Chapter 4

Data Analysis Methods

As outlined in section 2.3, the essence of time-distance helioseismology is the measurement of acoustic wave travel times. This chapter describes the various processing steps involved in measuring travel times from the MDI Dopplergrams described in chapter 3.

4.1 Remapping of Images

For the purposes of time-distance analysis (and indeed almost all helioseismic measurements) it is convenient to remap the images onto a new grid which is regular in some desired coordinate system. For reasons which will hopefully become clear, the remapping used in this work was the "cylindrical equal area" projection[1]. In this projection, the remapped image has points which are equally spaced in longitude and sine of latitude; all output pixels in a particular column have the same longitude, and all pixels in a row have the same latitude. The scale of the map is chosen such that at disk center it matches the scale of the input images. It follows that the image is oversampled at all other locations.

The remapping algorithm for each image takes into account the latitude and longitude of solar disk center, the effective solar position angle, and all other relevant

[1]The author is grateful to R.S. Bogart and the MDI software development team for the remapping software.

orbit parameters. The interpolation between input pixels uses a cubic convolution algorithm (Keys, 1981). In addition, the remappings used in this thesis were "tracked" in the sense that the remapped area was moved to follow regions as they rotated across the disk. The Sun, of course, does not rotate with a single angular velocity, so this tracking cannot exactly match the rotation at all latitudes.

The procedure for remapping, then, is to take a time series of images (typically several hours to several days) and produce a "data cube" where longitude is the first axis, sine of latitude the second axis, and time the third axis. This cube then forms the input to subsequent processing stages.

4.2 Filtering and Image Correction

Since the signal of interest in helioseismology is the acoustic wave (p-mode) signal, it is advantageous to remove other velocities which do not contribute to the desired output and therefore fall into the general category of "noise." Also, there are some instrumental effects which can be removed or corrected in order to enhance the final results.

4.2.1 Rotation gradient

It can be seen in figure 3.1 that a strong gradient is present in every Dopplergram. Most of this gradient does not vary with time, so it can be removed by subtracting a fixed pattern from each image. In almost all cases this was done during the remapping procedure. Typically, a background image was produced by computing the average of the input time series of Dopplergrams. This average image was then subtracted from each Dopplergram in the series before remapping. The signal in each pixel is thus forced to have a mean value near zero.

4.2.2 Supergranulation

Supergranules cause a problem for time-distance measurements because they are coherent velocity structures which cause nearby points on the solar surface to have

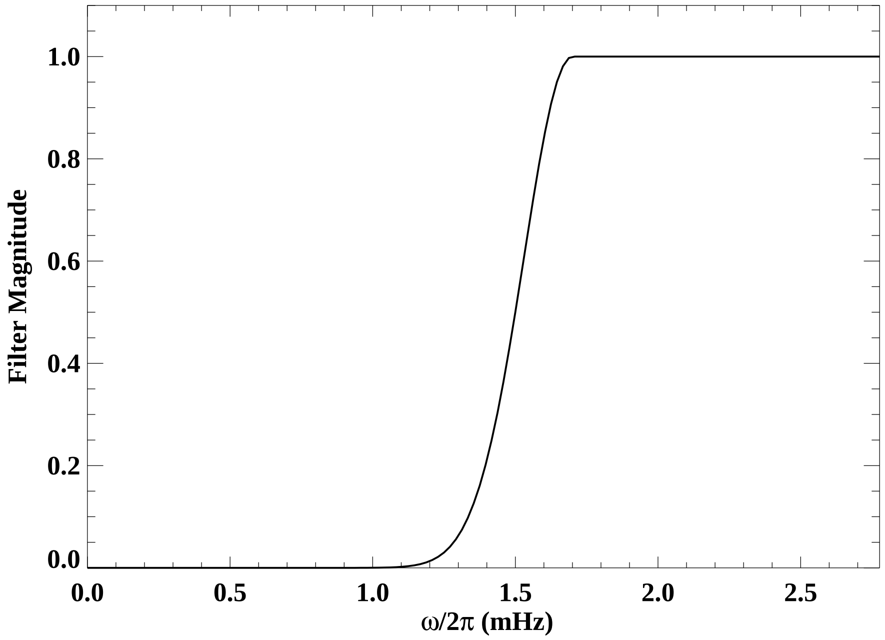

Figure 4.1: To remove supergranulation noise, a high-pass filter is applied to the data. The period at the cutoff shown here is roughly 10 minutes. Note that the horizontal axis does not extend all the way to the temporal Nyquist frequency ($\omega_N/2\pi = 8.33\,\mathrm{mHz}$). The filter shown is applied by multiplying with the power spectrum of the input data.

highly correlated velocities. Since time-distance measurements rely on the correlation of signals due to acoustic wave propagation from one point to another, it is necessary to remove as much of the supergranular signal as possible. Fortunately, supergranulation can be largely removed from the input data by use of a simple high-pass filter (see figure 4.1). Typically the filter used for this purpose was a gaussian roll-off with full transmission above 1.7 mHz (periods corresponding to about 10 minutes) and a full width at half maximum of 0.40 mHz.

4.2.3 Surface gravity waves

Although the supergranulation is easily separated from the acoustic signal in the temporal frequency domain, some other sources of noise are not. One such source

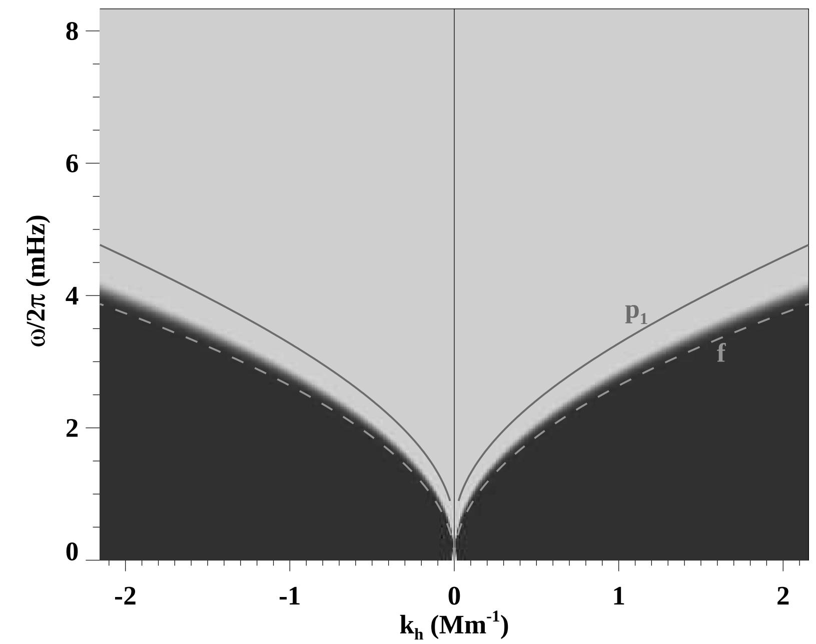

Figure 4.2: A depiction of the filter used to remove the oscillation due to the fundamental or f mode. The darkest greyscale indicates a filter magnitude of 0.0; the brightest, a magnitude of 1.0. The three-dimensional filter is multiplied by the power spectrum of the input data. The solid line shows the location of the p_1 ridge (see equation 4.2) in the k-ω diagram, and the dashed line the position of the f-mode (equation 4.1). The filter is constructed so that the f-mode is attenuated by a factor of 100.

of coherent velocity signals is the so-called "fundamental" mode (f-mode) of solar oscillation. The waves which make up this mode are surface gravity waves (analogous to waves on the surface of the ocean). Since these are propagating disturbances which do not follow acoustic ray paths, they pose a problem in the interpretation of measured travel times if not isolated from the acoustic signal.

The solution to this problem can be found in the k-ω diagram of figure 2.1. At each temporal frequency, the power of the f-mode appears at a larger value of horizontal spatial frequency k_h than for the p-modes. Thus it is possible to design a filter that removes all power beyond a certain value of k_h, where that value depends on ω. The locations of the f- and p_1-mode ridges in k-ω space were approximated with polynomial forms:

$$l_0 \approx \mathrm{R}_\odot k_{0h} = 100\nu^2 \tag{4.1}$$

$$l_1 \approx \mathrm{R}_\odot k_{1h} = \sum_{k=0}^{4} c_k \nu^k, \qquad c = \{17.4, -841, 95.6, -0.711, -0.41\} \tag{4.2}$$

where the cyclic frequency $\nu \equiv \omega/2\pi$ is measured in milliHertz. The filter was then constructed at each frequency as a low-pass filter in the spatial domain, with gaussian roll-off, full transmission halfway between k_{0h} and k_{1h}, and 1% transmission at $k_h = k_{0h}$ (see figure 4.2).

4.2.4 Phase velocity filtering

Since the travel time between two points is computed as the cross correlation of their respective velocity signals, it is beneficial to consider only those waves which travel between the two points. As described in section 2.3.2, the travel distance Δ between photospheric reflections is determined by the ratio of the temporal and spatial frequencies, ω/k_h. All waves with the same value of ω/k_h travel approximately the same distance, and all other waves travel different distances. If we are only interested in a finite range of distances, then, we can increase the signal-to-noise ratio of the measurement by filtering out waves which do not have the appropriate horizontal phase speed. In the k-ω diagram, the waves of interest have power which lies within a wedge-shaped region (see figure 4.3); in the three-dimensional power spectrum the surfaces are cones of constant ω/k_h. Again the filter is constructed to have a gaussian rolloff from each surface.

It should be noted that this filter should be used with caution when distances become large. It has been my experience that using a phase speed filter for large distances (large values of ω/k_h) can cause systematic errors and reduce the signal-to-noise ratio of the cross correlations. This may be due to projection effects on the power spectrum; since the remapping used is equally spaced in sine of latitude, the spatial wavelength scale in that direction is difficult to disentangle in an FFT.

In addition, for the work described here the measurement is always of a flow

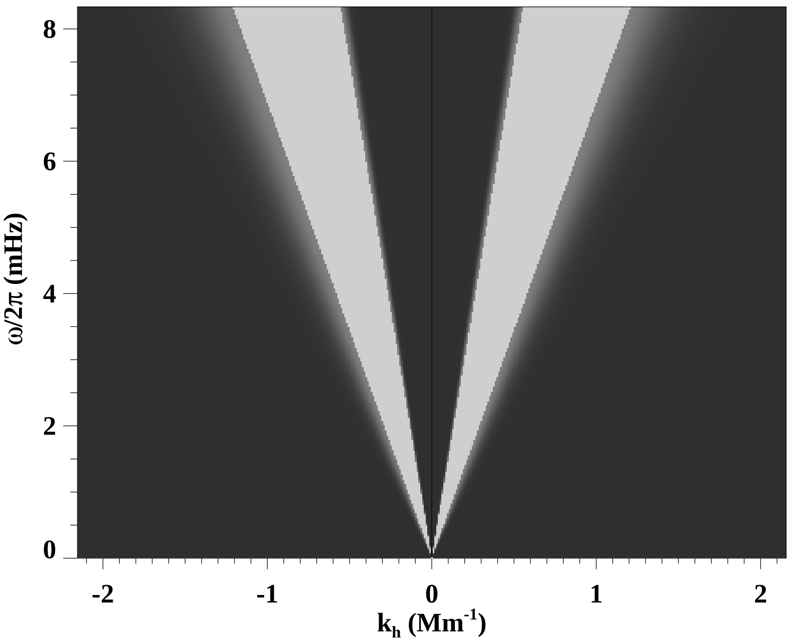

Figure 4.3: An example of a phase speed filter used for the anlysis. The horizontal phase speed is given by the ratio of the temporal frequency ω to the spatial frequency k_h. This figure shows a section of the power spectrum of the filter which is applied to the data; the Fourier transform of the filter is multiplied by the Fourier transform of the data. Here the lightest greyscale represents a filter value of 1, and the darkest a value of 0. The upper cutoff is the phase speed corresponding to a travel distance $\Delta = 10°$, and the lower corresponds to $\Delta = 3°$.

in a particular direction. In this case it is also beneficial to remove those waves whose horizontal phase velocities are in different directions. If I denote the horizontal wavenumber in the longitude coordinate as k_x, and in the latitude coordinate as k_y, then the waves of interest lie in a wedge-shaped region in the k_x-k_y plane centered on either the k_y axis (to measure the meridional flow: see figure 4.4) or the k_x direction (to measure the rotation). In three-dimensions the surfaces are planes of constant $\phi_k = \arctan(k_y/k_x)$.

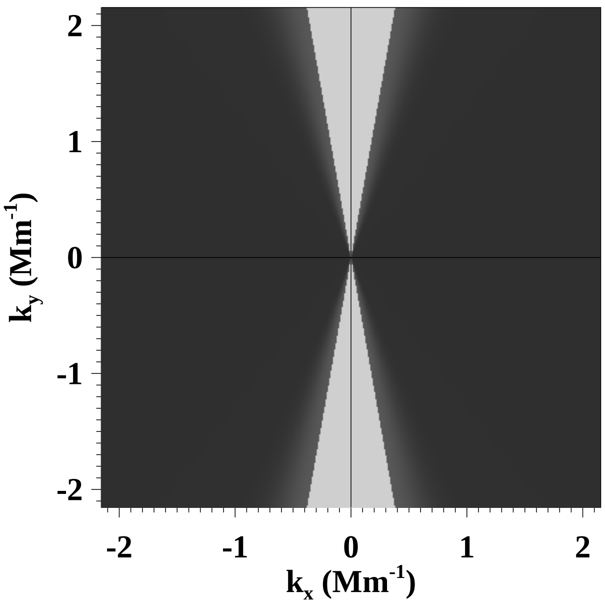

Figure 4.4: An example of a directional filter used for the anlysis (see text). The axis coordinates are the horizontal wavenumber in the longitude and latitude. The lightest greyscale represents a filter value of 1, and the darkest a value of 0.

4.2.5 Instrument modulation transfer function

The modulation transfer function (MTF) of the MDI instrument has the property that signals of high spatial frequency are attenuated relative to signals of lower spatial frequency (Scherrer et al., 1995). A typical measured power spectrum is shown in figure 4.5. For the purposes of time-distance helioseismology, the low power at high wavenumber means that waves with short travel distances are more difficult to observe and hence that the surface layers are less easily resolved. To correct for this effect, the input power spectrum is divided by the measured MTF. In practice, I have

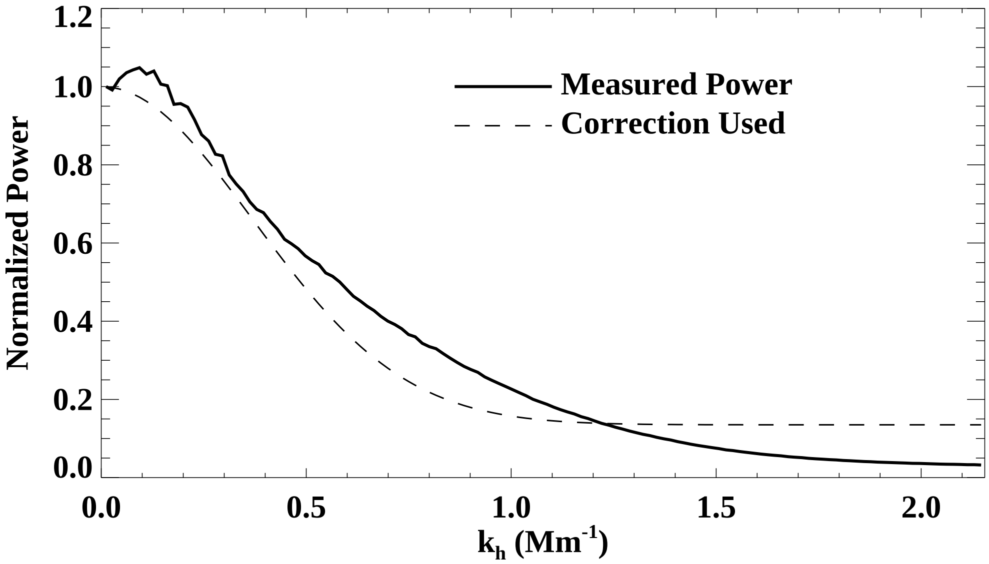

Figure 4.5: The MDI instrument attenuates the power at high spatial frequencies compared to the power at low spatial frequencies. The solid curve shows the power, averaged over temporal frequency and direction, from eight hours of Dynamics images. The power at low wavenumber has been normalized to 1. (Compare to Scherrer *et al.* (1995), figure 5.) The dashed curve is described by equation 4.3.

approximated the MTF by a smooth function:

$$M(k_h) = 0.865 \exp\left[-(\mathrm{R}_\odot k_h/364)^2\right] + 0.135, \tag{4.3}$$

This function describes the attenuation of the oscillation *power* as a function of wavenumber; to remove this effect, the Fourier transform of the input data is divided by the square root of M. The function M is constructed with a flat response at the highest values of k_h, since in this region the noise (non-oscillatory signal) in the solar velocity spectrum becomes large and it is not desirable to overamplify the power in this region.

In addition to the expected attenuation of high-k oscillation power, it has been observed that the power measured in MDI Dopplergrams varies with the direction of

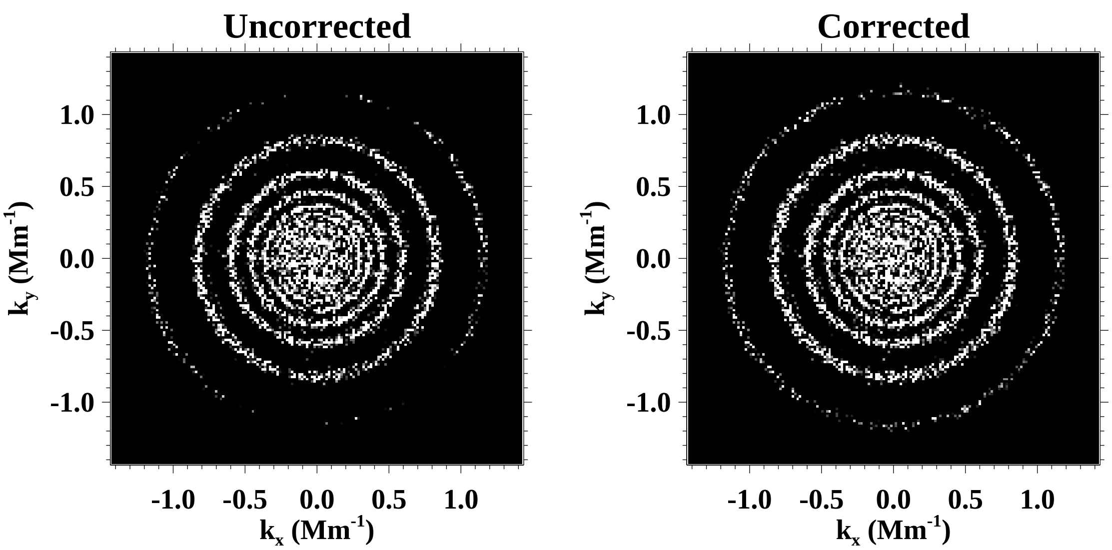

Figure 4.6: "Ring diagrams" for eight hours of Dynamics images from MDI. Each figure is a slice through the three-dimensional power spectrum of the series of images, at temporal cyclic frequency $\omega/2\pi = 4.16\,\mathrm{mHz}$. The horizontal axis is the horizontal wavenumber in the longitude direction and the vertical axis is the horizontal wavenumber in the latitude direction. The greyscale in each case is saturated to emphasize the variation of power with angle in the outermost ring, where the effect is most pronounced. The left figure is the uncorrected image; on the right, the power spectrum has been corrected with the function described in equation 4.4.

propagation; there is less apparent power in the north-south direction than in the east-west direction. This astigmatism can, like the MTF, be measured and then corrected for in the input power spectrum (see figure 4.6). The measurement is done at disk center and then applied to data from any position on the disk. Note that near the limbs there is some variation of acoustic power with direction due to foreshortening and the fact that the acoustic signal is mainly radial; this effect is not corrected for. The measured astigmatism is given by

$$
\begin{aligned}
A(k_h, \phi_k) &= [1 - Q(k_h)]\, P(\phi_k) + Q(k_h) \\
Q(k_h) &= \begin{cases} 0.56, & \mathrm{R}_\odot k_h > 1100; \\ \exp\left[-(\mathrm{R}_\odot k_h/1457)^2\right], & \mathrm{R}_\odot k_h \le 1100 \end{cases}, \\
P(\phi_k) &= \begin{cases} \cos^4(\phi_k - 0.153), & k_h \ne 0; \\ 0, & k_h = 0 \end{cases}.
\end{aligned}
\tag{4.4}
$$

where $\phi_k \equiv \arctan(k_y/k_x)$ is the direction east of north on the disk. Equation 4.4 describes the power as a function of direction and degree; to remove the astigmatism, the Fourier transform of the input data is divided by the square root of A. Note that for small values of the horizontal wavenumber k_h the astigmatism is quite small; in particular, the correction is not very important for images from the Structure program, which contain information only for small and moderate wavenumbers.

4.3 Computing Cross Correlations

Once the Fourier transform of the three-dimensional data cube has been filtered and corrected as described in section 4.2, the inverse Fourier transform is computed in the two spatial dimensions. This leaves an array of temporal Fourier tranforms, one for each input pixel. These are used to compute temporal cross correlations between pairs of points, taking advantage of the special relationship between cross correlation and multiplication in the Fourier domain (see equation 2.7). This relationship greatly simplifies the computation of cross correlations. It remains to discuss how pairs of pixels are chosen.

Since for this work it is desired to measure the flow in a particular direction (either east-west or north-south), it is necessary to measure travel times for waves which propagate parallel to that direction. For example, to measure a meridional flow, pairs of points are chosen which have the same longitude. For any particular latitude, there are a number of pairs of pixels which have the same longitude and which are the same distance apart. These cross correlations are computed and then averaged together to improve the signal-to-noise ratio in the final result. Furthermore, pairs with similar — but not exactly the same — latitudes and separations are also added to the average, since they are expected to have similar travel times. Finally, cross correlations are computed for pairs of points which are only approximately on the same line of longitude.

The analysis for measuring solar rotation proceeds similarly, except that target pairs of points have the same latitude instead of the same longitude.

4.3.1 Observing modes and computational burden

As described in section 3.2.2, two kinds of Dopplergrams have been used for the measurements in this thesis. Let me make a few comments on the two cases.

First, the increased spatial resolution of the Dynamics images is obviously an advantage when it comes to measuring cross correlations for short distances. This is not only because the centers of the pixels are closer together, therefore permitting measurements for short distances, but because the larger number of pixels leads to an increase in the number of pairs and an increase in the number of cross correlations which can be computed. This is equivalent to an increase in the number of measurements, and leads to an improvement in the signal-to-noise ratio.

For longer distances, however, it is not clear that a great gain is made by using higher-resolution images. Waves which travel longer distances have a larger spatial wavelength and therefore a larger coherent wavefront. Taking more samples of this large wavefront should not improve the noise characteristics of the measurement. Therefore I expect that for large distances the resolution of the Structure images should be adequate. I will show some evidence for this claim in section 4.5.

This is also fortunate from a computational point of view. Since the Structure images have many fewer pixels (by a factor of 25), the number of cross correlations to be computed is greatly decreased. This means that one month's worth of Structure images can be analyzed in a few days, up to a distance of 45°, whereas the same amount of time is required to compute cross correlations from about 6 days of Dynamics images, for distances up to only 6°.

On the other hand, the Dynamics images have a great advantage over the Structure images at high latitudes. This is partly due to the increased spatial resolution, which

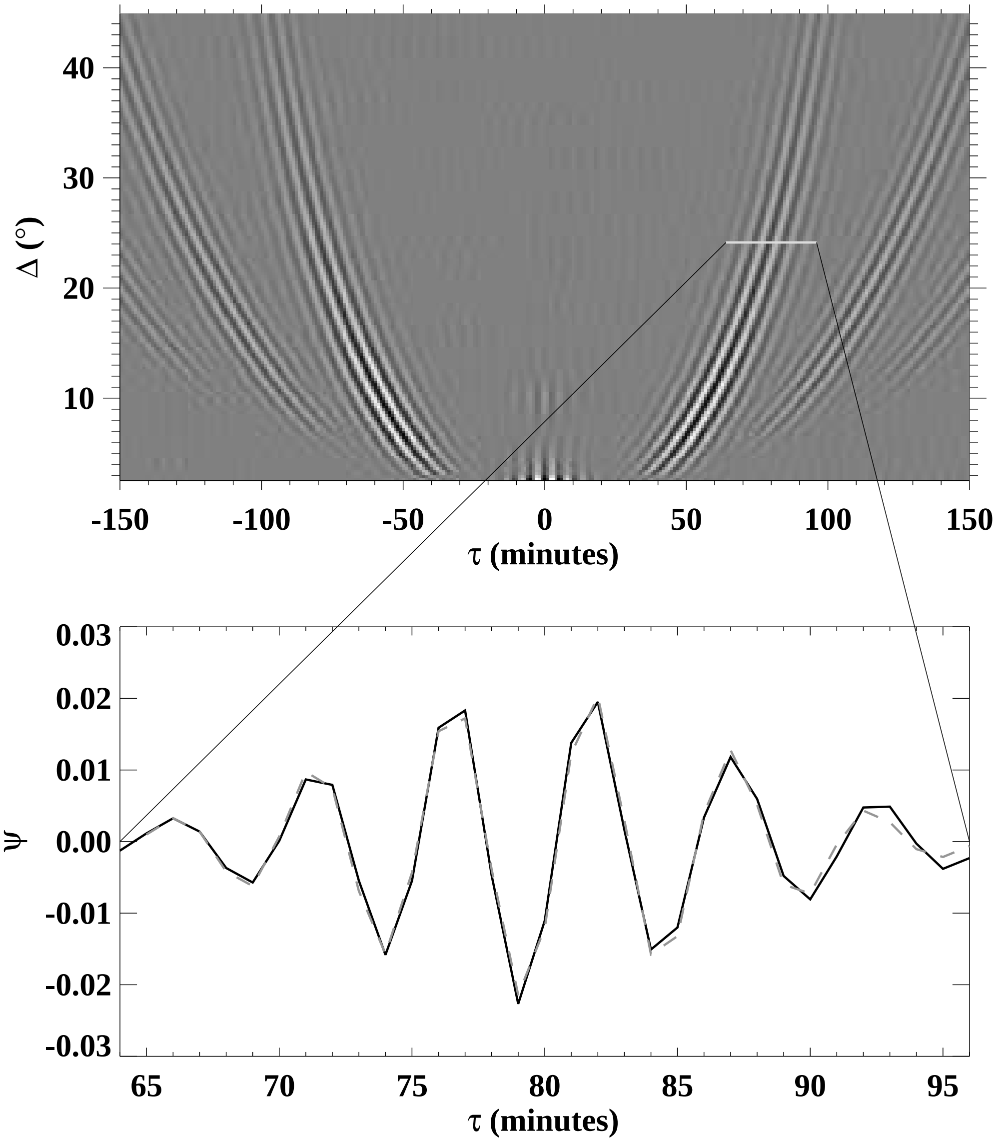

Figure 4.7: This "time-distance diagram" shows some cross correlations typical of those used in this work. In the upper plot, the greyscale denotes the cross correlation amplitude as a function of the time lag τ and the distance Δ. The lower plot displays the cross correlation for a particular distance ($\Delta = 24.1°$), near the maximum. The dashed line in the lower plot shows the function $G(\tau)$ (equation 4.5), where the free paramaters have been determined by a non-linear least squares technique.

allows more pixels in regions near the limb. In addition, the Structure images are cropped by the on-board processor to allow them to fit into the telemetry stream.

4.3.2 Averaging of cross correlations

In order to improve the signal-to-noise characteristics of the measurements, it is often necessary to average the cross correlations together over certain ranges in distance, latitude, and time. The averaging in distance is fairly minimal, generally over a range of 0.25 degrees. It is not possible to average the cross correlations very much in this dimension because the position of the peak depends strongly on distance (figure 4.7) and excessive averaging tends to smear the result.

Averaging in latitude is more flexible, since the dependence of the cross correlation on latitude is fairly weak. Generally, the results presented here have been averaged over latitude bands of five degrees. Some care is taken to ensure that the contribution of each cross correlation to the average is properly weighted. The weight is proportional to the total number of pairs of pixels used in each component cross correlation. This is especially important at high latitudes, where the number of available pixels drops off sharply as the region of interest approaches the limb.

Finally, since the cross correlations are generally computed on relatively short time series of images (a few hours to a few days), it is sometimes desirable to average together many such series. For example, to examine the axisymmetric component of the meridional circulation, it is obviously important to observe the Sun for at least one rotation period. The results presented here have been averaged over different time spans, depending on the object of the measurement; the results for the flows very deep in the convection zone come from averages over more than 700 days. Many measurements are averaged over one or a few solar rotations. In any case, some care is again taken to compute a properly weighted average. In the case where the observations do not each span the same length of time, the contribution of each cross correlation to the sum is proportional to the length of the observation.

4.4 Travel Times From Cross Correlations

Section 2.3.1 illustrates that the peak in the cross correlation function can be interpreted as the wave travel time between points on the solar surface. The second peak at negative time lag τ is interpreted as the travel time in the opposite direction. Each travel time is located by fitting a function of the form (compare to equation 2.21)

$$G(\tau) = A \exp\left(-\frac{\delta\omega^2}{4}(\tau - \tau_g)^2\right) \cos(\omega_0(\tau - \tau_p)) \tag{4.5}$$

to the cross correlation near its peak. Here, A is the cross correlation amplitude, ω_0 is the central frequency of the traveling wave packet, and $\delta\omega$ is the frequency bandwidth. The two parameters τ_g and τ_p are the group and phase travel times, respectively. Figure 4.7 shows the function $G(\tau)$ and a typical measured cross correlation. Because the phase time is more accurately determined in the fitting procedure, I have used the phase time in all cases and will drop the subscript p.

The end result of the fitting is a collection of travel times τ^+ in the "positive" direction (southward or eastward) and τ^- in the "negative" direction (northward or westward). From these, the time differences are computed:

$$\delta\tau(\lambda, \Delta) = \tau^+ - \tau^-. \tag{4.6}$$

These can then be interpreted as a measure of the subsurface velocity.

4.5 Measurement Uncertainties

The fitting procedure described in section 4.4 does not, unfortunately, give a realistic estimate of the uncertainty in the fitted parameters and hence in the measured travel time τ.

In order to estimate the uncertainty in the travel times, I examine the scatter of the measurements about their mean. Fortunately, there are a large number of measurements and it is possible to get meaningful statistics from their distribution. However, some care must be taken in the process and it is worth describing in some detail.

First, it is obviously important to weed out any outliers from the measurements. These are principally points where the fitting did not converge because the function 4.5 is a poor fit to the cross correlation. The most reliable way to remove these points is by visual inspection of the fits and the cross correlations. Fortunately, this is generally only a problem for the shortest distances. The problem is accentuated at high latitudes, but there are still relatively few fits (a few dozen out of a few thousand, say) which need to be inspected by eye.

Next, I compute the travel time difference for each pair of travel time measurements. These can be used to compute the mean travel time difference, and the scatter about that mean, as a function of latitude and distance. This is accomplished by selecting small subsets of points, with similar distances and latitudes, and computing the statistical moments for that subset of measurements. Then each of the measurements in the subset is assigned an uncertainty which is equal to the standard deviation of the subset. Finally, the uncertainties are smoothed in both latitude and distance to remove any discontinuities in the distribution of errors.

Figures 4.8 and 4.9 show the distribution of measurement errors as a function of latitude and distance. These measurements were made from the first two years of Structure images and make use of almost the entire disk for averaging of cross correlations. The two measurement directions give somewhat different distributions of errors because of the geometry of the remapped regions.

Earlier I stated that the Structure program Dopplergrams, despite their lower spatial resolution, could be used to measure travel times for large distances with no degradation of the signal-to-noise ratio. Figure 4.10 shows a comparison of the Dynamics and Structure measurements for time series of equal length which were nearly contemporaneous and covered comparable areas of the solar image. Clearly the Dynamics images have a significant advantage for measuring time differences at short distances; in fact, it is quite difficult to make measurements from Structure images for distances less than three degrees. However, the Structure images are quite adequate for distances larger than about ten degrees.

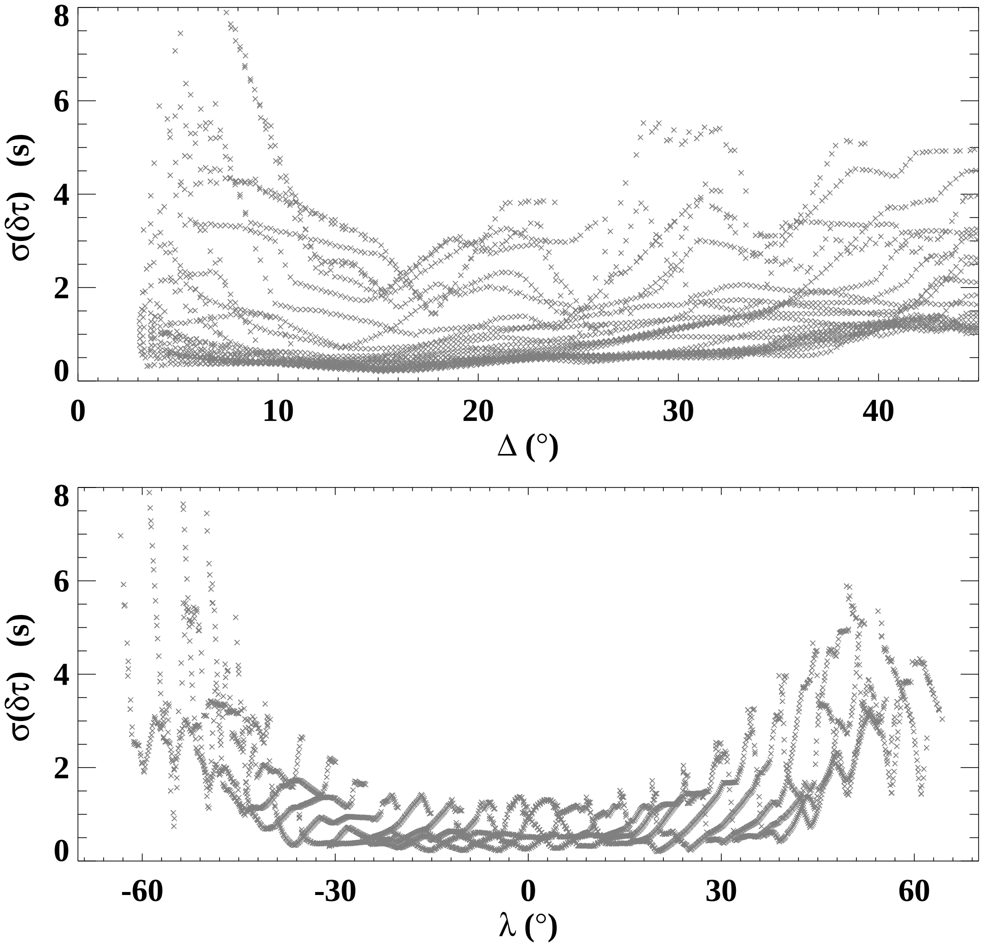

Figure 4.8: Measurement errors for the south-north time differences are shown as a function of distance and latitude. The measurements are from 792 days of Structure images. Each $\times$ represents a single measurement with unique coordinates (Δ, λ).

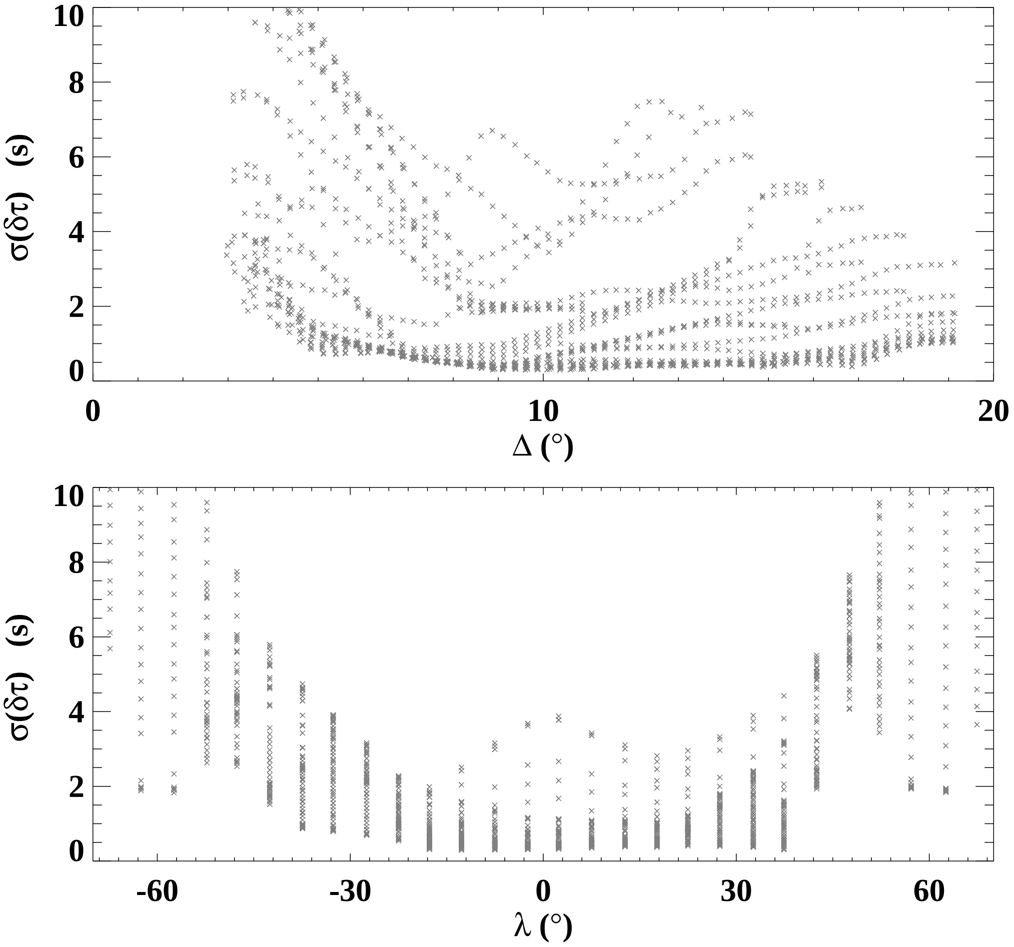

Figure 4.9: Measurement errors for the west-east time differences are plotted as a function of distance and latitude. The measurements are from 792 days of Structure images. Each $\times$ represents a single measurement with unique coordinates (Δ, λ).

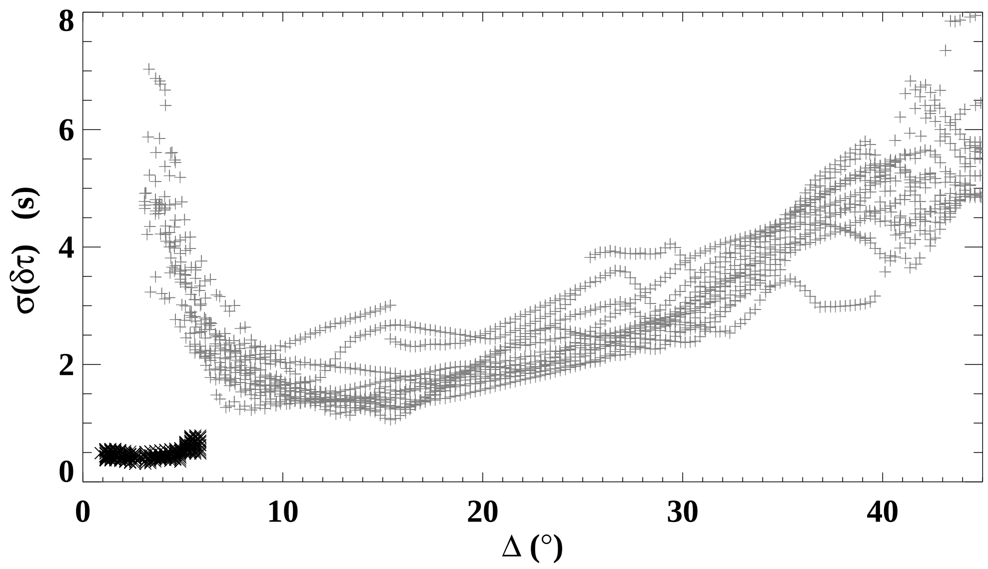

Figure 4.10: The measurement errors σ are shown as a function of distance. The errors σ are for a selection of south-north time differences, measured at locations within 30° of the equator. The points in the lower left corner, marked with a black $\times$, were made from one month of Dynamics images from 1998. The points marked with a gray + were made from one month of nearly contemporaneous Structure images. The spatial averaging was also comparable in the two datasets.

4.6 Typical Measurements

The set of time difference measurements $\delta\tau(\lambda, \Delta)$ contains information about the solar velocity field $\mathbf{u}(\lambda, r)$. One interesting way to display the measurements is to simply average them over different distance ranges and display the results as a function of latitude. Figure 4.11 shows a typical set of measurements of south-north travel time differences. This plot includes a particular range of distances and the results are averaged in latitude bins of five degrees. When the distances Δ are not too large compared to the latitude bins, a plot like figure 4.11 can be said to describe an average of the meridional flow over a certain depth range near the surface.

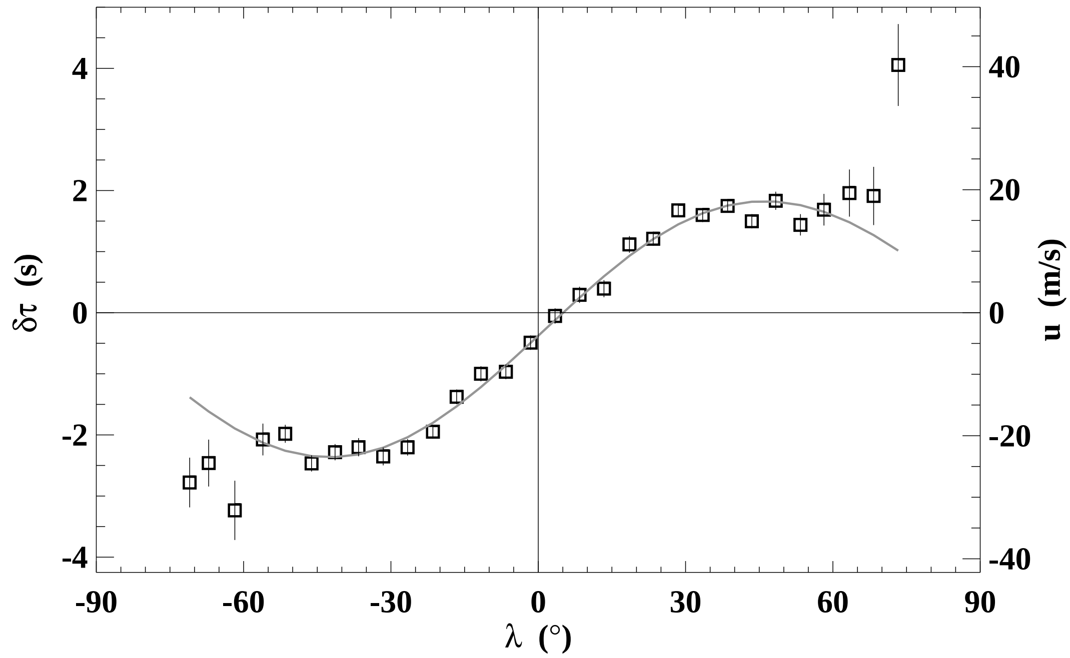

Figure 4.11: A typical set of measurements, from the Dynamics period of 1996. The time difference $\delta\tau$ is plotted as a function of latitude λ; the measurements have been averaged in distance for $(1° \le \Delta \le 6°)$. The sense of the time difference is south − north, so that a positive time difference indicates a northward flow. The solid curve is a least-squares fit of the form 4.7. The right-hand axis indicates the corresponding meridional velocity in m/s. This velocity was calculated by assuming that the velocity does not vary with depth over the range in question; a ray with $\Delta = 6°$ penetrates to $r = 0.963\,\mathrm{R}_\odot$.

Following some past observers, I sometimes will use a smooth function to approximate the latitude dependence of the meridional circulation:

$$\delta\tau = \sum_{k=1}^{N} a_k \sin(k[\pi/2 - \lambda]). \tag{4.7}$$

Generally the fit is made with $N = 4$ and with coefficient a_3 set to zero.

Figure 4.12 is a similar plot for the east-west time differences, and thus represents a measurement of the solar near-surface rotation. In the case of the rotation,

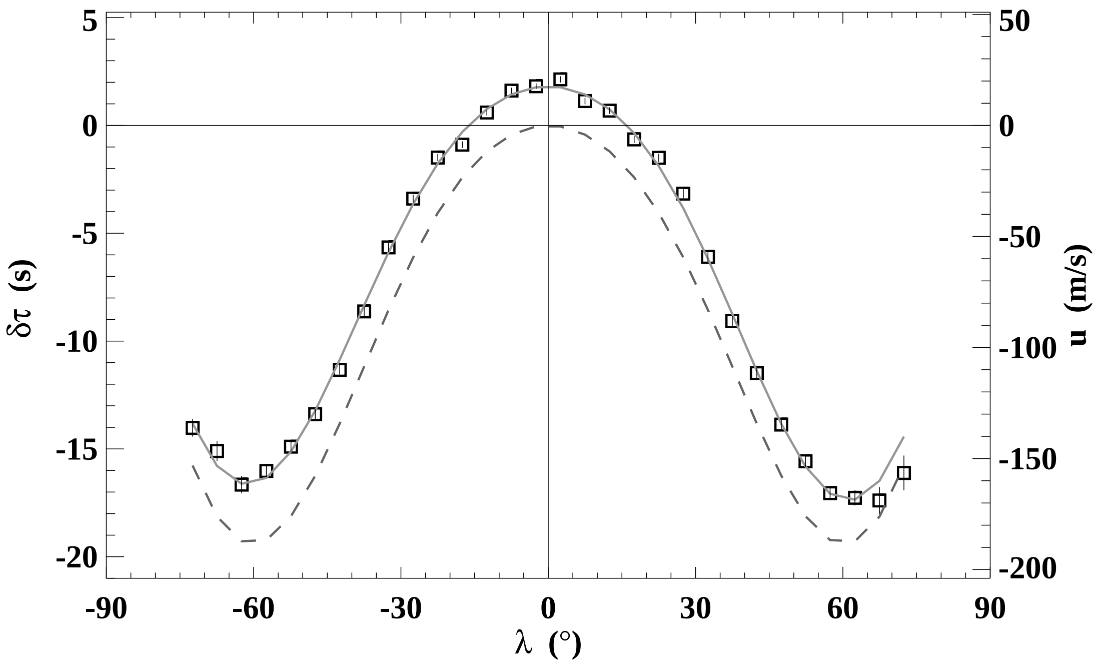

Figure 4.12: A typical set of measurements, from the Dynamics period of 1996 (see also figure 4.11). For these measurements, the sense of the time difference is east − west, so that a positive time difference indicates a prograde flow. The solid curve is a least-squares fit of the form 4.8. The right-hand axis indicates the corresponding zonal velocity in m/s, calculated as described in figure 4.11. The dashed curve shows the rotation velocity at the solar surface as measured by Snodgrass (1984), plotted with respect to the right-hand axis. The tracking velocity at each latitude has been subtracted, to compare with the measurements.

as described in section 1.2.1, both the surface and interior values are already quite well known, so we are able to check our results against those from other methods. Figure 4.12 shows a measured surface rotation rate for comparison with the time-distance measurement. Clearly the time-distance measurements give a larger velocity than the surface measurements; this is in agreement with other helioseismic measurements which indicate an increasing angular velocity in the layers just below the solar surface. I will have more to say about this when I discuss inversion results in section 6.1.2.

The solid curve in figure 4.12 is a least squares fit of the following expansion:

$$\delta\tau = \sum_{k=1}^{N} a_k \sin(\pi/2 - \lambda) \cos^k(\pi/2 - \lambda) \tag{4.8}$$

Comparing this to equation 1.1, we see that the two forms are equivalent if $N = 4$ and $a_1 = a_3 = 0$.

4.7 Instrumental and Systematic Errors

4.7.1 Roll angle uncertainty

One of the most obvious features of the meridional flow measurement in figure 4.11 is the southward flow across the equator (about 5 m/s for that particular measurement) which suggests that the meridional circulation is not antisymmetric about the equator. It has been pointed out that such a velocity pattern could result from an error in the orientation of the MDI camera of a few arcminutes (Giles et al., 1997). This would mean that the nominal "north" direction would not be truly orthogonal to the solar rotation velocity; a small component of this velocity (which is roughly 2000 m/s at the equator) could leak into the meridional circulation measurement.

Although the MDI instrument had its camera aligned very precisely with the SOHO spacecraft reference direction prior to launch, it is possible that some shift occured during the launch. Furthermore, thermal stresses may cause the spacecraft itself to deform. One way to evaluate this misalignment would be to assume that the cross-equator meridional circulation was entirely due to such an effect. However, this assumption may not be entirely justified. Another way to evaluate the orienation of the MDI camera is to compare with observations from ground-based instruments where the absolute orientation can be carefully measured. Evans (1999) has performed a comparison between images from MDI and Dopplergrams from the Mount Wilson Solar Observatory and finds that the MDI images are rolled by $0.3°$ with respect to the Mount Wilson images. The sense of this roll is the same as that required to explain a southward flow across the equator; however, the magnitude is large enough to produce a flow of 10 m/s, which is larger than the average value observed. Evans estimates that the random error in the determination of the roll angle is about 0.05°.

Another possibility along this same line is that the coordinates used by SOHO to maintain the correct attitude do not accurately describe the direction of the solar rotation axis. Interestingly, this error would cause the cross-equator component of the meridional circulation to vary with time as SOHO moves in its orbit. This possibility offers an intriguing way to measure the direction of the Sun's rotation axis, as outlined in appendix A.

It is possible, of course, that the equator-crossing flow is not an error at all. There is no requirement from first principles that the meridional circulation has to be zero at the equator. Other measurements of the meridional circulation using different techniques have also found similar effects (Schou and Bogart, 1998; Basu et al., 1998; Meunier, 1999); but all of these results were obtained using MDI images and so are subject to the same intrumental errors. An independent observation of an equator-crossing flow has been noted by Komm *et al.* (1993b), using magnetic field observations from the Vacuum Telescope at the National Solar Observatory on Kitt Peak. These researchers noted that there was a northward flow across the equator which averaged 10 m/s during the period from 1978 to 1990, but attributed this to telescope misalignment. It is probably fair to say that the question of the equator-crossing meridional circulation is still unresolved. Further information can be gathered by looking at the depth dependence of the flow and comparing to the profile which would be expected if the flow was due to a roll angle. I will comment on this again in later sections.

4.7.2 Plate scale uncertainty

Another effect which is common to a large number of measurements from MDI has been labelled the plate scale error. The observational effect is that the solar rotation rate appears to be larger on the receding (westward) side of the disk than it does on the approaching (eastward) side. The solar plasma thus appears to be moving faster and faster as it traverses the disk. (see figure 4.13). This is seen in measurements

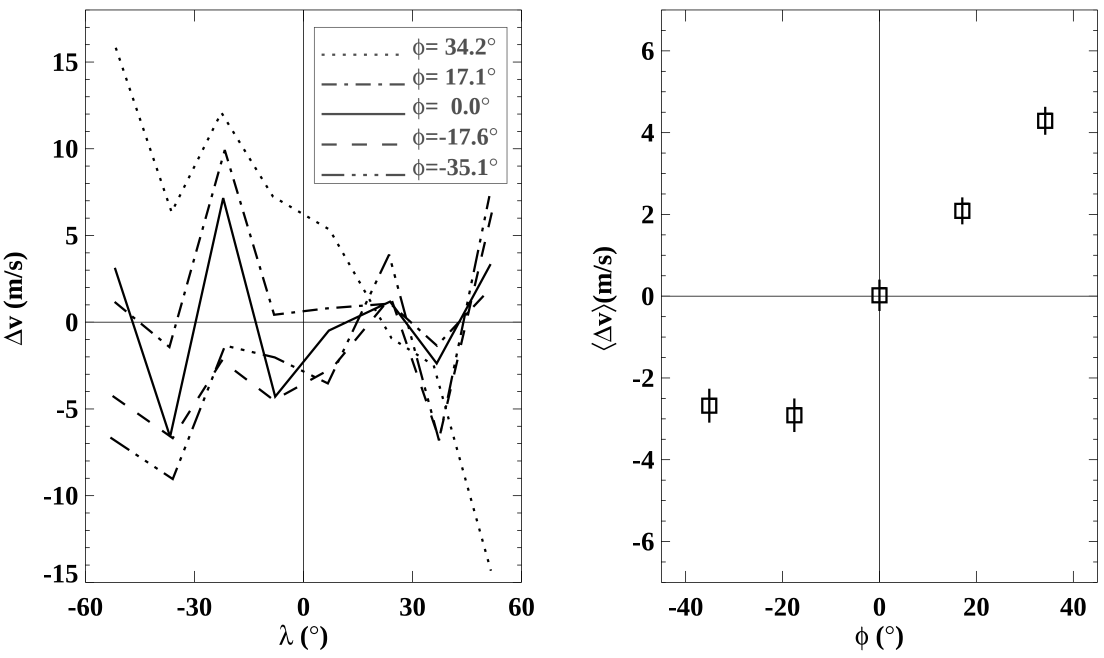

Figure 4.13: The left-hand figure shows the zonal velocity as a function of latitude for different positions on the disk. The measurements were made for one complete Carrington rotation during the Dynamics run of 1998. The different curves are for different longitudes ϕ, measured westward from the center of the disk. A smooth fit of the form of equation 4.8 has been removed from each of the five curves, and the residuals are plotted. The right-hand plot shows the latitude-averaged residuals as a function of ϕ, with error bars to give an idea of the significance of the effect.

from time-distance helioseismology, ring diagram analysis, and correlation tracking, although the magnitude of the effect seems to vary from method to method. The fact that this is seen even when the measurements are averaged over several months indicates that it cannot be a real solar phenomenon.

The cause of this effect is probably the fact that image focus is not uniform over the MDI camera. This means that the plate scale, defined as the heliocentric distance between pixels, is also not uniform. In remapping the MDI images as described in section 4.1, the plate scale is assumed to have some uniform value; if the plate scale is in fact smaller on the west side of the image than it is on the east side, then a feature

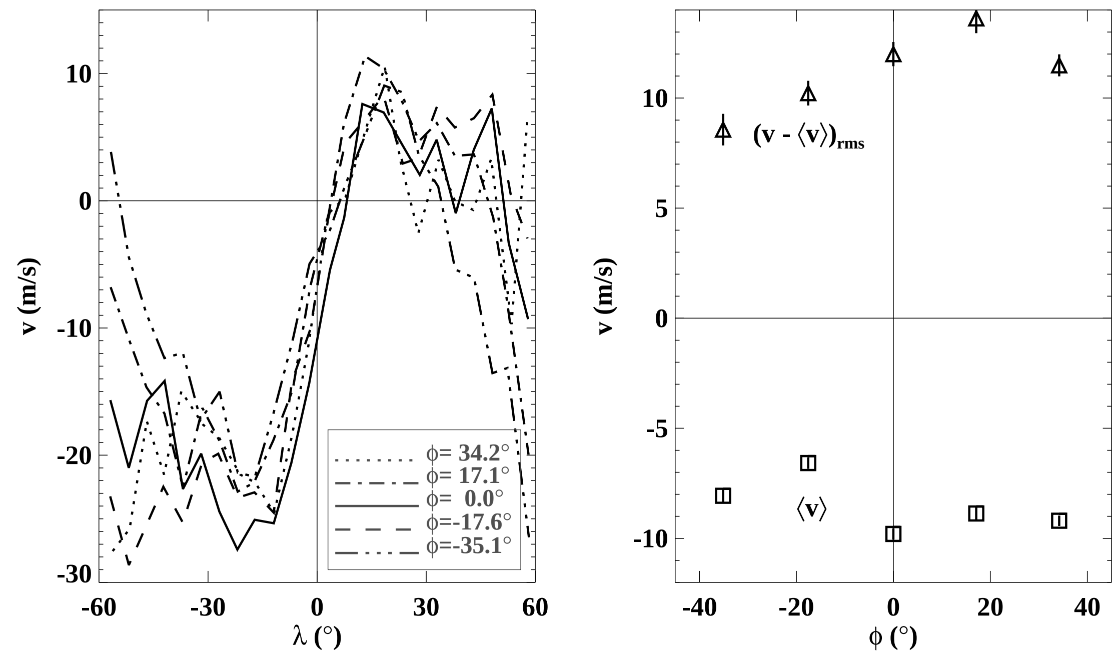

Figure 4.14: The left-hand figure shows the meridional velocity as a function of latitude for five different positions on the disk. The measurements were made for one rotation during the Dynamics run of 1998. The different curves are for different longitudes ϕ, measured westward from the center of the disk. The right-hand plot shows the latitude-averaged mean flow (squares), and the deviation from the mean (triangles), As a function of longitude. Although the mean velocity does not show a strong systematic dependence on longitude, the shape of the profiles clearly varies systematically according to the position of the observed region on the disk.

on the solar surface will appear to accelerate as it rotates across the solar disk.

It should also be noted from figure 4.13 that the effect is clearly more pronounced at some latitudes than at others. If it is indeed due to a non-uniform focus of the image, then it also probably changes when the instrument focus is changed.

The meridional flow measurements also show a similar effect. Figure 4.14 shows the measured time differences as a function of longitude on the disk. While the mean flow does not show a strong systematic dependence on position, the poleward amplitude does increase from east to west. Note that although the spread between the curves is largest at high latitude, the pattern does not appear to be systematic; this scatter is probably due to the large uncertainties near the limb.

In order to account for the uncertainty in the plate scale, in this work some care has been taken to ensure that all of the measurements were made on the same location on the disk. In most cases, the measurements were made during a time period where the remapped region crossed the central meridian. This reduces the chance of relative errors between measurements, but the absolute plate scale is still somewhat uncertain. Thus any comparisons between these measurements and those from different instruments may suffer from a small absolute offset.