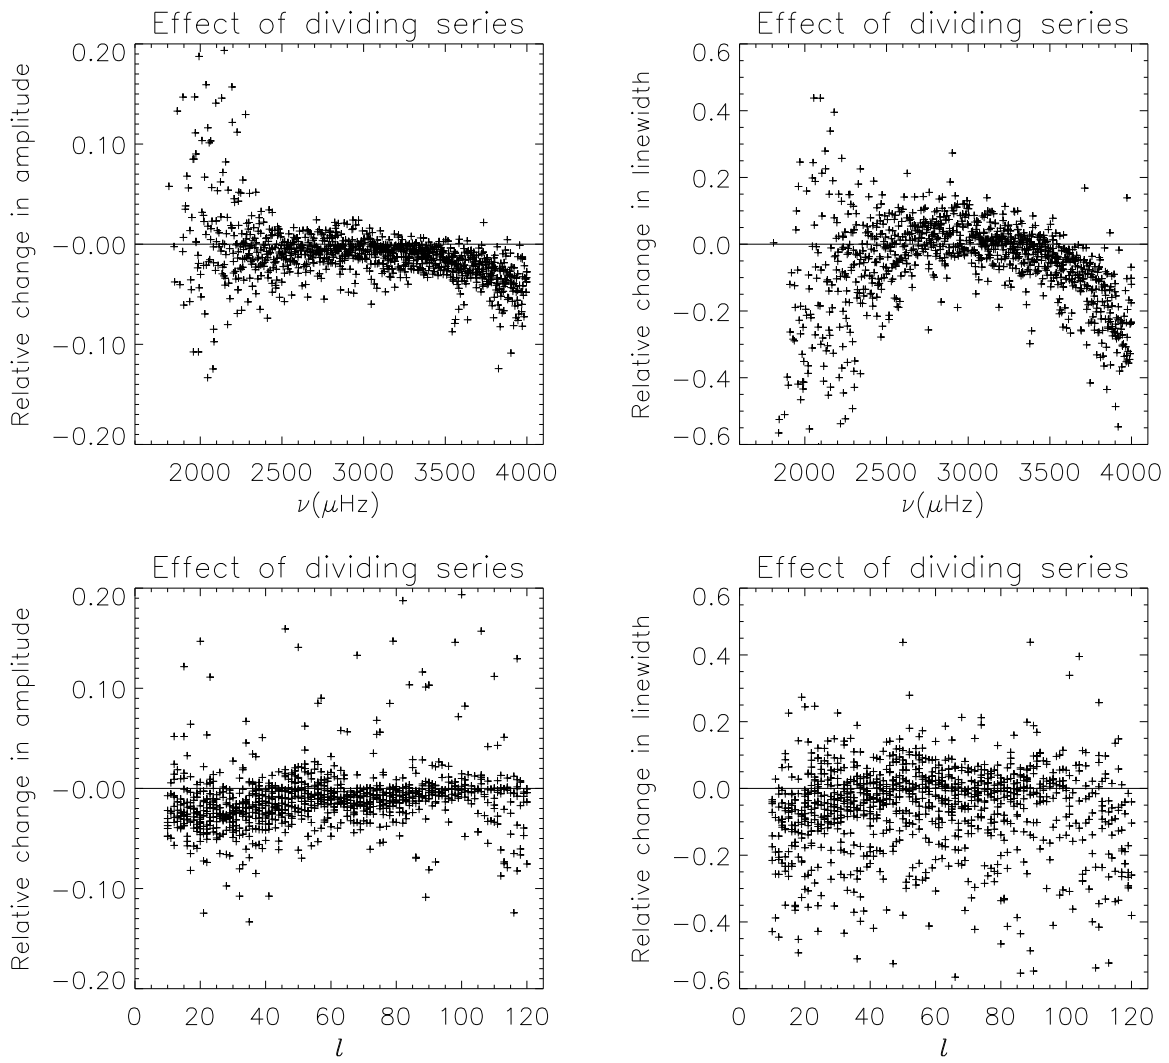

Figure 3.35: The differences in the amplitudes and linewidths between the average of the analysis of the three chunks and the analysis of the whole time-series.

As can be seen from this figure, both the mode amplitudes and the mode linewidths have changed. The changes are generally largest at high frequencies. Although there should preferably not have been any changes, it is not surprising that some occur at the high frequencies where the linewidths are comparable to the size of the fitting window and 1 day^{-1}. It is important to note that the size of the fitting window in frequency was not the same in the different runs, which may explain some of the differences. It is interesting that it appears that both the amplitudes

and the linewidths change more at low ls (less than say, 40) than at medium ls. The cause for this is unknown.

The linewidths also appear to be lower in the average than in the analysis of the whole time-series. A possible cause is that the frequency bins are of the order $0.4\mu Hz$ for the three chunks, which is about the same as the linewidths at low frequencies. In other words, the modes are coherent over the subparts.

Another problem is that it is well known that the mode frequencies change with time (see section 3.8 on comparisons of different years). This would be likely to appear as increased linewidths when analysing long runs compared to shorter runs.

The other parameters, such as the a-coefficients, only change by small amounts.

3.6.3 Geometric Transformations

A common problem with many optics/detector systems is various distortions of the images and misalignments. As it is known that we do have small distortions such as slightly oblate images and uncertain P-angles, we decided to try to deliberately distort the images in various ways and look at the changes in the various fitted parameters.

Generally, one would expect that various geometric transformations would lead to slightly wrong crosstalks (assuming that the current best values are correct). Generally, the amplitude would be expected to decrease as the masks fit the modes less accurately. The linewidths would most likely increase as the inaccurate masks spread out the sensitivity more than expected. The noise estimates are also likely to increase as some of the power spread out to other modes might be interpreted as additional noise. The changes in the a-coefficients are likely to depend on the details of the distortion and the sign should vary with the sign of the distortion.

For high resolution studies where only a fraction of the Sun is observed, it can be difficult to get the correct image scale. If the linewidths are also comparable to the frequency spacing, it is not possible to see individual modes (ls) and the spectrum turns into ridges instead of separate modes. In this case, the only way to tell l (or k in physical units) from k in pixels (which is what one can get from the images) is to know or assume a solar radius in pixels. If the radius is off by 1%, l and m will be off by 1% and the a_1 determined will be off by 1%. In view of this, one could naively expect to see some change in a_1, when the distortion causes a scale change. On the other hand, the linewidths are generally low enough for low ls that the individual ls can be separated. Note that the ms generally can't be separated, due to the close frequency spacing.

All the experiments in this section were done with only the first approximately one month of data (actually 40000 75s time steps were used) and only for $10 \leq l \leq 110$.

As the determination of the exact solar radius in pixels is somewhat uncertain, the first 'geometric transformation' tried was to change the assumed radius by ± 0.5 pixels (corresponding to roughly 1%). The changes in various parameters resulting from this are shown in figure 3.36.

The change in a_1 is rather small and far from able to explain the bump; after all, we assume that the radius is not off by substantially more than 0.5 pixels. It is interesting that for ν/L

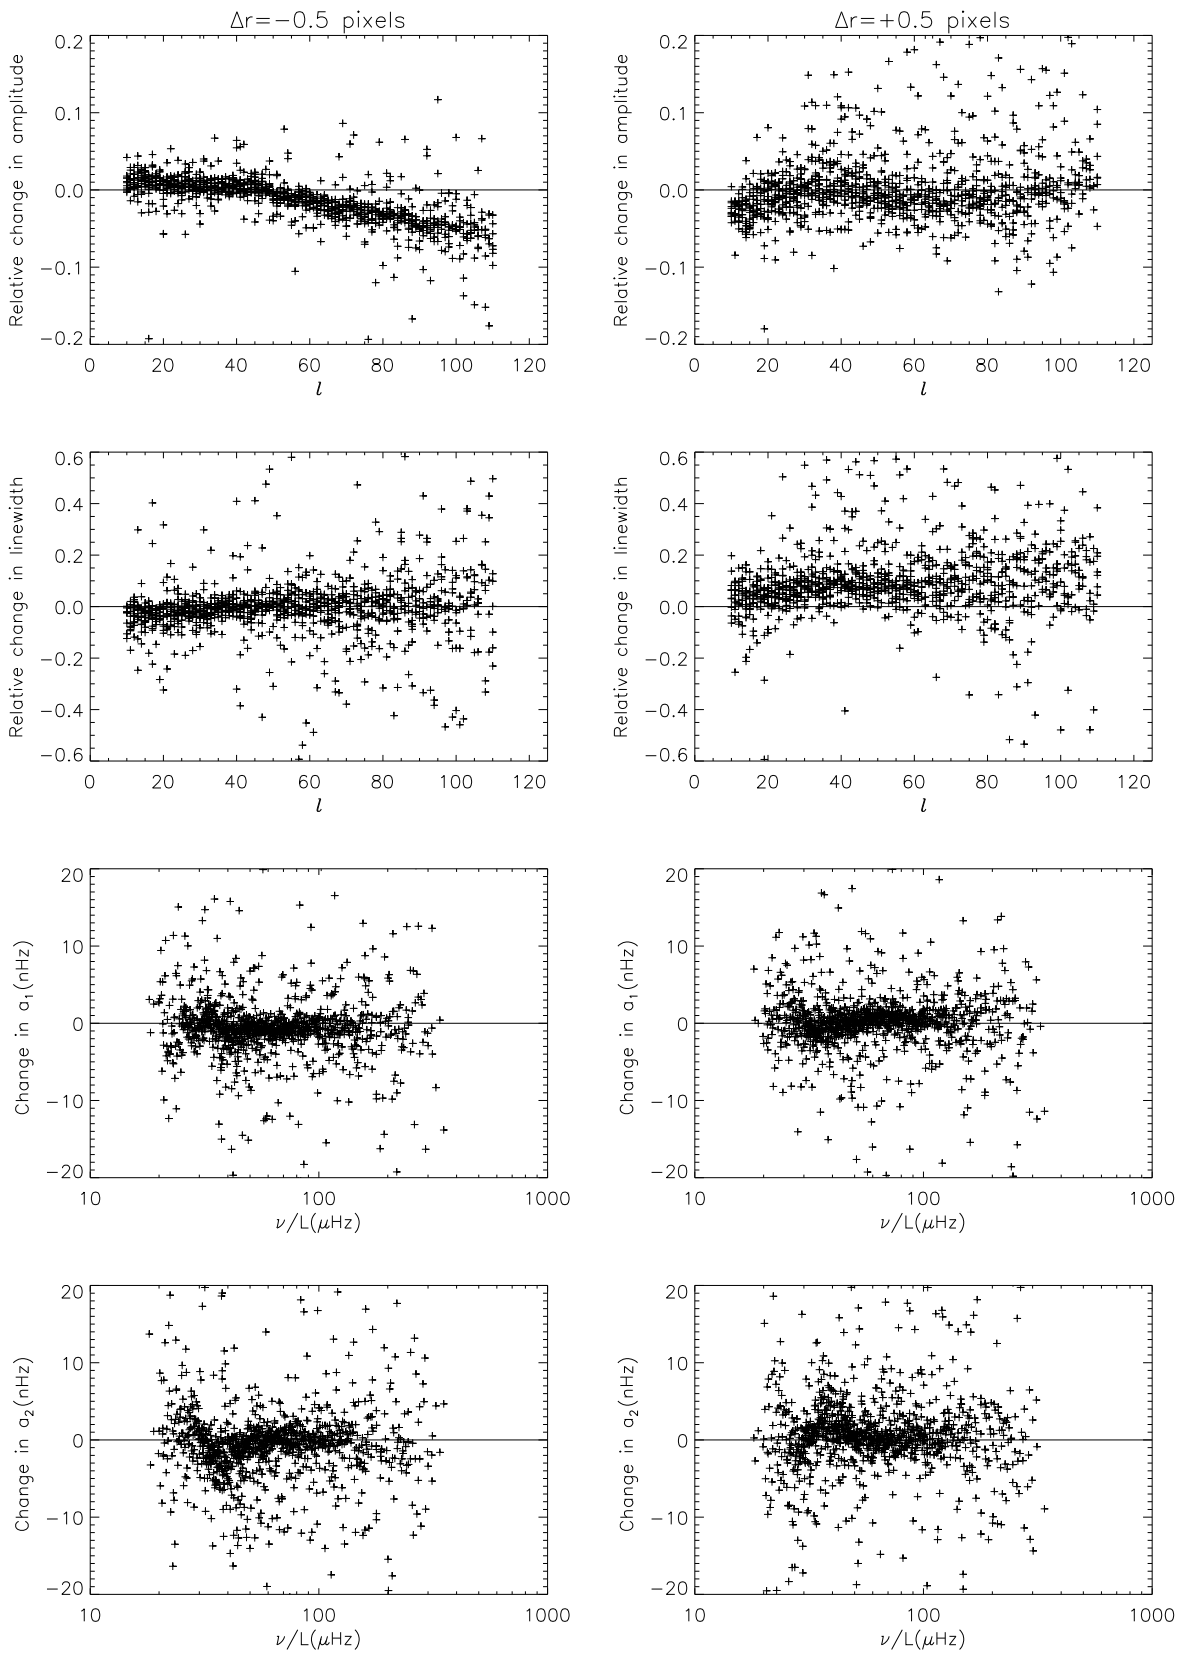

Figure 3.36: Changes in various parameters caused by changing the assumed radius.

above the bump, a_1 does seem to change systematically by approximately 1nHz. The reason that a small change in the region of the bump is visible may be that the sensitivity to incorrect initial mode parameters (in particular for modes other than the target mode, as shown later) varies slightly (in this case the initial a_1 was a constant for all modes).

Of the other a-coefficients only a_2 changes significantly. A possible explanation for this change is given under the description of the squeezing, where the effect is larger.

The changes for some of the other parameters are rather puzzling, in particular the mode amplitude has a rather unusual behavior. Not only is the trend with l different for $\Delta r = -0.5$ pixels and $\Delta r = +0.5$ pixels, but the scatter is also very different. It is not quite clear what exactly is happening. Changing the radius slightly could bring in some very noisy pixels, thereby increasing the noise, although this should generally have a fairly small effect on the estimated parameters (Schou 1993b). The trend difference is more difficult to understand, and it could be argued that it would be interesting to try more radius changes, but this has not been done, as it requires reruns of the whole reduction code, which are very computationally expensive. Naively, one would expect the amplitudes to decrease, given that the masks should fit the modes less well when the radius changes away from its correct value. For high ls the decrease is expected to be larger as the spatial wavelength gets smaller, as is indeed seen for $dr = -0.5$ pixels.

The linewidths only change slightly for $dr = -0.5$ pixels, while they increase systematically for $dr = +0.5$ pixels. The noise change is also different for the two radius changes.

The fact that the various parameters change differently depending on the sign of the radius change would seem to indicate that the current assumed radius may be somewhat off, though the direction is not quite clear. Fortunately, most of the changes are rather small.

As has been discussed before, another known problem is that of determining the effective P-angle. To test the sensitivity to this we changed the P-angle by 1.5°, which should be more that the uncertainty in the determination. The resulting change in a_1 is shown in figure 3.37. Again, the change is much smaller than what is needed to explain the bump.

None of the other parameters changed noticeably.

Although we have tested that the solar images are round to better than a pixel, we decided to squeeze the images deliberately (actually we assumed in the reduction code that the images were elliptical and corrected for the distortion). Among the things that could cause the solar images to be distorted are nonsquare detector pixels and misaligned optics. (For instance, the normal to the detector could make a small angle to the optical axis, which would of course also change the PSF as a function of position on the chip.) There are indications that the solar image is indeed somewhat oblate, as discussed in section 3.7, but only by a fraction of a pixel. In the test performed the images were compressed 1% in the vertical direction and stretched by 1% in the horizontal direction. The resulting changes in different parameters are shown in figures 3.38 and 3.39.

The change in the amplitude is fairly constant at a little over -2%, with some dependence on frequency and l. The cause of this change is presumably the general power lowering discussed earlier due to the effectively incorrect masks used. The linewidth seems to be mostly unaffected,

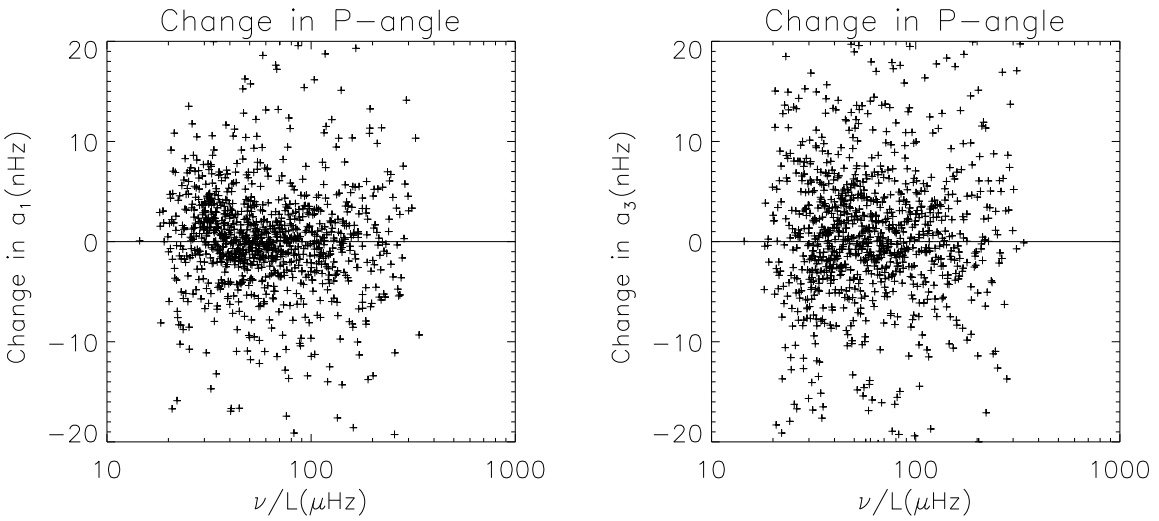

Figure 3.37: The change in a_1 and a_3 caused by a change in the P-angle of 1.5°, as a function of ν/L.

although there is a hint of a change.

The a-coefficients on the other hand change significantly. a_1 decreases by somewhat over 1% for most modes, and does show a small bump (compared to the original bump), possibly due to the scale change, especially in the horizontal direction. a_3 and a_5 both change by significant amounts too, also with some hint of a bump.

The interesting parameter, however, is a_2, where the change is very large and resembles the familiar bump in a_1. To see why such a distortion, or any similar distortion which shifts the sensitivity in l, will give rise to a change in the mode frequencies and even a-coefficients, an example of a side lobe spectrum has been shown in figure 3.40.

Considering figure 3.40, it is easy to see that the mode frequency will change, since the center of gravity of the central ridge changes. That the even a coefficients also change is due to the sensitivity to the perturbation depending on m. Given that the perturbation is an even function of m, it should only affect the even a-coefficients and not the odd a-coefficients. Also notice that the effect should disappear when the temporal sidelobes of the neighboring ls fall exactly on the main lobe, as is indeed observed, and that the perturbation should be roughly antisymmetric around this point.

It is however, difficult to see how such a distortion in the optics would go unnoticed.

A common problem with optical systems is distortions. In particular, there is often a cubic distortion term, such that a point, which, in the absence of any distortion would have distance from the optical axis of r, ends up in a position $r' = cr + c'r^3$. Such a distortion would be extremely difficult to see as the solar limb would stay circular. Although it is believed that any such distortion would be rather small in the Fourier Tachometer, we have tried to introduce such a distortion to test for possible effects. The original distortion was chosen to keep the solar radius constant and having a maximum value of $r' - r$ of 1 pixel. The modes successfully fitted

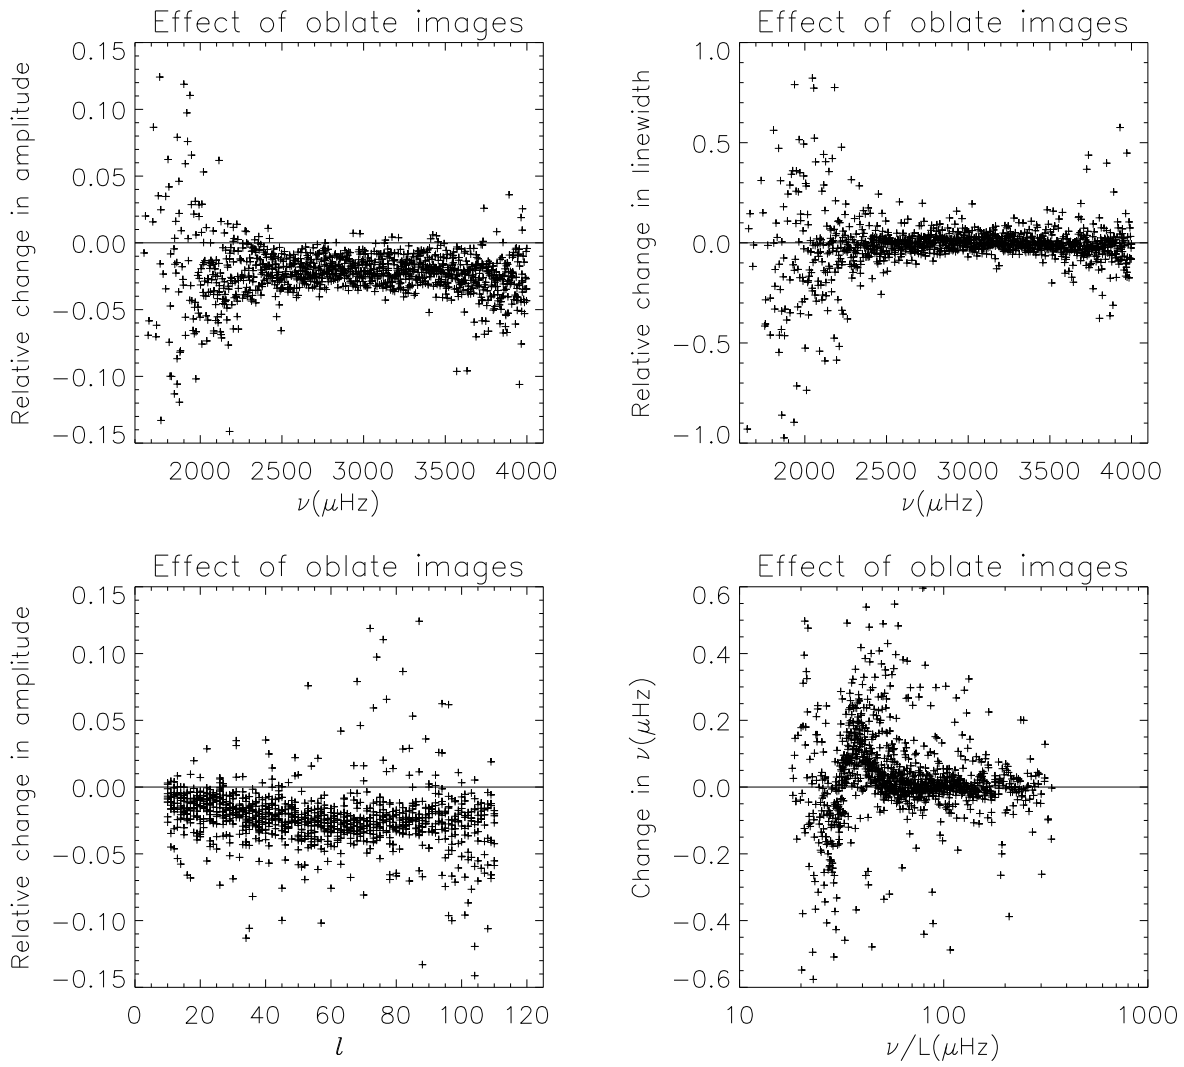

Figure 3.38: The changes in the amplitudes, linewidths, and mode frequencies caused by making the solar images oblate.

are shown in figure 3.41. As can be seen, a lot of modes were missed. The most obvious feature, apart form the missing modes, is that many parameters (not plotted) have a large scatter. The cause for large scatter is probably that the scale change close to the center of the solar images (where most of the useful signal is) is large, of the order 5%, which probably also explains the problem with the modes the fitting procedure could not converge (modes that either did not converge within a preset number of iterations or gave totally unrealistic parameters).

In view of this it was decided to redo the distortion with the opposite sign and half the magnitude. Results of this are shown in figures 3.41 and 3.42. The amplitude change is rather large and increases for higher ls, this is probably an effect of the scale change. The change in the linewidth and a_1 is small. a_2 shows a some change, again probably due to the scale change. Also notice the change in the mode frequency, around the location of the bump, probably due to the same effect as for the oblateness.

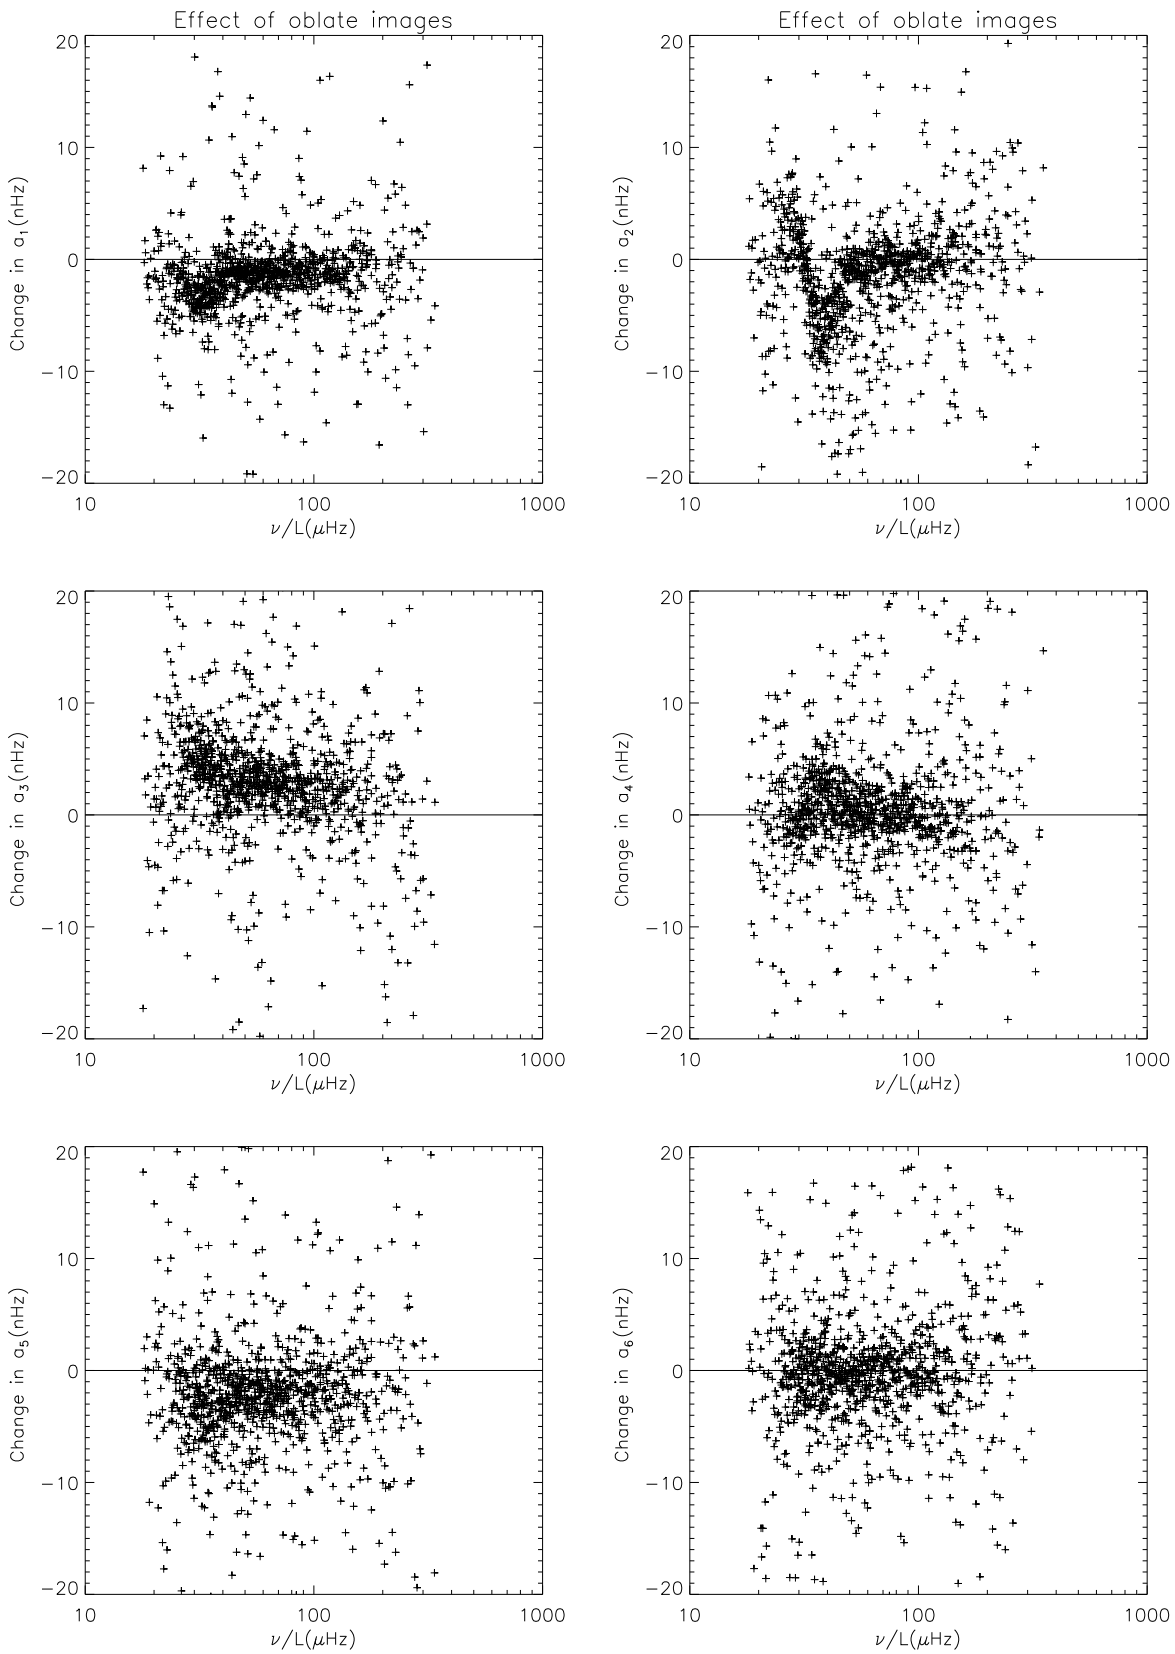

Figure 3.39: The changes in the a-coefficients caused by making the solar images oblate.

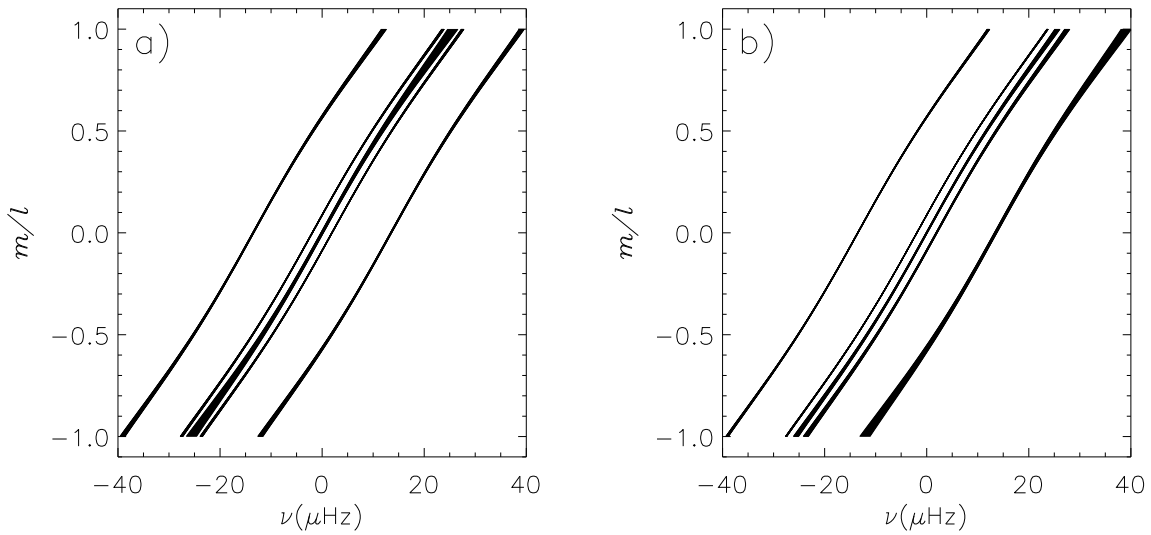

Figure 3.40: A schematic of the effect of distortions on the sidelobe structure. a) shows the structure before the distortion, while b) shows the structure after. The five lines shown are from left to right, $\Delta l = -1$, main lobe, $\Delta l = -1$, temporal sidelobe, $\Delta l = 0$, main lobe, $\Delta l = 1$, temporal sidelobe and $\Delta l = 1$, main lobe. The magnitude of the distortion was chosen such as to shift the peak of the sensitivity from $\Delta l = 0$ to $\Delta l = 1$. The thicknesses of the lines indicate the strength of the various peaks and are not from an actual calculation. Note that the typical fitting window used in the analysis does not include the main lobes of the neighboring ls. The leakage among the ms, as well as several temporal sidelobes and other ls and ns have been left out for clarity.

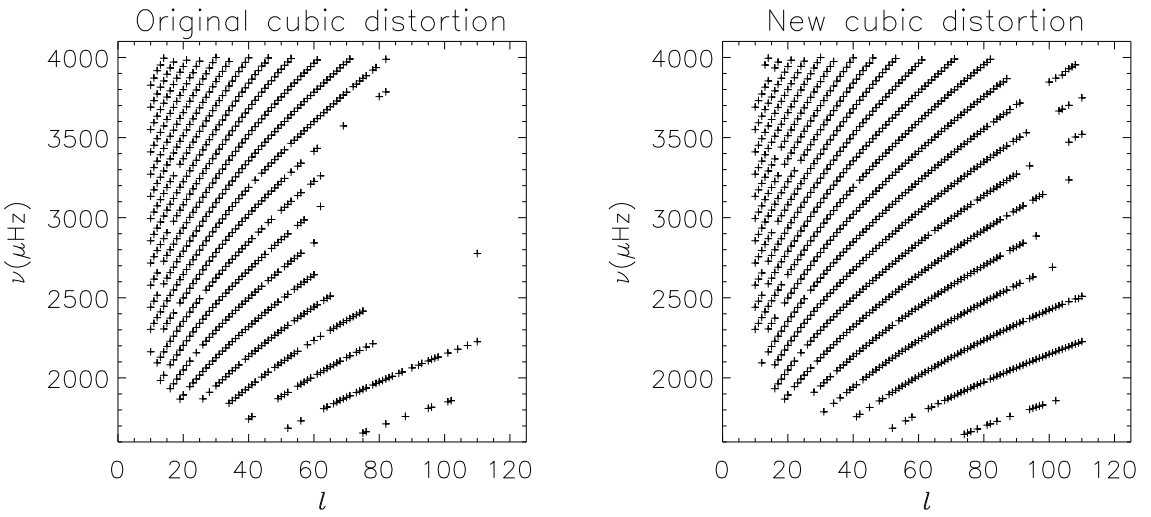

Figure 3.41: The modes successfully fit with the original magnitude of the cubic distortion (left) and with -0.5 times this magnitude (right).

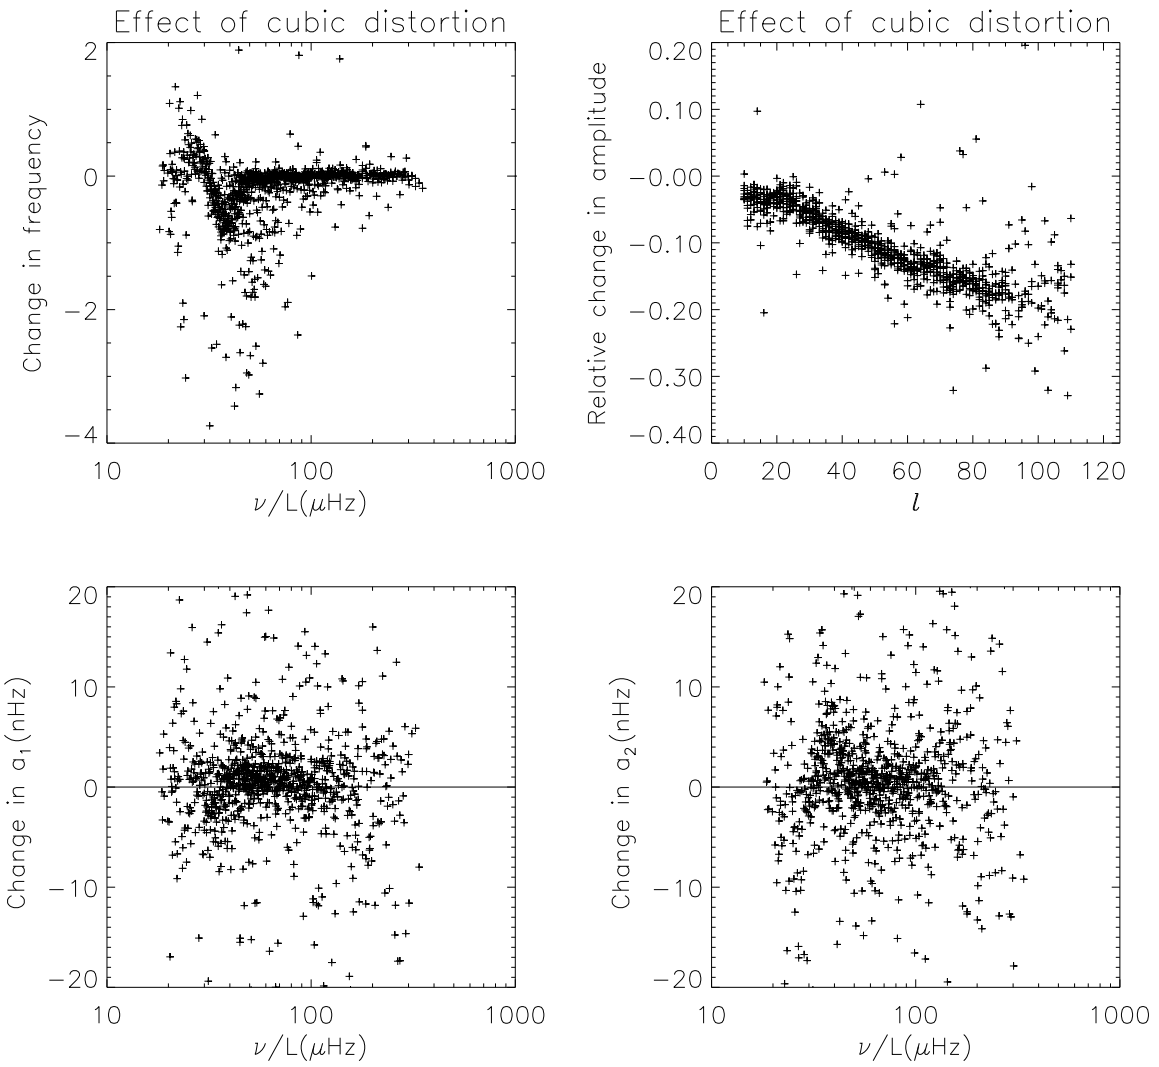

Figure 3.42: The change in the mode frequencies, amplitudes, a_1, and a_2, caused by the new magnitude of the cubic distortion.

3.6.4 Modification of Cut

The difference in the appearance of the bump between different years made us consider the various differences between the different years. Except for such trivial differences as the exact gap structure, the main differences are that the sets are taken with different time-cadences, and the noise problem in 1989 with its associated modification of the apodization. Also the effective P-angle was different, but this has been covered in a previous section.

The fact that the observations were done with different integration times is, for all practical purposes, impossible to change after the fact. Both sets could naturally be binned to one observation every five minutes, but this would make the Nyquist frequency, which is already lower than what one would like, so low that the p-modes would disappear. Interpolating one of the sets to the same cadence as the other is a possibility, but it would lead to a number of

problems with the statistical properties of the resulting time-series. Also, it would be somewhat unpleasant computationally.

The cut on the other hand is easy to change. Given the high noise in the part of the Sun that was cut away, it was decided to make the cut worse, instead of removing it. The original cut was at $r \approx 0.6R$ and it was decided to put the new cut at $r \approx 0.5R$. This obviously modifies the crosstalk, which thus had to be recomputed. The result of this run is shown in figure 3.43. The change in a_1 is extremely small. There are small changes in some of the other a-coefficients, again this may be significant or just an artifact of the initial estimates being off.

Notice that since the temporal sidelobe problem is really due to a combination of spatial and temporal sidelobes, modifying the spatial sidelobes, as done here, should change a_1 if it is indeed an analysis problem. On the other hand, changing the apodization in the horizontal direction should mainly affect the m crosstalk, as opposed to the l crosstalk which is more likely to be the cause of the problem.

Several of the other parameters change here, as shown on figure 3.43. The amplitudes should not change, as the smaller area observed is reflected in the crosstalks. As is clear from figure 3.43, the amplitude does change by roughly -2%. Furthermore, the change seems to be larger at low and high frequencies. The explanation for this is not immediately obvious, but could be related to an inaccurate estimate for the PSF. The change in linewidth does not appear to be significant, although there is a hint of a small decrease for high frequencies.

Not too surprisingly, the noise level decreases by about 20%, which is roughly the reduction in the observed area. It is, however, interesting that the reduction depends on both ν and l. This would seem to indicate that there are different noise sources with different spatial and/or temporal properties. Remember that the discussion in the section on basic statistical properties of the data was mainly concerned with frequencies below those considered here, or suggested that the data were 'contaminated' by p-modes.

Since the time-series and crosstalks were available, I decided to try to analyze the time-series with the modified cut using the original crosstalks. Assuming the wrong crosstalks leads to some problems for the fitting procedure, resulting in more modes failing, that is, the number of successful modes is lower. The result is shown in figure 3.44. Again, the change in a_1 is negligible.

Several of the other parameters do change considerably, however. The amplitudes are lower, as less of the Sun is observed than assumed. Again there is a strong l dependence. It is perhaps not surprising that high ls are affected least, as a combination of foreshortening and the PSF has already reduced the amplitude away from the center of the disk.

The linewidths on the other hand are significantly off. It appears that the change is fairly constant in terms of the absolute change, which is $\approx 0.2\mu Hz$. This difference is likely to be due to the assumed m crosstalks being too small. The additional crosstalk is then interpreted as additional linewidth (notice that the neighboring ms are spaced by $\approx 0.82\mu Hz$).

As mentioned earlier and in Schou (1993a, cf. App. 5), one may construct certain diagnostics, which will in some cases tell when the crosstalks are incorrect. Examples of the changes

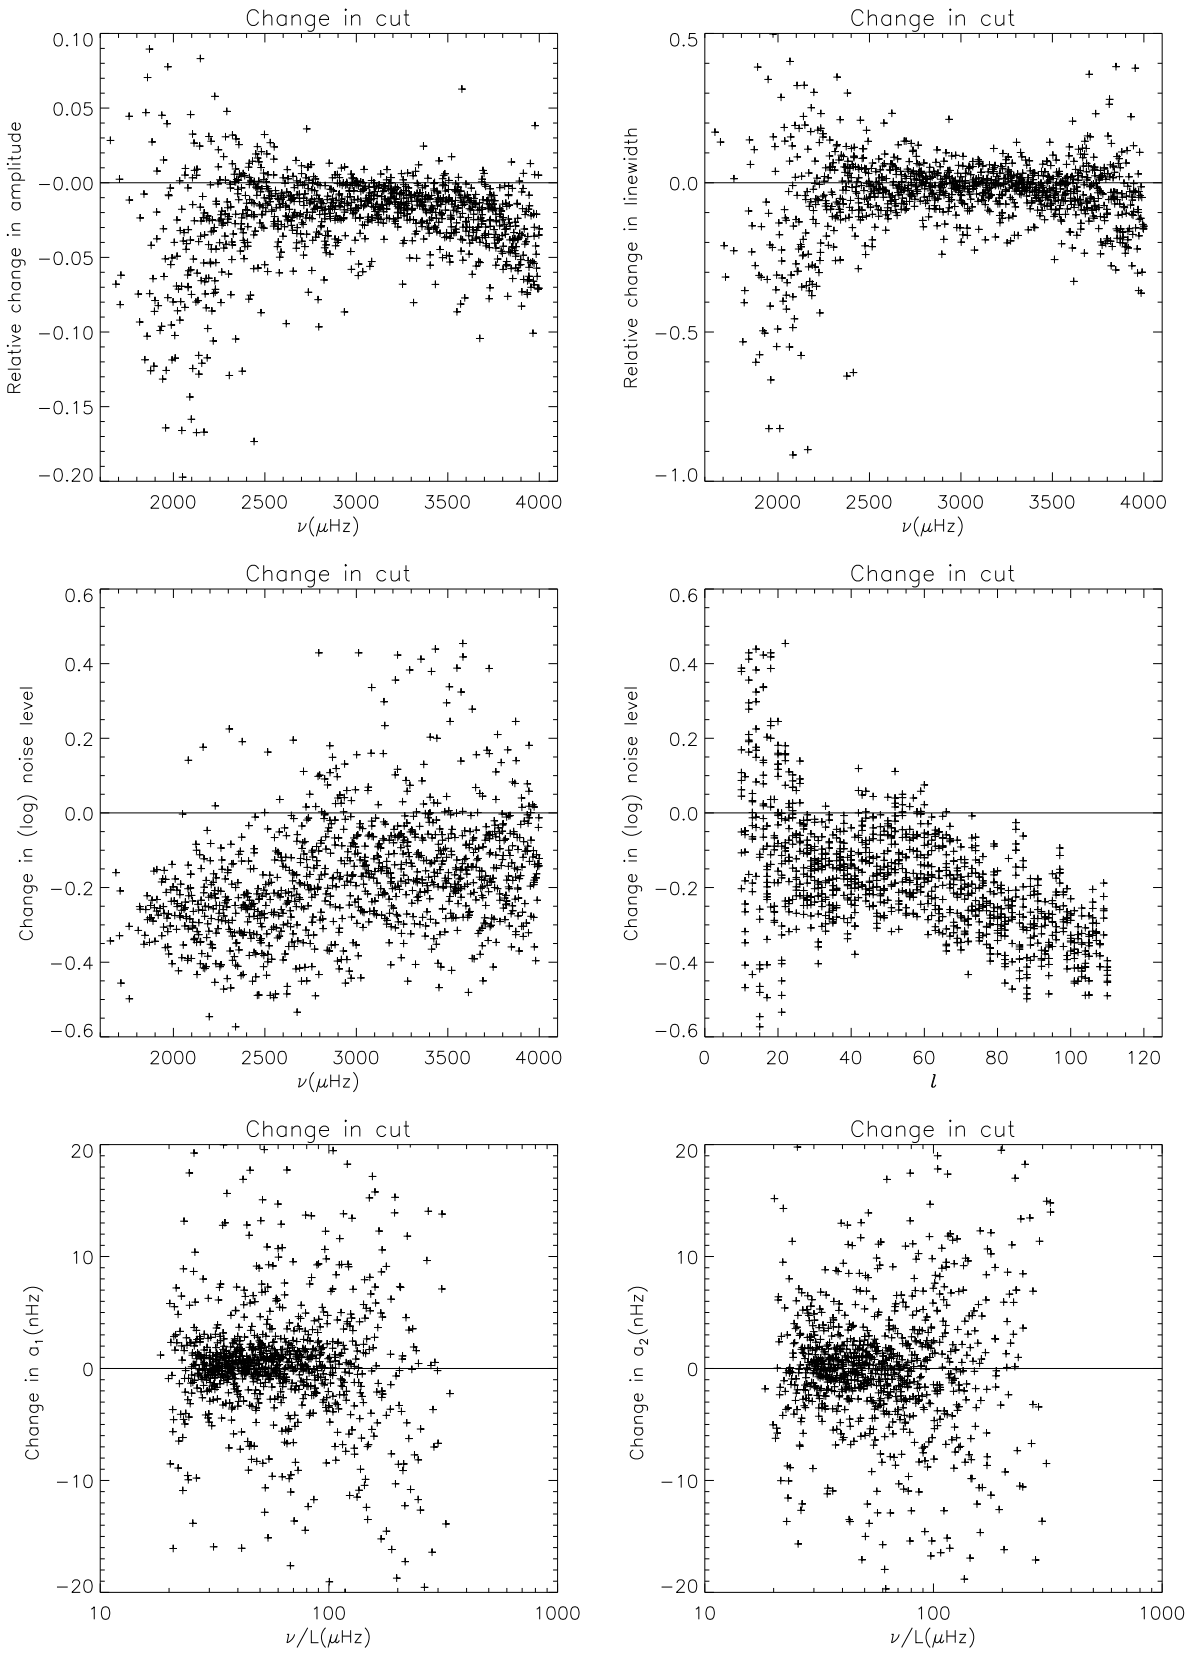

Figure 3.43: The change in various parameters caused by changing the cut.

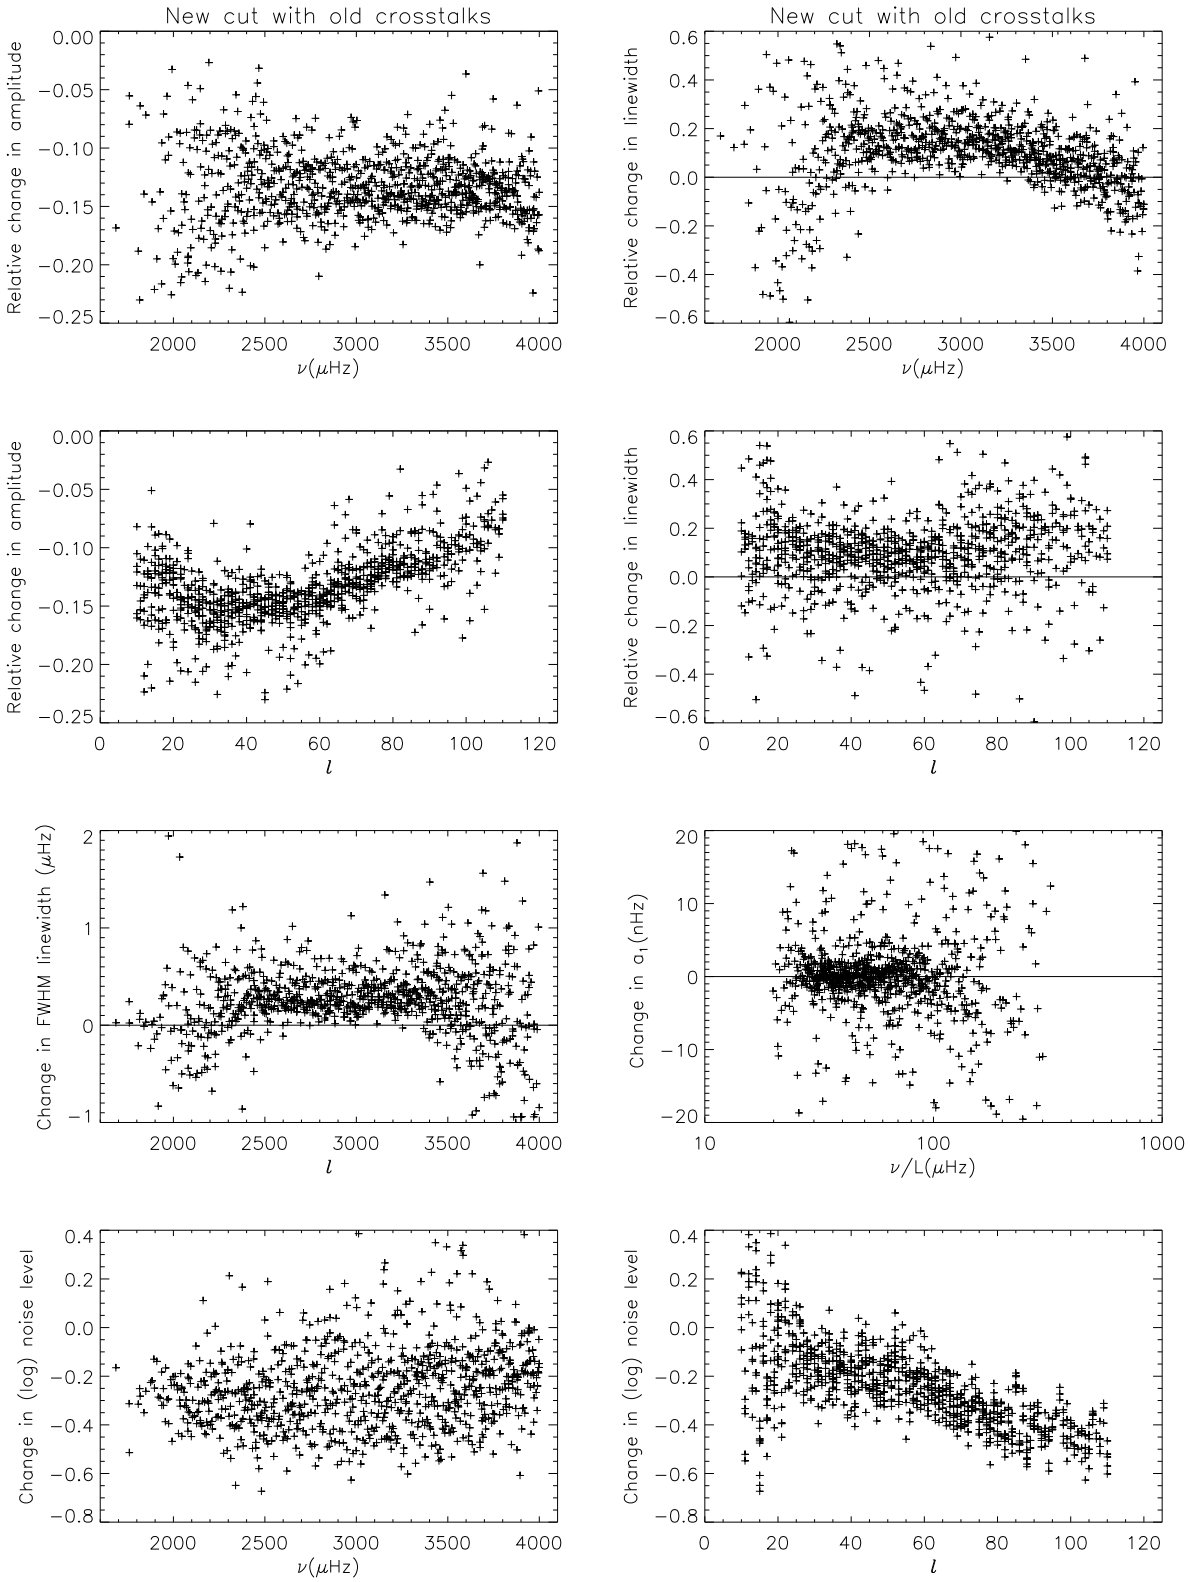

Figure 3.44: The change in various parameters resulting from reducing the time-series with the cut at 0.5R using crosstalks corresponding to a cut at 0.6R.

in two of these diagnostics are shown in figure 3.45. It is obvious that both diagnostics change considerably, and in a way very different from that caused by the incorrect PSF in the next section. Generally, one would expect this change to show up in the first diagnostic, as it affects the m crosstalks. Notice also that it affects low ls strongly and that there was no error in this diagnostic for low ls in, for instance the aj89 dataset, indicating that the cut was probably properly taken care of in the analysis of that dataset. Since the cut also affects the general dependence of the power on m, it also shows up in the second diagnostic. Notice that the effect is strongest at fairly low l, indicating again that this is not the cause of the incorrect diagnostics in the aj89 dataset. Also notice that the effect is much smaller than the one shown in the next section.

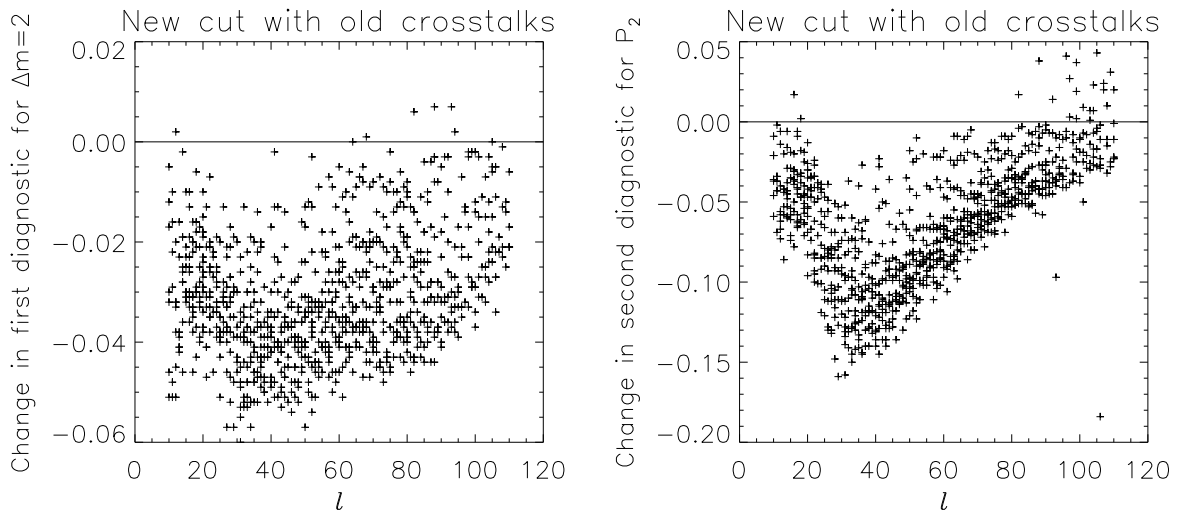

Figure 3.45: The change in a couple of the diagnostics caused by analysing the data assuming the wrong cut.

3.6.5 Change in the Assumed PSF

A known problem in the analysis is that of determining the PSF to use. For this reason it seemed natural to try to modify the assumed PSF. Unfortunately, it is expensive computationally to recalculate the necessary crosstalks. As the crosstalks for the aj89 run, which have been used for all the tests, have had the cut performed on the images, it is not possible to use crosstalks from another year, as the effects of the cut and the effect of the different PSFs would be mixed with each other. As a quick and dirty solution I analyzed the usual one month worth of data from aj89 with the crosstalks for both the aj87 and the pole87 runs. Recall that the aj89 data have been analyzed with an assumed PSF with $\sigma_x = 0.0160$ and $\sigma_y = 0.0144$. The pole87 PSF has $\sigma_x = 0.0195$ and $\sigma_y = 0.0160$, while the aj87 PSF has $\sigma_x = 0.0160$ and $\sigma_y = 0.0150$. Results are shown in figure 3.46.

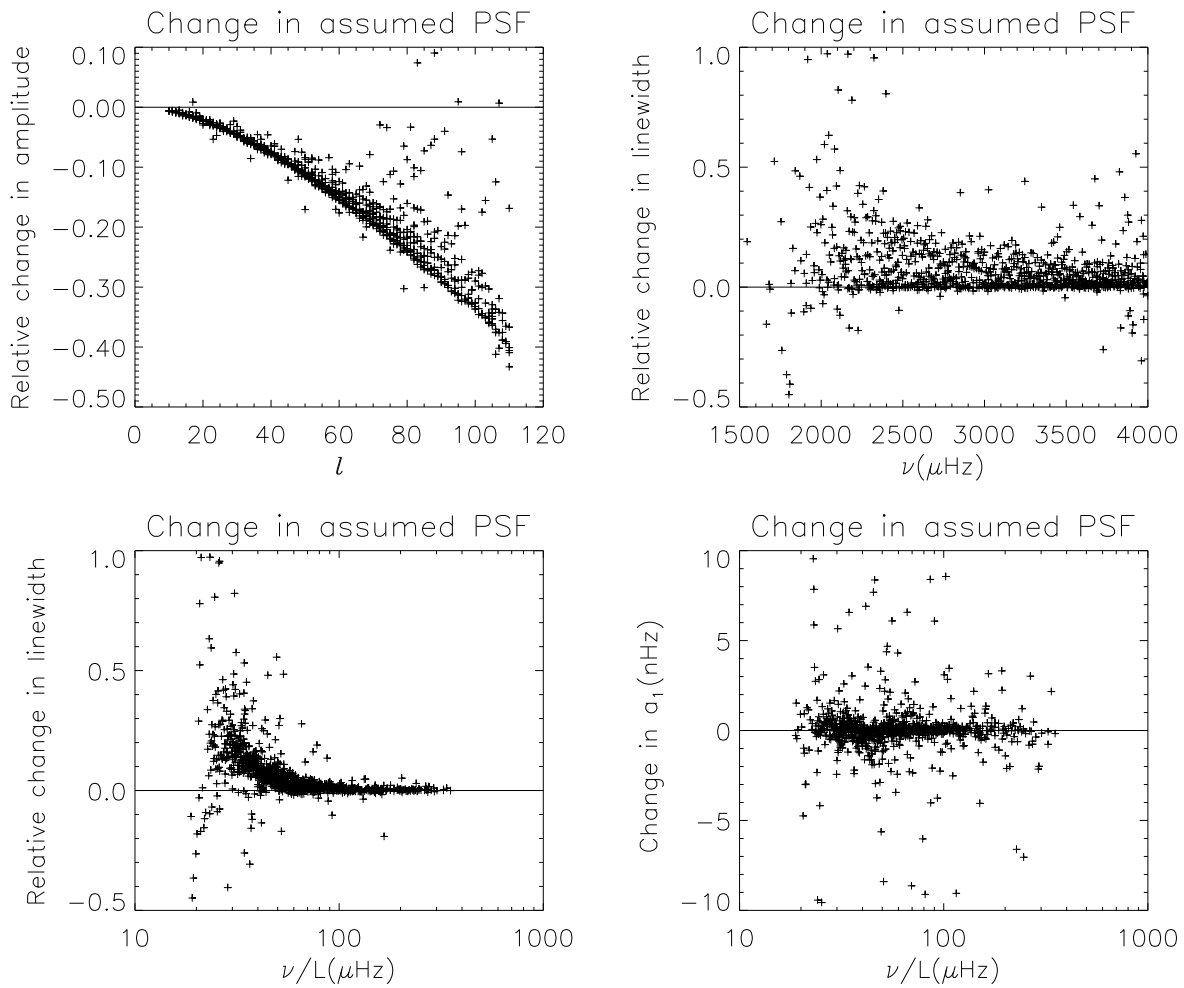

Figure 3.46: The change in the mode amplitudes and linewidths between the runs with the two different assumed PSFs.

A curious fact (not shown) is that the number of modes successfully fit is considerably different in the two cases and that it is generally the *low* l modes that are failing to converge. The cause for this has not been investigated.

From figure 3.46 it can be seen that the amplitudes change significantly as a function of l and that the linewidths change somewhat. Other parameters change by very small amounts. The changes in the linewidths are likely due to the usual changes in the crosstalks giving problems. The amplitude changes on the other hand are very large and systematic and not too surprising. The main effect of changing the assumed PSF is to generally change the sensitivities as the spherical harmonics get more or less washed out. As high ls have a higher wave number they are more affected than lower ls leading to the obvious trend with l. Since the difference between the two PSFs is mainly in the horizontal direction, high ms will be affected more than low ms at the same l, leading to a different dependency of the mode power on m. Given that the second diagnostic discussed in Schou (1993a, cf. App. 5) is supposed to test for exactly this problem, it has been shown in figure 3.47.

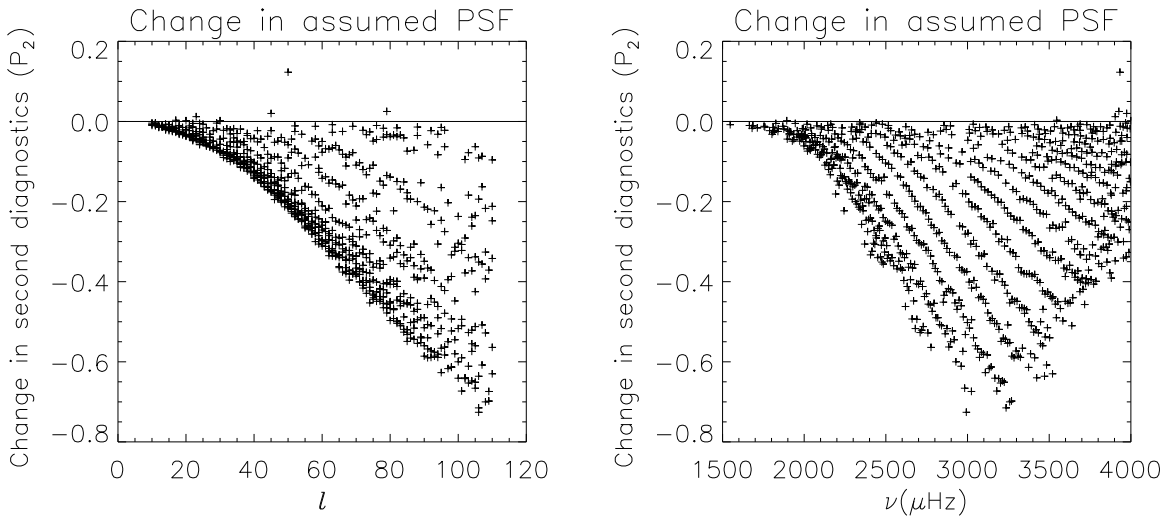

Figure 3.47: The change in the second diagnostic for the P_2 case caused by the change in the assumed PSF.

As can be seen the change in this diagnostic does indeed show a variation with l similar to the change in the mode amplitude. The frequency dependence is due to the noise covariance matrix being better estimated than the crosstalks and that the error is therefore largest when the S/N ratio is largest (that is, in the middle of the p-mode band).

This signature in this diagnostic clearly indicates the usefulness of these diagnostics. In this particular case the change was most likely caused by the change in the asymmetry of the PSF. A generally degraded PSF is likely to have less effect as it will degrade all ms at a given l by roughly the same amount. The first diagnostic changes in a way similar to the second, but by a considerably smaller amount.

To use these diagnostics to actually correct the PSF, further experiments are clearly necessary. Also note that this particular diagnostic relies on the assumption that the true power is independent of m (that is, the horizontal direction of propagation of the waves), which although likely to be true, remains to be proved.

3.6.6 Change in the Initial Parameters

While the sensitivity to the initial parameters for the target mode and the values for the other modes should be irrelevant when things are converged, it is nevertheless interesting to see how sensitive the resulting parameters are to these things. For instance, it indicates the magnitude of some of the problems one has in iterating the parameters and how sensitive the final results are to the iteration not being totally converged.

As the number of input parameters is rather large, it is difficult to test all possible perturbations, so only a few possibilities will be explored here. In particular, the 4 main parameters ν, A, w and a_1 have been changed for other modes than the target mode. a_1 for the target mode has also been changed to check for the stability of the fitting method.

In the first experiment, a_1 for the modes other than the target mode was increased by $5nHz$. Results are shown in figure 3.48. The only parameter to change significantly is a_1, where the change is a bump in the region where the temporal/spatial sidelobes of the neighboring l interferes, with the largest (negative) change where the interfering sidelobes cross the main lobe. That the change is negative where the sidelobes cross is not surprising as the program will try to compensate for the increased slope of the interfering mode by decreasing the slope of the target mode. It is interesting that the change is much smaller than the perturbation, by about an order of magnitude, which is convenient when it comes to iterating the fit.

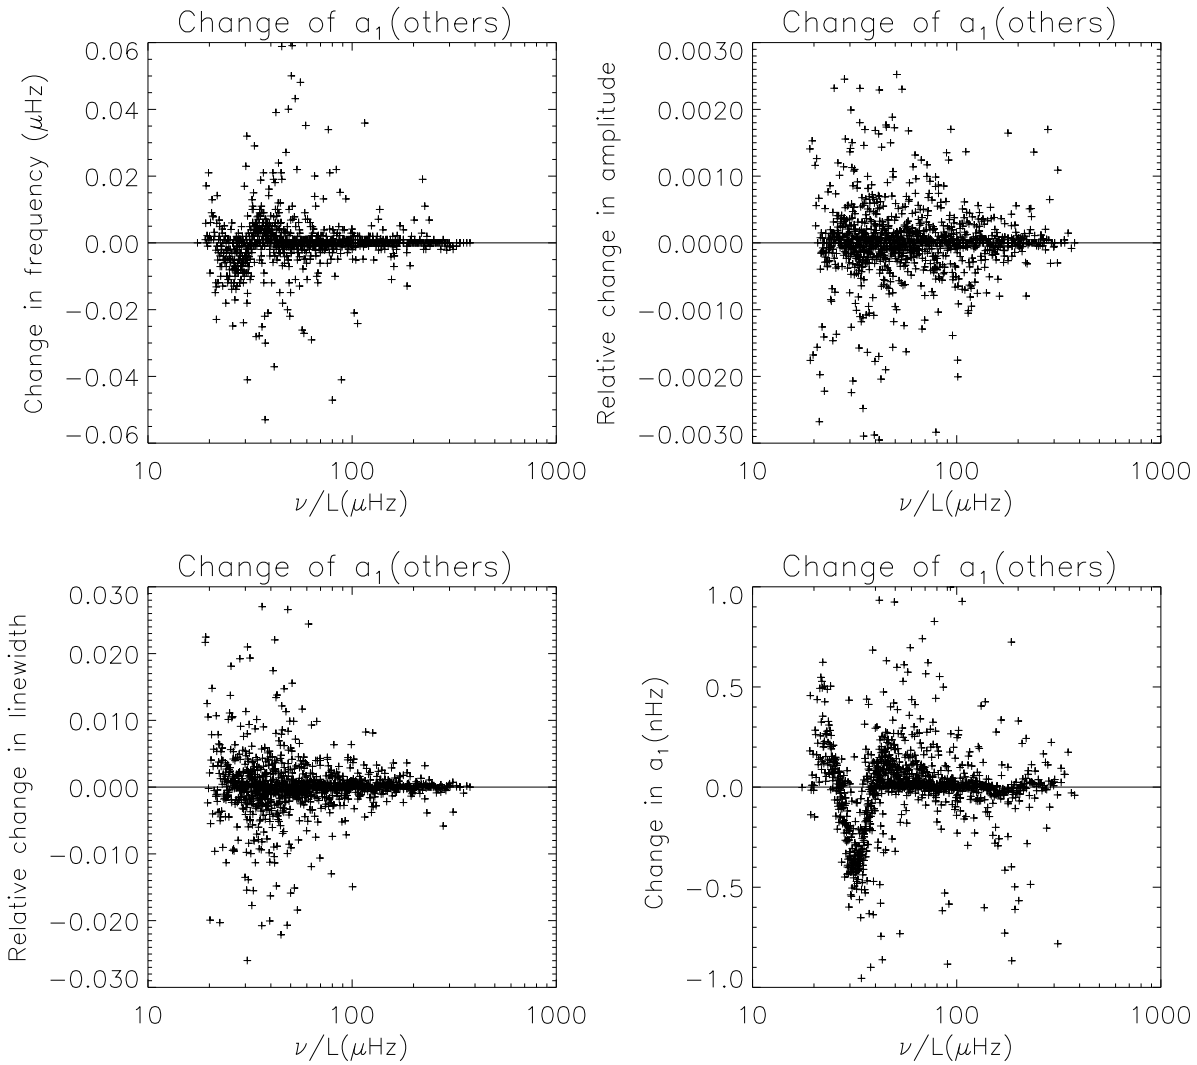

Figure 3.48: The change in various parameters caused by a change of $+5nHz$ in a_1 for the modes other than the target mode.

Notice, however, that the bump does not have the same shape as the original bump in a_1, and that the magnitude is very small.

The other parameters, including the other a-coefficients, change very little.

In the second test the amplitudes of the modes other than the target modes were held 20% higher than their expected values compared to the mode of interest. Results are shown in figure 3.49. Note that the level of the other modes was not left floating, but kept fixed with respect to the mode of interest. This change in the amplitudes is, in a sense, equivalent to a change in the crosstalks between ls, and thus also tests the sensitivity to those.

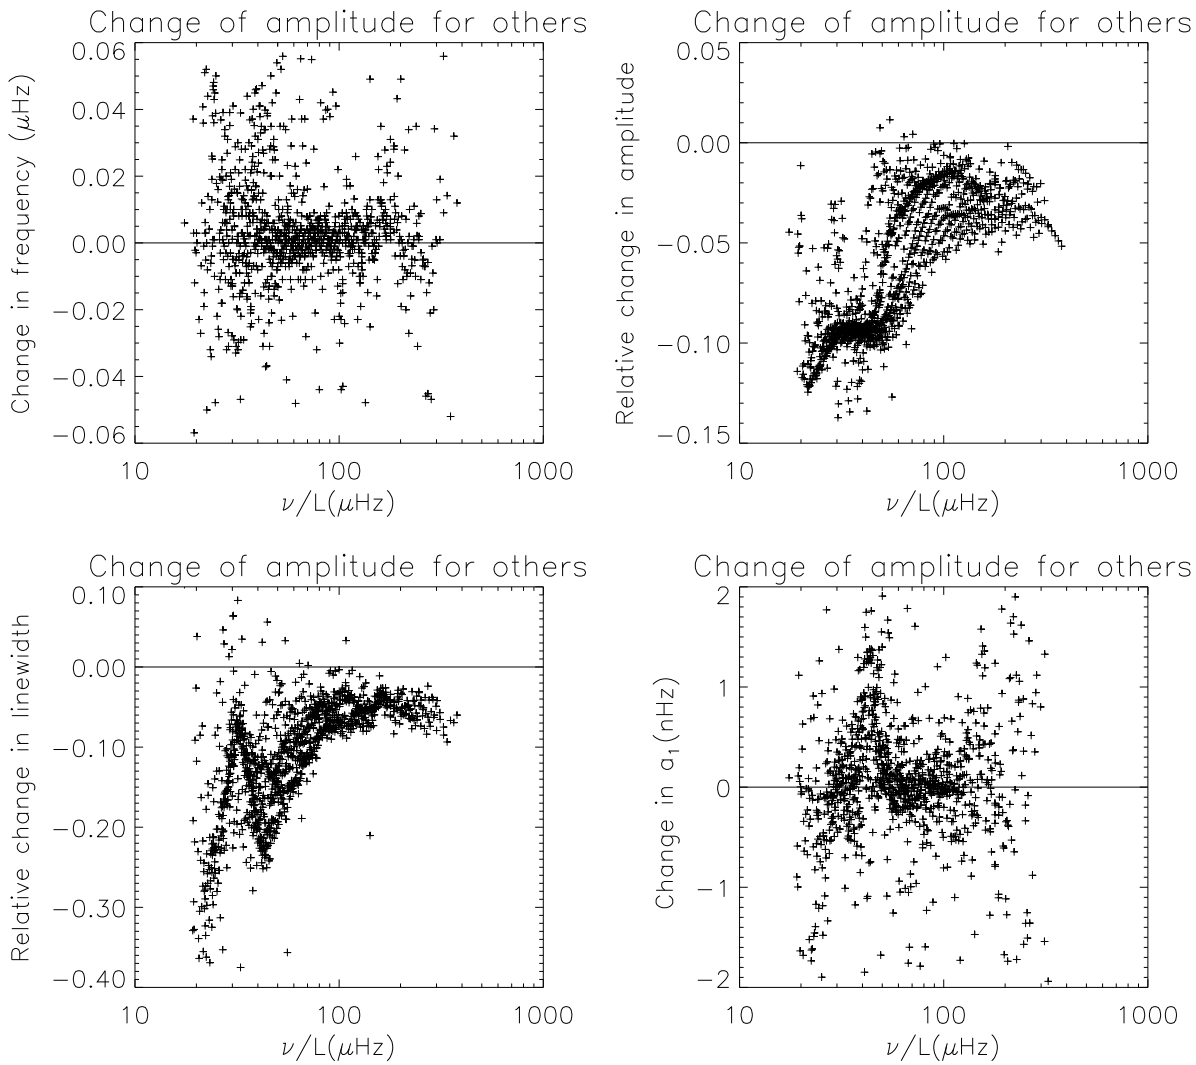

Figure 3.49: The change in various parameters caused by a change of 20% in the amplitudes for the modes other than the target mode.

Not surprisingly, the amplitude of the target mode becomes significantly lower for almost all modes. As the sidelobes of the other modes are off by a constant ratio compared to the target mode, it would be expected that the amplitude changes by an amount depending on the total power from other modes inside the fitting window. The amount of power in this window depends upon a combination of the size of the fitting window, the magnitude of the relevant crosstalks, and the spacing of the neighboring ls, and it is thus not surprising that the amplitude

change depends strongly on ν/L. The rapid change around $\nu/L \approx 50$ corresponds fairly well to the point where the temporal sidelobes for the neighboring ls are at the edge of the fitting window, that is when the spacing is 1 day$^{-1} + 5 * \text{HWHM} \approx 11.57\mu Hz + 5 * 1\mu Hz \approx 17\mu Hz$, where a typical FWHM linewidth of $2\mu Hz$ has been used, corresponding to $\nu/L \approx 50\mu Hz$. The scatter is likely due to the variable window size (given by the linewidth) and to a smaller degree the variable crosstalks.

The change in the linewidths is more interesting. It is also a strong function of ν/L and negative except for a few modes. The change is fairly large, of the order 10%. The reason for the change is probably that in trying to compensate for unexpectedly large interfering sidelobes, the program tries to squeeze the main lobe. Notice that the local minimum in (the absolute value of) the change around $\nu/L \approx 30\mu Hz$ corresponds to the point where the temporal sidelobes cross the target mode, as would be expected from this explanation.

a_1 changes somewhat, but not in a way resembling the bump, and again the magnitude is very small compared to the observed bump. Most other parameters do not change significantly, although there are small changes in some of the other a-coefficients.

Increasing the linewidths by 20% changes the amplitudes by a fairly small amount, of the order 1%, as shown in figure 3.50. The linewidths change considerably, of the order 10-20%. The maximum (negative) change occurs almost exactly where the temporal sidelobes cross the target mode, the cause being that the program compensates for the widened sidelobes by narrowing the main lobe of the target mode.

The a-coefficients again change by fairly small amounts.

It was decided to change the frequency by a constant fraction (0.25) of the FWHM linewidth instead of a constant amount, because a change of several linewidths would cause the interference with the target mode to change considerably. For this reason the changes have been scaled by 4/FWHM, such that they show the extrapolated change for a frequency change of $1\mu Hz$. Results are shown in figure 3.51.

In this case the frequencies change substantially, unlike for the other changes. As expected, the frequencies change by a large (compared to the perturbation) negative amount around the place where the temporal sidelobes of the neighboring ls cross the target mode. Again it seems like the program is trying to compensate for an incorrect position of the interfering modes by moving the target mode in the opposite direction. The fact that the feedback (the ratio of the change to the input perturbation) is so large makes it difficult to make the iterations converge.

a_1 also changes by a substantial amount around the crossing point, again leading to some problems for the iterations. The other a-coefficients change by various amounts.

Tests such as these, where the values of the parameters for the modes other than the target mode are changed, also test the sensitivity to the temporal/spatial sidelobes. If both of these were absent the values of the parameters for other modes would be of no relevance. Given the different characters of the changes they should supplement each other. On the other hand, these problems should be insignificant after a sufficient number of iterations.

To check that the fits are actually stable the initial estimates for a_1 for the *target* mode

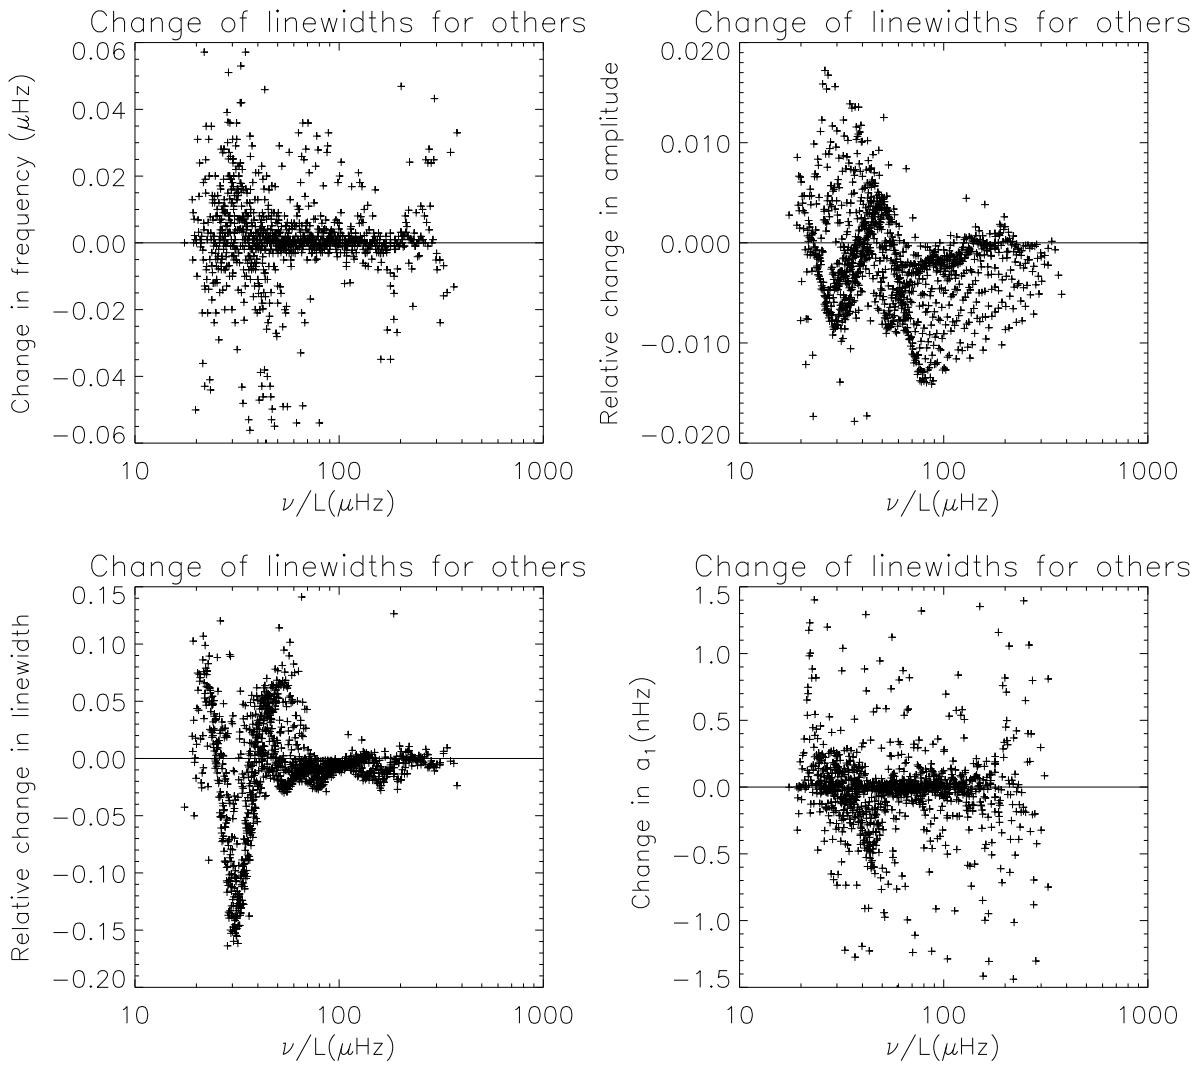

Figure 3.50: The change in various parameters caused by a change of 20% in the linewidths for the modes other than the target mode.

were increased by $5nHz$. The results are shown in figure 3.52. The changes are generally very small except for a few unstable modes.

A reason to be concerned about changing especially a_1 is that the position of the window used for the fitting is determined by the initial parameters. For l of 100, say, the change in the position of the window for $m = 100$ is about $0.5\mu Hz$, which is small but not totally negligible compared to the window size.

3.6.7 Conclusion

Although most of the changes were as expected, some were not, and for others some of the details still need to be worked out. In particular, the changes in the various parameters, when the assumed radius is changed, is somewhat puzzling and could indicate problems with the

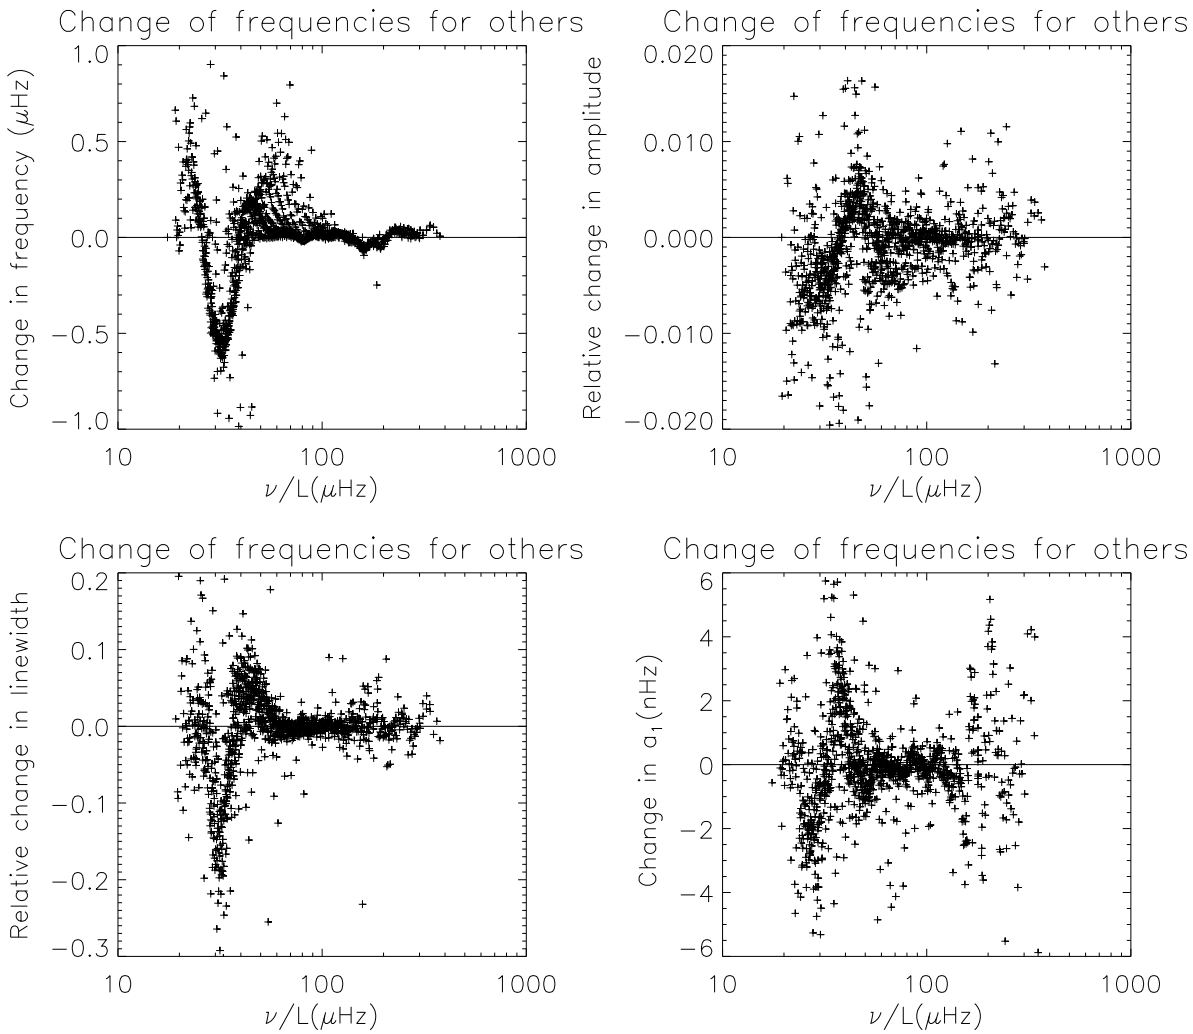

Figure 3.51: The change in various parameters caused by changing the frequencies for the modes other than the target mode by 25% of the modes linewidths w. The changes have been multiplied by $4/w$, to show the scaled change corresponding to a constant frequency change of $1\mu Hz$.

radius determination. For this problem it is at least convenient that even a very large change in the radius would be unable to explain the bump in a_1.

The change in a_2 when the images are squeezed is also a reason for some concern, as the magnitude is very large and the shape very close to that of the bump. As will be explained in the section devoted to the bump, this could offer some clues as to the cause of the bump.

Of course, the bump is not the only thing one can learn about doing these tests. In particular, they give valuable information about the possible systematic errors in all the parameters from wrong assumptions.

An area where more work would be useful is in the design and interpretation of the various diagnostics. It would be useful if it was possible to tell what assumptions are likely to be wrong if a particular diagnostic shows some particular feature.

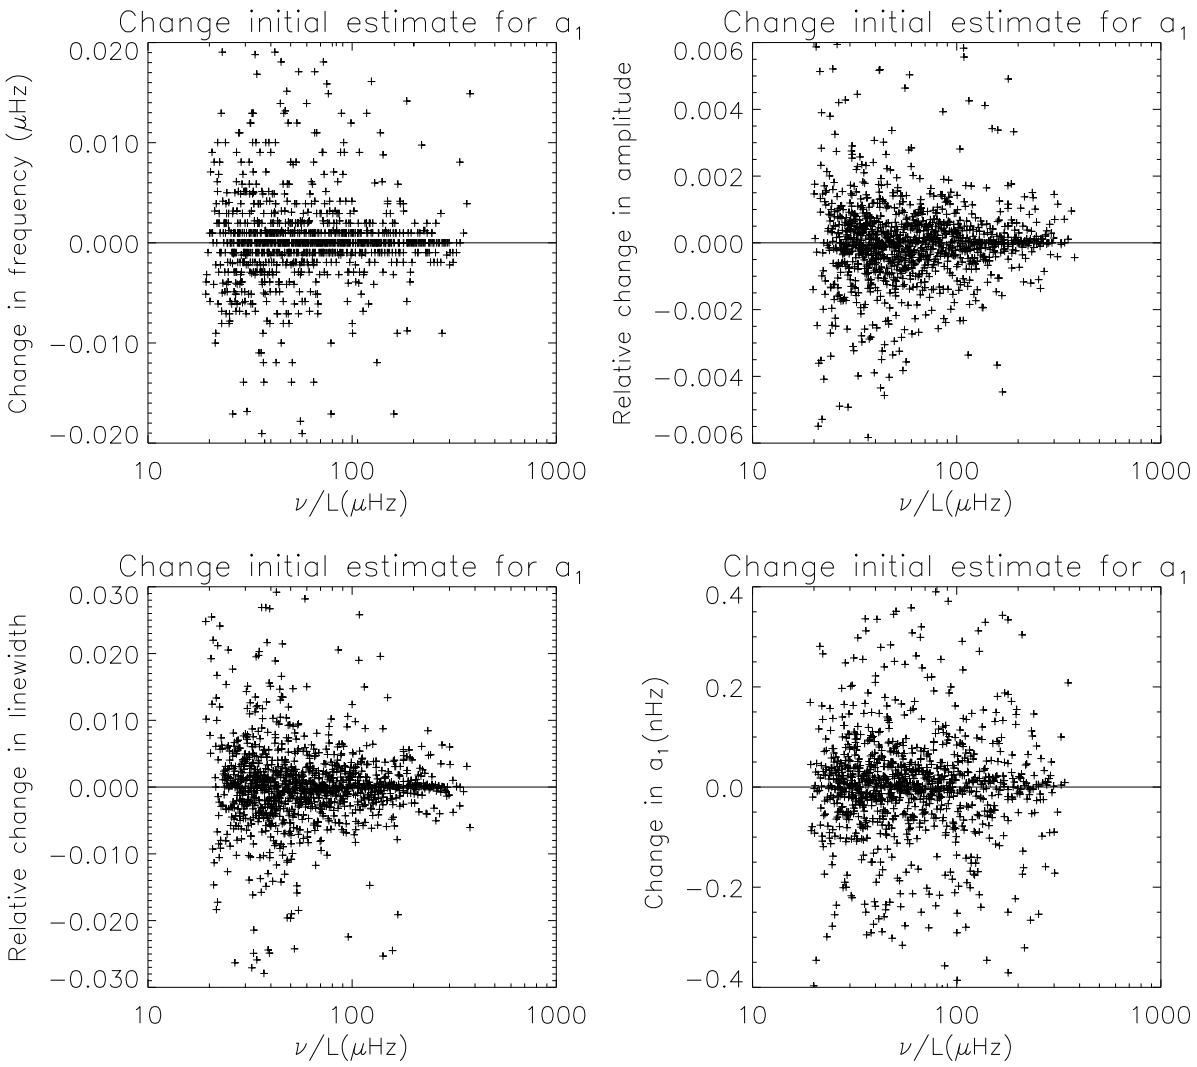

Figure 3.52: The change in various parameters caused by a change of $5nHz$ in the initial estimate for a_1.

Another way to look for possible problems in the analysis is to analyze artificial data, as described in Schou and Brown (1993, cf. App. 4) and Schou (1993a, cf. App. 5). The tests described in this section have in a sense been differential, in that the change in various parameters caused by different perturbations have been studied. Using artificial data on the other hand, is an absolute test.

3.7 The Bump

Since the bump in a_1 is probably the most surprising feature in the results from the data reduction, I will try to summarize the properties of the bump and describe some of the possible explanations in this separate section. Some of the plots here are copies from other sections, but have been collected here for clarity.

3.7.1 Description

The bump itself has been shown earlier in the sections in the individual runs. A combined plot of a_1 for the aj87, pole87, and aj89 runs is shown in figure 3.53.

The features of the bump that have been attracting the attention generally fall in the following categories:
The size
The location
The shape
The variability/lack thereof

While the size of the bump is only a few percent of a_1, and is therefore, in a sense, not very large, it is highly statistically significant. The magnitude is considerably larger than both the estimated standard errors on a_1 and the scatter around the bump. It is also the dominant feature in a_1. As a matter of fact, the effect is large enough to be seen in the power spectra, as illustrated in Bachmann, Schou, and Brown (1993, cf. App. 6). It can be seen that the slope as a function of m varies as a function of frequency (or equivalently n), as one goes across the bump. Note that as the same l was used for all the plots, the crosstalks are the same. As shown, the effect does not appear to be due to some subtle problem in the Fourier transform analysis program, although, as explained later, it could be caused by some incorrect assumption in that part of the reduction relating to previous parts of the reduction or the underlying physics.

Another indication that the problem is not in the analysis part of the code is that an independent analysis performed by K. Bachmann (see Bachmann, Schou, and Brown 1993, cf. App. 6) on the same data (the reduction was the same up to the time-series output by the YLM code) using a simpler analysis code, gives very similar results. a_1 from his code has been shown in figure 3.54 for comparison. While it is clear from 3.54 that the two methods do not give identical results, the results are very similar. Also using various variations of this other method does not seem to affect the results significantly, at least not when it comes to a_1.

The location of the bump is probably the single most important fact making it suspect. The point on the bump where a_1 crosses the value above the bump (in ν/L) is right at the point where the first temporal sidelobes of the neighboring ls are on top of the main lobe of the target mode. As discussed in section 3.6, the distance between neighboring ls at a given n is mainly a function of ν/L. As discussed in chapter 2, this is also the parameter one would expect a real feature in the solar rotation to show up as a function of. While this does not mean that the bump is an artifact of the analysis, it is highly suspicious, as this is exactly the location where problems were expected beforehand.

The shape of the bump is obviously connected to the location. The interesting fact is that the zero crossing is at the position just discussed and that the bump is roughly antisymmetric around this point. The point of this was obvious in the discussion of some of the possible explanations in the section (3.6) on the tests.

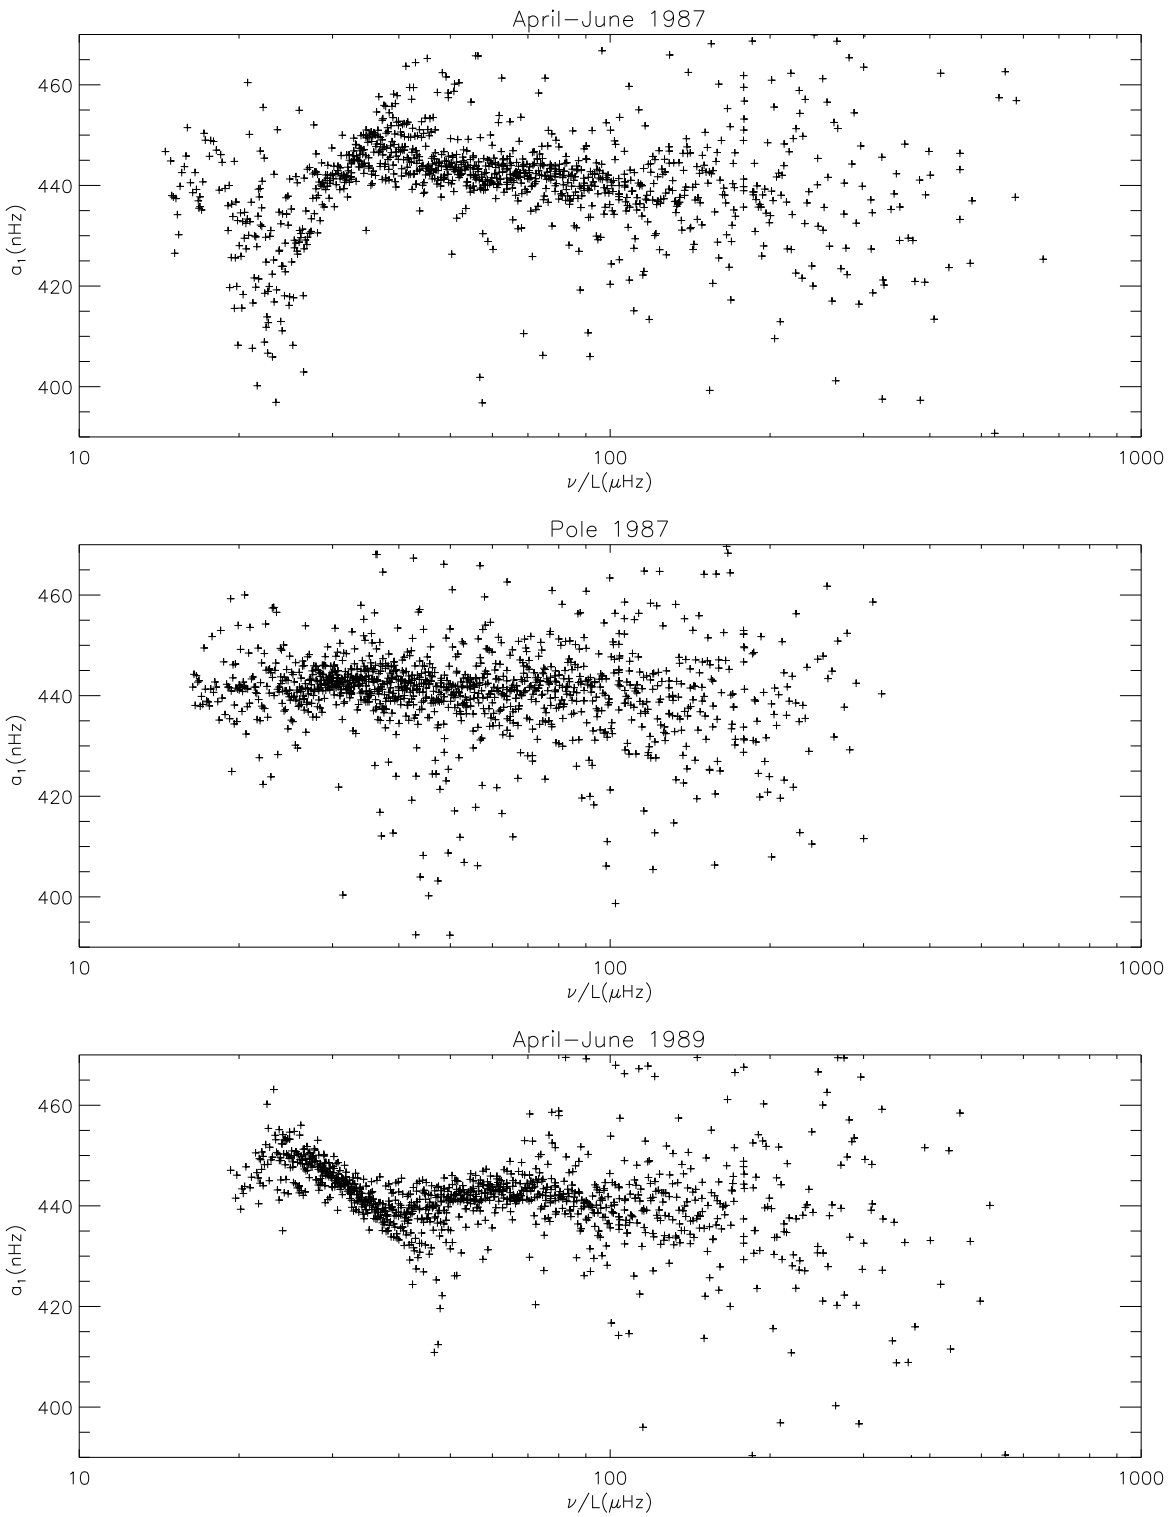

Figure 3.53: a_1 for the aj87 run (top), POLE87 (middle), and aj89 (bottom). The unweeded modeset was used for the aj87 set. Numbers from K. Bachmann were used for the pole 87 dataset. Notice that the exact modesets used for the three plots are different. Also notice that the POLE87 run is shorter than the other runs, leading to a larger scatter.

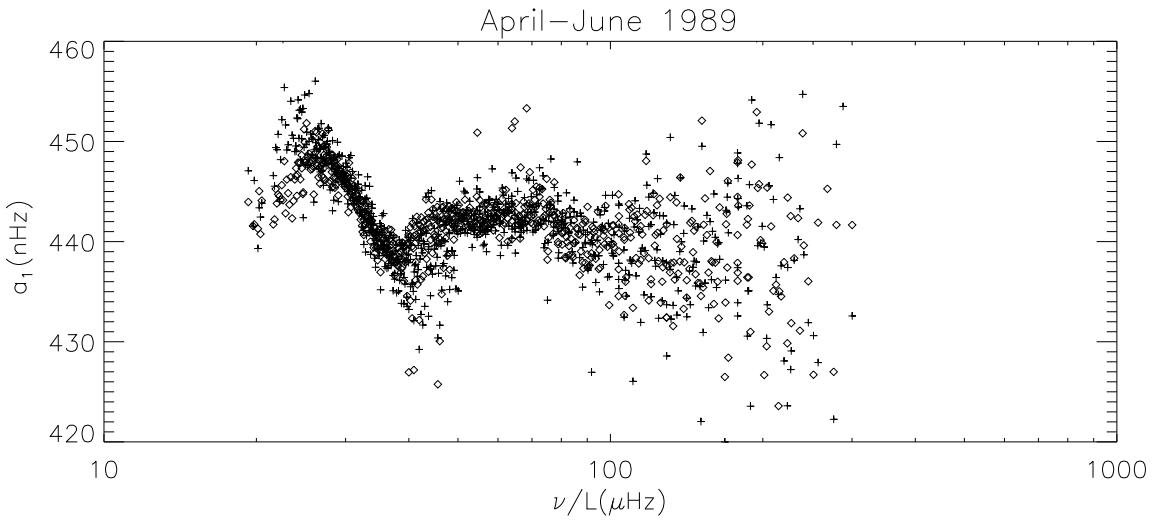

Figure 3.54: a_1 as a function of ν/L. Crosses indicate results from the analysis program used here, diamonds results from an analysis of the same data by K. Bachmann using a different analysis method. Only modes common to the two sets were used.

Also note that a_1 is not just a function of ν/L, as illustrated in figure 3.55. It is clear that the magnitude depends on the frequency (or possibly the linewidth). If the bump is indeed caused by some analysis effect it is not surprising if there is a dependence on the linewidths (as they determine how close modes must be to interfere), on the other hand, a real physical effect could be related to the frequency. Given the quality of the data it is difficult see which of these possibilities is the case.

It is also interesting that the shape of the bump is almost identical to the bump induced in a_2 by squeezing the solar image (see figure 3.56). Notice the striking similarity between figures 3.54 and 3.56. It is highly suspicious that the bump looks so much like a known systematic error, even though the error is in a_2 (which shows no bump) and not a_1. As discussed in the section on the tests performed, any geometric distortion of the images combined with the 1 day^{-1} is likely to only lead to problems for the even a-coefficients.

A final but very important point is the variability of the bump. In the aj87 run which corresponds to the aj89 run in terms of the time of year (which means that it, among other things, had the same B-angle), the bump has roughly the same magnitude and shape as in aj89, but the opposite sign. In the run from November-December 1987, there is no bump (see figure 3.53). On the other hand, it is interesting that the change between the three months of the 1989 run are small, as shown in figure 3.57.

Another interesting fact concerning the variability is that none of the tests performed (see the section on those) were able to change the size significantly, even with, in some cases, very significant perturbations.

Clearly, any candidate explanation for the effect should be able to explain at least most, and ideally all, of these properties.

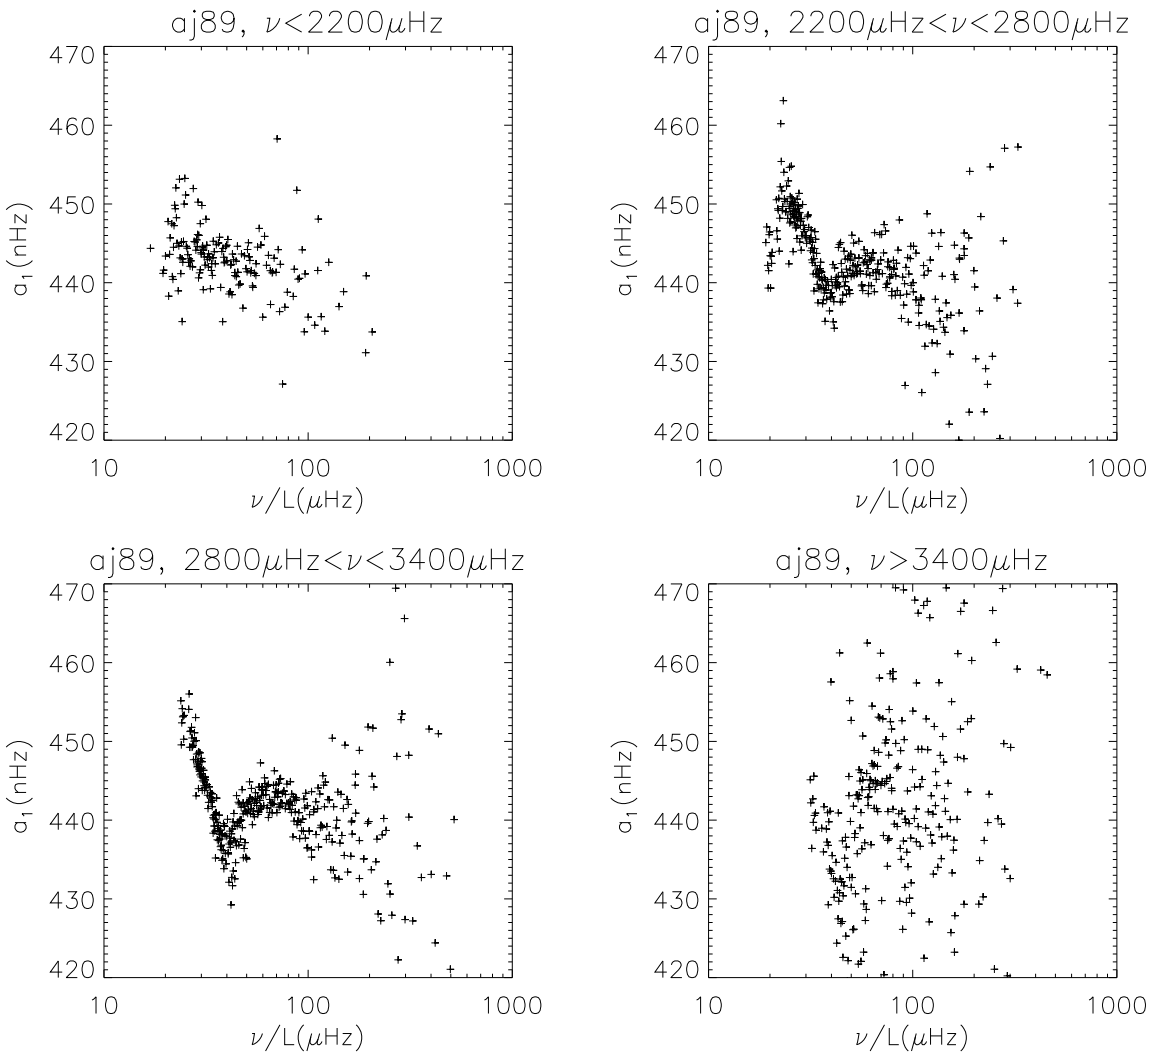

Figure 3.55: The bump for 4 different frequency ranges. Notice the increasing scatter with frequency due to the increasing linewidth.

3.7.2 Possible Explanations

The number of suggested explanations for the bump is very large. Most of the explanations have been rejected, either because they obviously would not work or because tests showed that they did not work. Although some of the explanations did not work, they will still be described here, as the reasons for some of them not working are fairly general. Some of the general explanations were also described in the section on the tests. A number of the possible explanations have not been tested yet, either because the tests are difficult to perform and not sufficiently promising or simply because new ones keep coming up. In particular, the tests requiring a rerun of the whole code from the raw images are very costly in both computing time and real time.

The simplest explanation from the point of view of the analysis is obviously that the bump is real and due to the internal solar rotation having a bump. If the splittings are taken as they are, and inverted, one gets an inverted rotation rate as shown in figure 3.58. While this

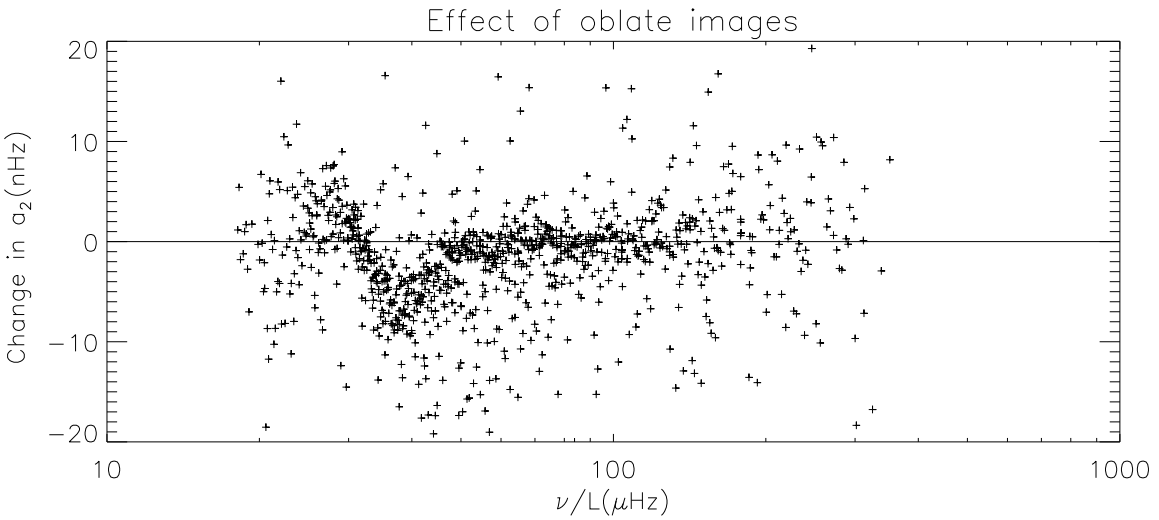

Figure 3.56: The change in a_2 induced by squeezing the solar images. See the section (3.6) on the tests performed for details.

inverted rotation rate seems rather extreme due to its variations with radius, there are no good physical reasons why it is impossible. On the other hand it seems to be such an unusual result that one is forced to consider other less drastic solutions before accepting it. Also the temporal variation would seem to indicate very large redistributions of angular momentum on fairly short timescales.

Another easy solution, from the point of view of the analysis, is that the bump is real, but due to some other physical mechanism than the solar internal differential rotation. However, since the effect only affects odd a-coefficients, it has to 'know' about the fact that the sound waves making up modes with $+m$ and $-m$ propagate in opposite directions. This rules out almost all perturbations to the structure of the Sun, such as sound speed dependencies on latitude. The only thing that comes to mind violating the east/west symmetry is the rotation. Apart from the direct effects of the rotation on the mode frequencies already discussed, it is difficult to think of a secondary effect of the rotation with this magnitude. Again, there are no hard physical reasons why no such effects could exist, but caution seems appropriate also in this case.

Most of the solutions thought of have, however, been in terms of various neglected, ignored or improperly treated properties of the oscillations, the instrument or some part of the data analysis. Many of these proposed solutions have been discussed in section 3.6. Most were based on there being some geometric problem, but as pointed out earlier in the section on the tests performed, such problems are most likely to only affect the even a-coefficients, as was indeed seen in the experiments. If the bump is due to a combination of temporal and spatial sidelobes, one would also expect the the magnitude of the temporal sidelobes to affect the size of the bump, but as shown in connection with the test where the window function was degraded, this does not seem to be the case.

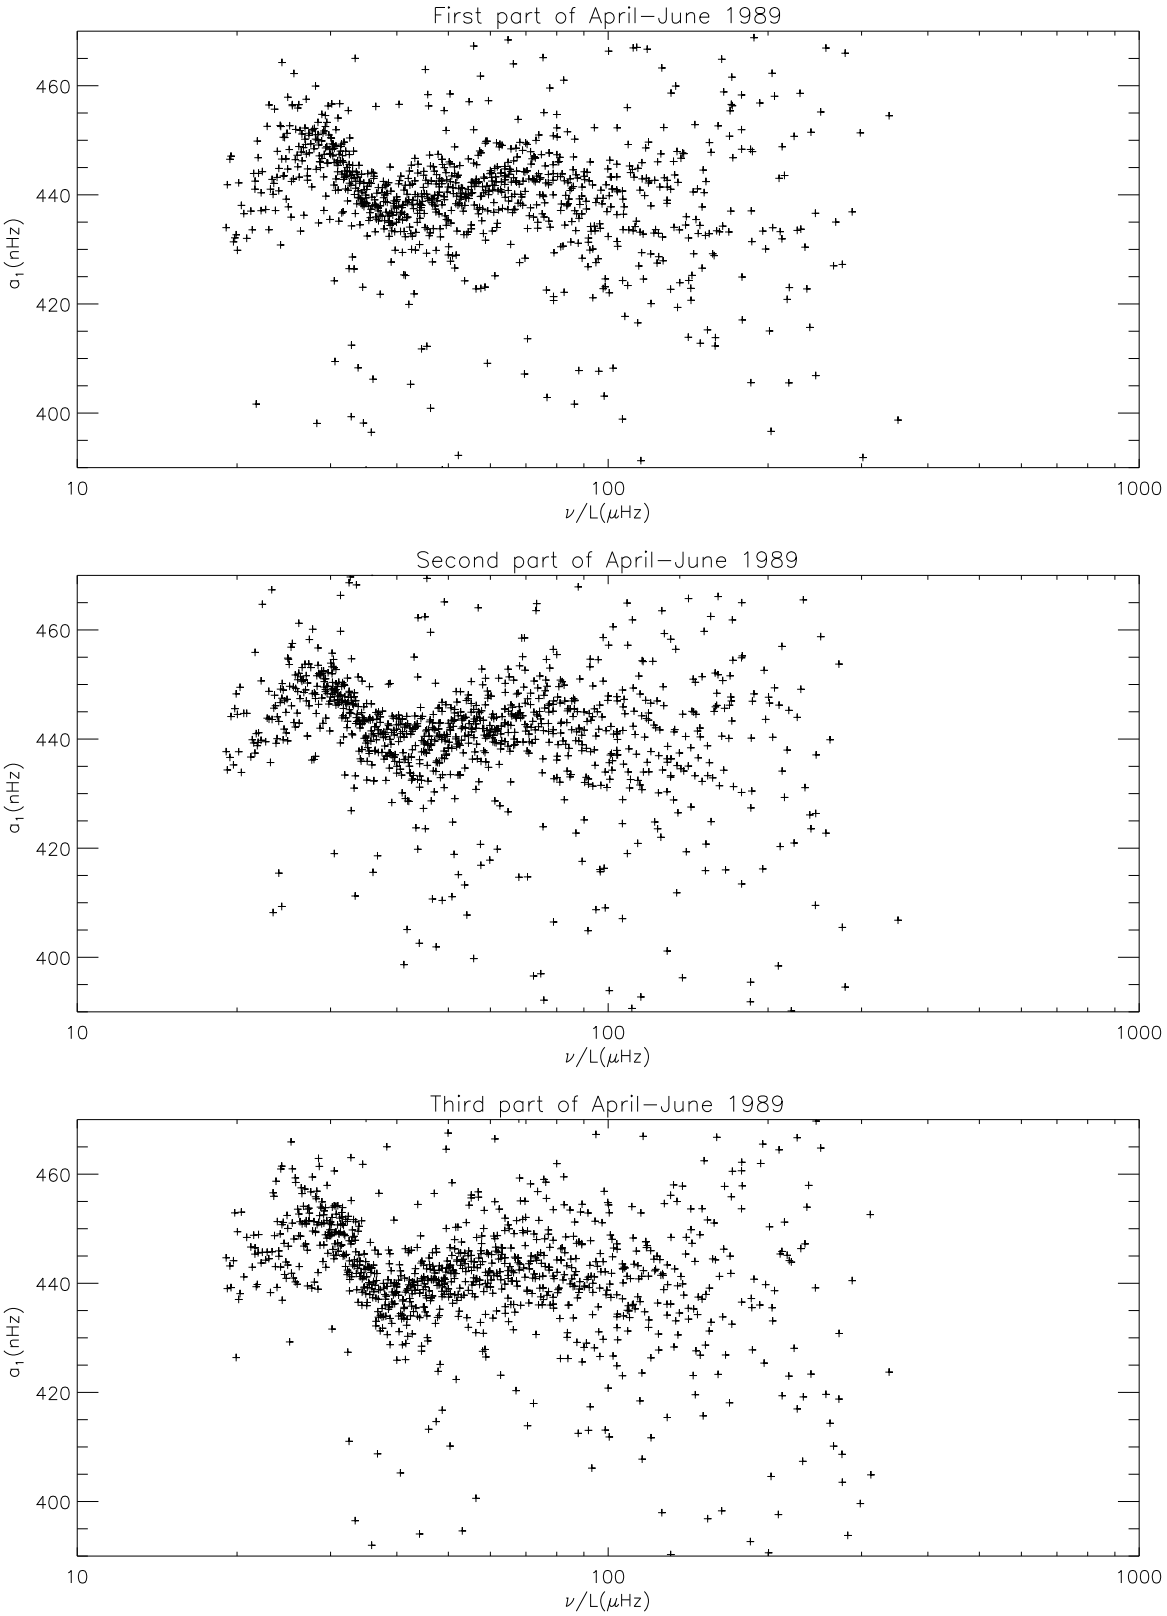

Figure 3.57: a_1 for each of the 3 subparts of the aj89 dataset used in section 3.6.2.

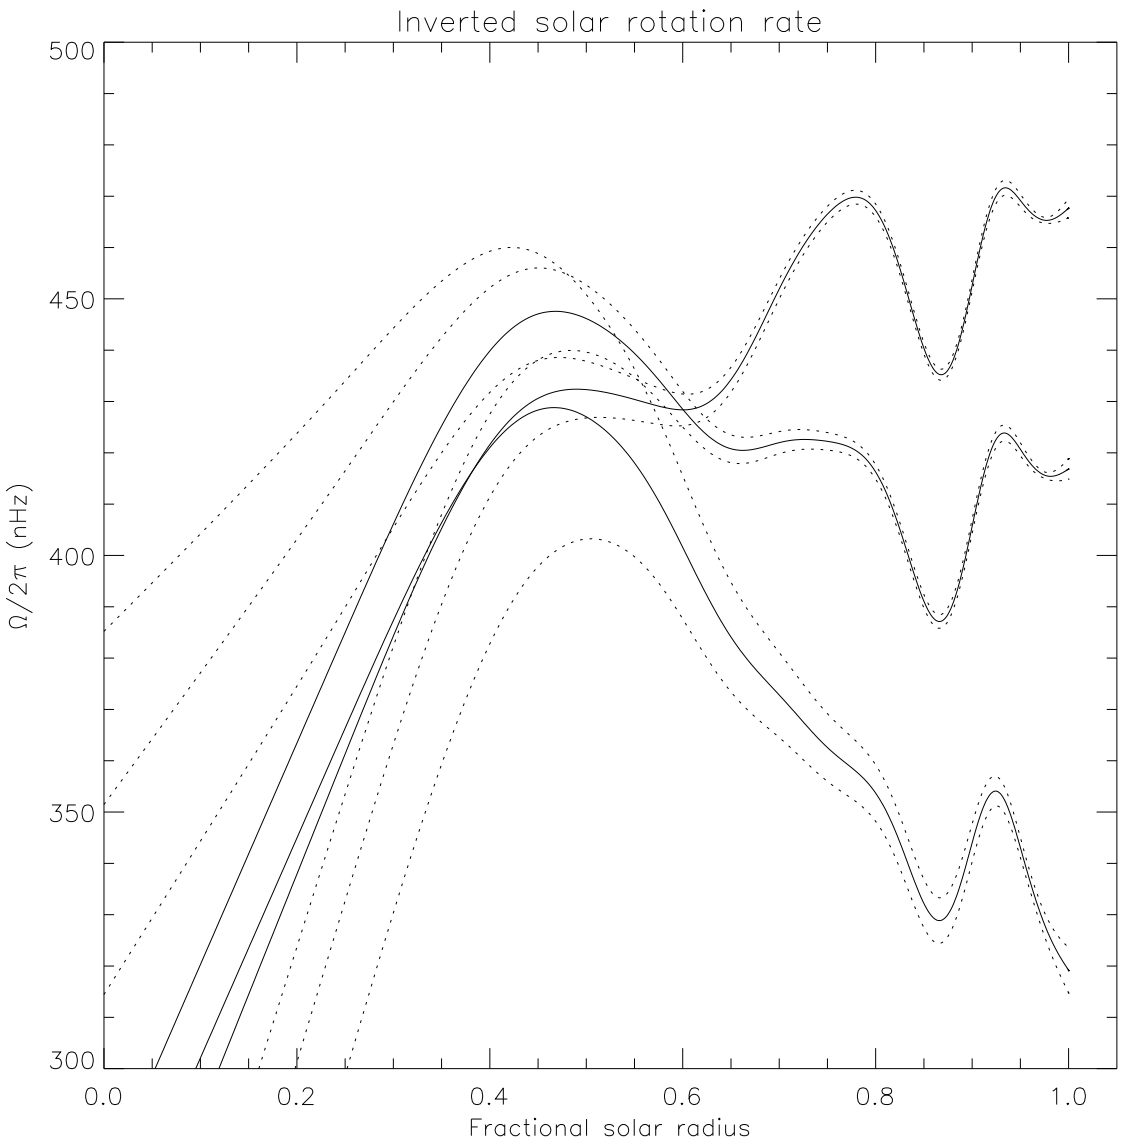

Figure 3.58: The result of inverting the 1989 splittings. The results are shown at 3 latitudes: 0°, 45° and 90°. A regularized linear least squares method with second derivative smoothing, and $\mu = 10^{-6}$ was used for each of the one-dimensional inversions. The solid lines are the inverted rotation rates, the dotted lines indicate $\pm 1\sigma$ formal errors. The inversion is somewhat oversmoothed, in particular close to the surface. The inverted rotation rate inside of $\approx 0.4R$, is not reliable, as it is mainly given by the regularization.

The fact that the bump in a_1 looks so much like the bump in a_2 induced by squeezing the solar image, could indicate that the cause of the bump is some sort of squeezing. However, as this a geometrical distortion, it is likely to affect only even a-coefficients, as is indeed observed. What would seem to solve the problem is some effect that will distort $+m$ and $-m$ the opposite way. Clearly, a geometrical distortion in, say, the optics is unable to do this as the eigenfunctions for $+m$ and $-m$ look the same at any given instant in time and are hence affected in the same way.

There is one fairly well-known, but generally neglected physical effect with the desired properties, namely the so-called 'Woodard effect', that was pointed out in Woodard (1989). This effect is based on the fact that as a wave changes latitude, it sees a gradient in the effective horizontal sound speed due to the differential rotation. This gradient will refract the wave, making it turn slower or faster in latitude depending on the sign of the gradient compared to the direction of propagation. As the modes with $+m$ travel in the opposite direction of those with $-m$ this leads to the one set being distorted in the opposite direction of the other. Unfortunately, Woodard estimated that the effect should be negligible for modes in the l range of concern here. Recently, however, we have become concerned that the effect might not be quite as negligible as originally thought. As all the other effects previously considered, this effect is unable to explain one of the properties of the bump, in this case the time variability. Why does this occur with different signs in spring 1987 and spring 1989 and why not in the fall of 1987? The latitude dependence of the rotation (which is reflected in a_3) on which the effect depends, appears to be almost the same for the three periods.

This effect is also unpleasant in other ways. The assumption that $+m$ and $-m$ have identical spatial eigenfunctions is built into many parts of the analysis, making it very difficult to take into account. Given that the effect is also unable to explain all the properties of the bump, we have not yet taken the time to test it. The proper way to take it into account would be to use different masks for $+m$ and $-m$. A quicker and dirtier way is to correct only the assumed crosstalks in the analysis program, which should work well if the effect is not too large, that is, if the peak of the power distribution is not moved too far away in l and m (less than 1 in l and maybe a bit more in m, say).

Notice that this effect is different from most of the other effects considered in that it is due to a known physical effect on the Sun and not some suspected effect or an instrumental/analysis effect.

Other explanations considered include temporal variations of the PSF and the effects of this on the temporal sidelobe spectrum. It is unlikely that this is the explanation though, as changing the magnitude of the sidelobes had no effect.

There are also some concerns as to the possible effects of the known geometrical distortion of the images. In particular, it appears that the images are somewhat oblate, and that the direction of this oblateness rotates through the day. It also appears that there are some problems with the centering of the images, possibly connected to the oblateness, leading to the 'centered' images drifting systematically during the day. Also the fact that the images would occasionally hit the sides of the detector array, was cause for some concern. These effects were tested by K. Bachmann using his code, and the effects on a_1 were small. Nonetheless, this problem has been corrected, and will be used for the analysis of the remaining datasets and the reanalysis of the datasets considered here.

Again, it should be emphasized that although this search for systematic errors was, to a large extent, motivated by the appearance of the bump, it should be useful for estimating the size of possible systematic errors in general.

3.8 Comparison of the Different Years

In this section some of the differences and similarities between the different years will be discussed. As a whole section has been devoted to the bump, it will not be discussed here directly, but only if it seems to have a connection to some other features. Some of the properties of especially the 1989 dataset have been discussed in Bachmann, Schou, and Brown (1993, cf. App. 6).

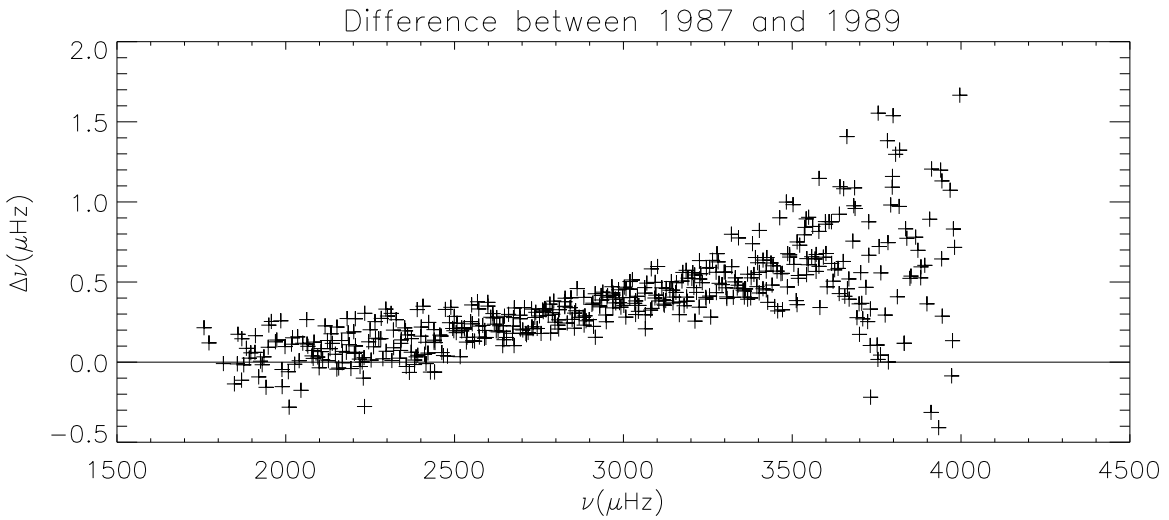

Figure 3.59: The frequency change from April-June 1987 to April-June 1989 as a function of frequency. To reduce the scatter, only modes with $20 \leq l \leq 60$ were used.

As has been pointed out by others (Libbrecht and Woodard 1990; Woodard et al. 1991), the mode frequencies seem to change from year to year, apparently related to the 11-year solar cycle. Comparing the aj87 and aj89 runs, a change similar to that reported by Libbrecht and Woodard is seen, as is obvious from figure 3.59. The change appears to be mainly a function of mode frequency rather than the turning point depth (which is related to ν/L). For this reason the frequency change has been plotted versus frequency.

Unfortunately, Libbrecht does not have any data from 1987, making it difficult to compare the details of the change. Libbrecht's 1989 dataset was, however, taken at almost the same time as ours (March to August 1989) the frequency differences between those frequencies and those from the present analysis are shown in figure 3.60. As can be seen, the differences are generally rather small.

The even a-coefficients also seem to change in a manner similar to that reported by Libbrecht, as illustrated in figure 3.61. Unfortunately, Libbrecht has not published splittings for any of these two years (only frequencies for 1989) making it difficult to compare details.

As mentioned in the description of the analysis of the aj87 dataset, the linewidths from that dataset seem to have some problems, which are likely to also affect the mode amplitudes. Also, the difficulties in properly estimating the PSF present a problem. Fortunately, these problems

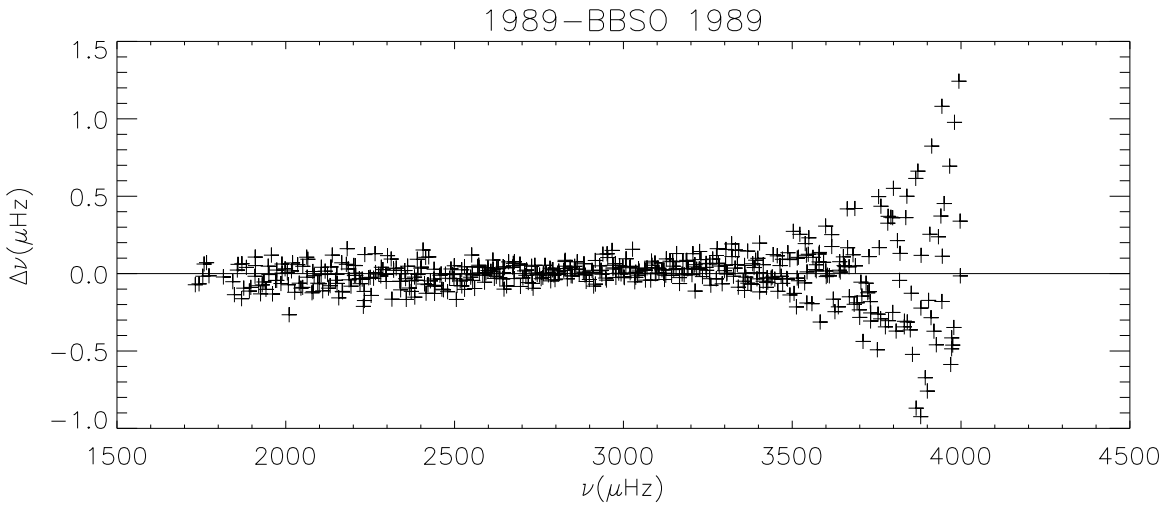

Figure 3.60: The frequency difference between the frequencies for April-June 1989 and frequencies from Libbrecht, March-August 1989. Again only $20 \leq l \leq 60$ was used.

appear to affect only high l modes. Restricting the l range such as not to include the affected modes, it appears that the amplitudes and linewidths are fairly similar, as shown in figure 3.62.

The change in the mode amplitudes between the two years is likely to be inaccurate when it comes to the l dependence (at least at high ls), due to the problems in determining the PSF. The frequency dependence, on the other hand, should be more reliable. It appears that the change has a frequency dependence such that the modes with frequencies around the peak of the p-mode band have lower amplitudes in 1989 compared to 1987, while those at low frequencies appear to have more power in 1989. Whether this is real or due to problems with the PSF or the different noise levels is difficult to tell. The l dependence for low ls, where the PSF should not pose too many problems, is also interesting. It appears that the mode amplitudes have decreased by approximately 5% at l less than ≈ 20. The fact that the amplitudes appears to have increased around $l = 60$ is probably an artifact of the PSF determination.

The linewidths seem to have increased systematically from 1987 to 1989 by about 20%. Again, there is a hint that the change is larger for the modes close the the peak in the mode amplitudes. There is a dramatic change in the linewidths above $l \approx 60$, but as mentioned, this is probably due to the unresolved problem concerning the linewidths, mentioned in the description of the 1987 dataset. While the problems with estimating the PSF are known to give problems with the estimation of the linewidths, as shown in the section on the tests performed, a change of 20% seems difficult to explain this way and could thus be real and connected to solar cycle variability.

The noise properties also seem somewhat different, as illustrated in figure 3.62. Given the different time cadences used the two years, different length of the time series, different apodization, and different estimation of the noise covariance matrix, the absolute size of the change is somewhat uncertain, and may not be real. Both the l and the frequency dependence show considerable structure. The frequency dependence seems to indicate that there is a larger

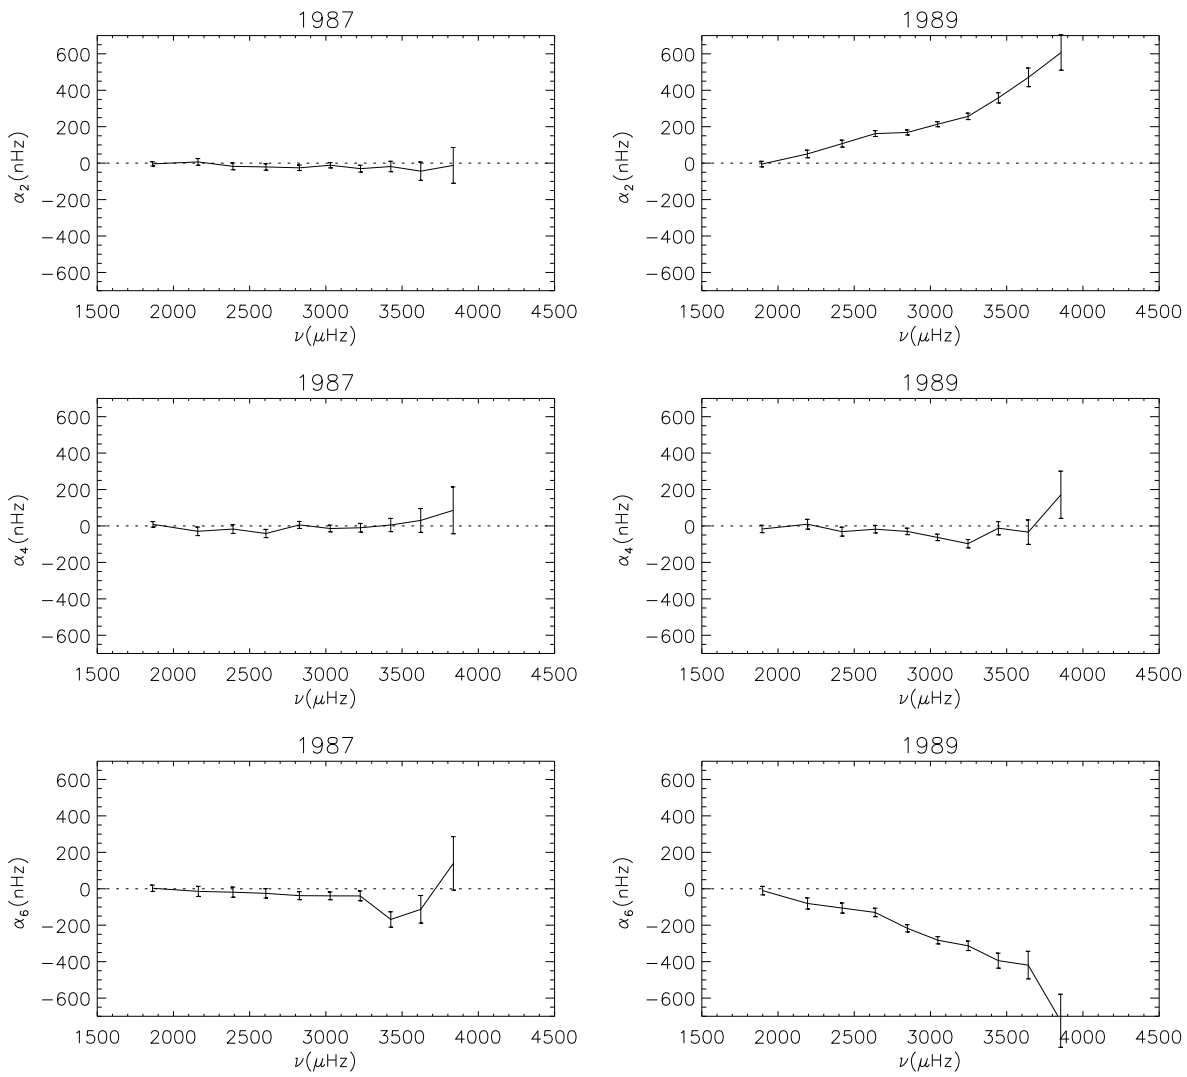

Figure 3.61: The even α-coefficients ($\alpha_i = l * a_i$) for 1987 and 1989. Only modes with $20 \leq l \leq 60$ were used. The αs were binned in frequency for clarity. The error bars are 1σ error bars on the average using the internally estimated errors.

problem with unresolved modes in the 1987 dataset compared to the 1989 dataset (the increase of the noise in the p-mode band is higher). A possible cause is the fact that the duty cycle is considerably lower in 1987 ($\approx 25\%$) compared to 1989 ($\approx 35\%$), leading to increased power far away from the main lobe in the temporal sidelobe spectrum.

The l dependence of the noise could indicate that some of the noise has different origins in the two years. One difference is that the 1989 noise problem on the left side of the solar images seems not to be present in 1987. The fact that there was this problem in 1989 could also be the cause for the frequency dependent change, in that it could swamp the contribution from unresolved modes, if the general noise level is considerably higher. The problems with estimating the PSF could be part of the cause too, as could the general change in the PSF between the two years.

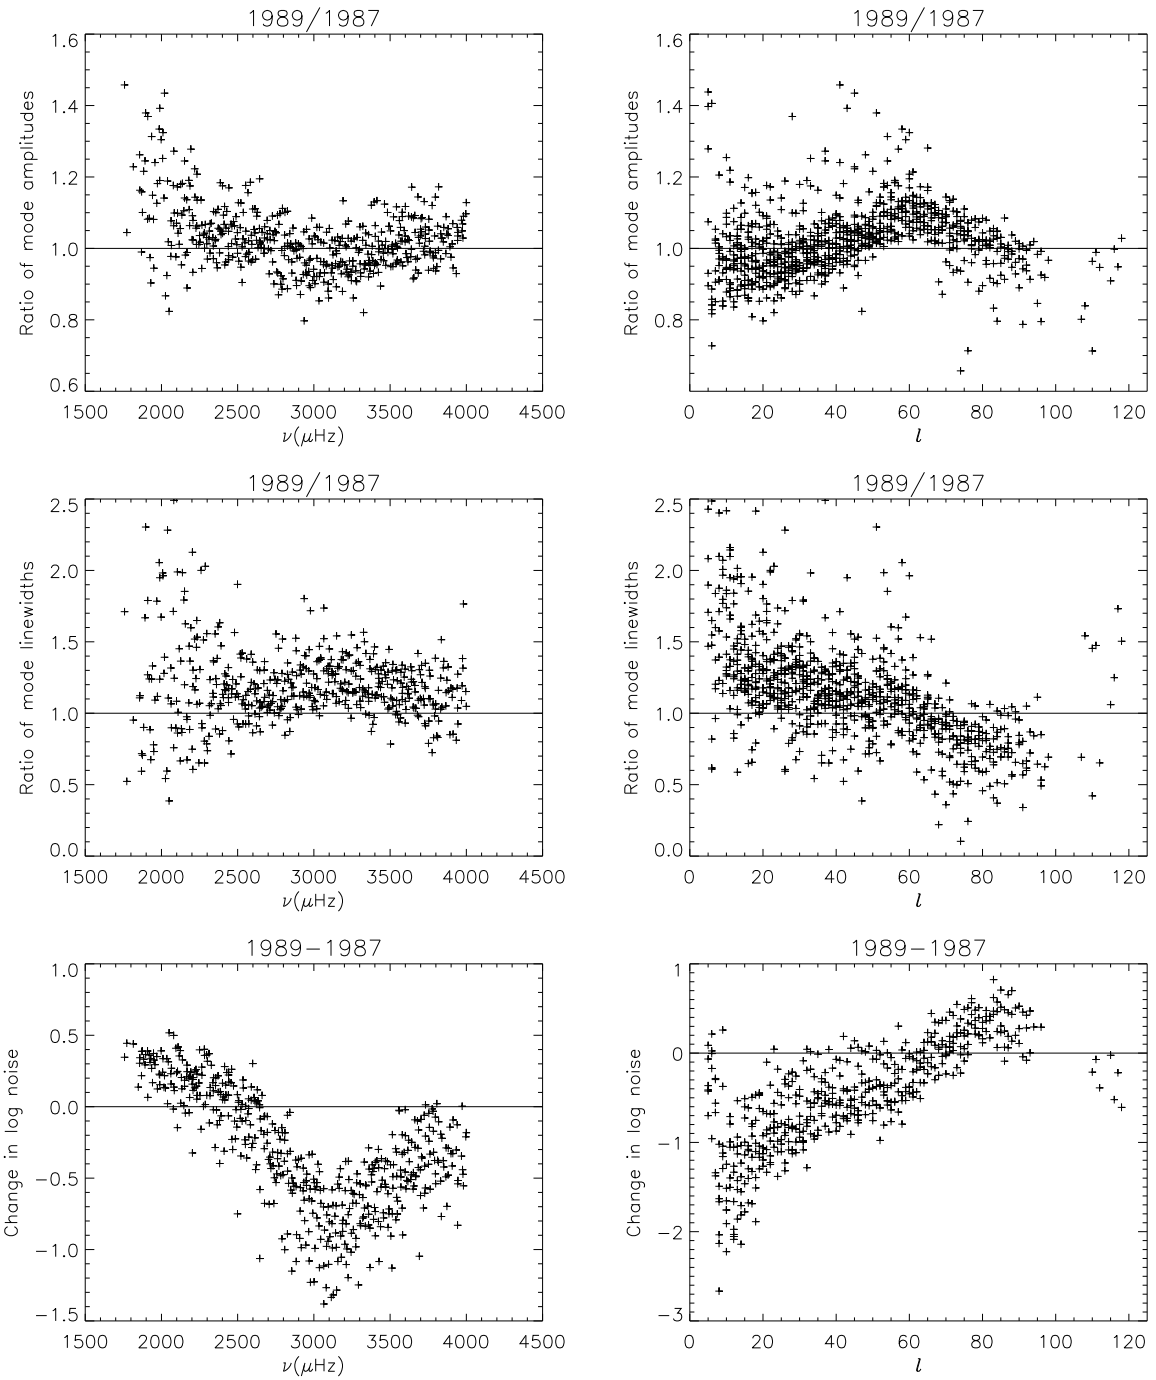

Figure 3.62: The ratio of the mode amplitudes and linewidths and the difference in the background noise level for the aj89 and aj87 runs. Only the range $20 \leq l \leq 60$ was used for the plots versus frequency to reduce the problems with the incorrectly estimated PSF and the problems with the linewidths in the 1987 run. For the plots of the amplitudes and linewidths versus l, all frequencies were used, for the plot of the change in noise versus l, only modes with $2500\mu Hz \leq \nu \leq 3500\mu Hz$ were used. Notice that the absolute value of the noise difference is partly caused by different lengths and time-cadences of the time-series for the two years.