

**Exploring the Dynamics of Near-Surface Solar Convection
with Helioseismology**

by

Benjamin J. Greer

B.S., Carnegie Mellon University, 2010

M.S., University of Colorado, 2012

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Astrophysical and Planetary Sciences
2015

This thesis entitled:
Exploring the Dynamics of Near-Surface Solar Convection with Helioseismology
written by Benjamin J. Greer
has been approved for the Department of Astrophysical and Planetary Sciences

Prof. Bradley Hindman

Prof. Juri Toomre

Prof. Mark Rast

Prof. Benjamin Brown

Prof. Keith Julien

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Greer, Benjamin J. (Ph.D., Astrophysics)

Exploring the Dynamics of Near-Surface Solar Convection with Helioseismology

Thesis directed by Prof. Bradley Hindman

I present a new implementation of local helioseismology along with observations of near-surface solar convection made with this method. The upper 5% of the solar radius (35 Mm) is known as the Near-Surface Shear Layer (NSSL) and is characterized by strong rotational shear. While the physical origin of this layer remains unknown, current theories point to convective motions playing an important role. In this thesis I investigate the properties of convection in the NSSL using a newly-developed high-resolution ring-diagram analysis. I present measurements of the speeds and spatial scales of near-surface flows and from these infer that the degree of rotational constraint on convective flows varies significantly across this layer. In-depth analysis of the convective patterns reveals the pervasive influence of coherent downflow plumes generated at the photosphere. These structures link the convective pattern of supergranulation seen in surface observations with the deeper motions found within the NSSL and further hint at the importance of rotation in this layer.

These observations of transient, small-scale convective motions are enabled by the use of improved local helioseismic techniques. Local helioseismology relies on observations of the solar wavefield to produce measurements of plasma flows beneath the surface. In general, this has the capability to map out the subsurface convective flows in three-dimensions, but is often limited in accuracy, resolution, and depth range by the specifics of the analysis procedure. Here, I focus on a particular implementation of local helioseismology called ring-diagram analysis that involves analyzing small patches of the solar surface to build up three-dimensional maps. I will present a new analysis scheme for ring-diagram helioseismology that produces maps of the subsurface flow fields with higher fidelity and vastly higher resolution than previously possible. This is achieved through a combination of novel tools including a robust nonlinear fitting procedure and a highly efficient linear inversion technique. I present these new methods and demonstrate how they enable a new class of high-resolution helioseismic observations. The scientific results made possible with these methods display the power of the new techniques and aid our understanding of near-surface solar dynamics.

Acknowledgements

I would like to thank my advisor Brad Hindman for providing guidance and continually urging me to write things down. Thanks to his dedication and patience, I have learned to be a better scientist. I would also like to thank Juri Toomre for explaining the big picture and asking the big questions. His encouragement has lead me to new and exciting areas of research. I thank Douglas Gough for the suggestion which initiated the expansion of my work into high-resolution helioseismology. Finally, I gratefully acknowledge the University of Colorado's George Ellery Hale Graduate Student Fellowship, which has supported my scientific endeavors for a large fraction of my graduate career.

Contents

Chapter

1	Introduction	1
1.1	Overview of Solar Dynamics	1
1.2	Near-Surface Convection	4
1.2.1	Supergranules	6
1.2.2	Convective Amplitudes	8
1.3	Observing the Solar Interior	8
1.3.1	Global Helioseismology	10
1.3.2	Local Helioseismology	11
1.3.3	Scientific Results from Local Helioseismology	14
1.4	Thesis Overview	17
2	Helioseismic Ring-Diagram Methods	19
2.1	Overview of Standard Processing Steps	20
2.2	Photospheric Observations of the Solar Wavefield	23
2.2.1	Dopplergram Data	23
2.2.2	Tile Tracking	24
2.2.3	Tile Apodization	25
2.2.4	Tile Sets	27
2.3	Representation in Spectral Space	30

2.3.1	Flow-Induced Perturbation to the Wavefield	30
2.3.2	Perturbations to the Power Spectrum	31
2.4	Measuring Shifts in Spectral Power	32
2.4.1	Sensitivity Kernels	37
2.5	High-Resolution Data Sets and Their Properties	38
3	Multi-Ridge Fitting	42
3.1	Introduction	43
3.2	Fitting Methods	45
3.2.1	Single-Ridge Fitting (SRF)	45
3.2.2	Multi-Ridge Fitting (MRF)	47
3.3	Comparison	49
3.3.1	High Phase-Speed Modes	50
3.3.2	Approaching the Solar Limb	52
3.3.3	Frequency Shift Accuracy	57
3.3.4	Post-Processing of Frequency Shift Measurements	61
3.4	Discussion	66
3.4.1	Improvements in Depth for Inversions	66
3.4.2	Spatial Variability and General Implications	67
4	Helioseismic Inversions	69
4.1	Introduction	69
4.2	Mathematical Background	72
4.2.1	Frequency Shift Measurements	72
4.2.2	Methods of Solution	73
4.3	One-Dimensional Depth Inversions	81
4.3.1	Two Box Kernels	81
4.3.2	Solar Wave Mode Kernels	84

4.3.3	Summary of Depth Inversion Behavior	88
4.4	Horizontal Inversions with Spatial Invariance	88
4.4.1	Solution in Spectral Space	90
4.4.2	Regularization Tuning	92
4.5	Three-Dimensional Helioseismic Inversions with MORDI	94
4.5.1	Multi-Dimensional Localization	97
4.5.2	Regularization Tuning	98
4.5.3	Selected Averaging Kernels	106
4.5.4	Conclusions	108
5	Convective Amplitudes and Rotational Influence	110
5.1	Introduction	111
5.2	Methods	114
5.2.1	Measurement Procedure	114
5.2.2	Inversions	115
5.3	Convective Amplitudes	118
5.4	Rossby Number	120
5.5	Discussion	123
5.5.1	Comparison of Convective Amplitudes	123
5.5.2	Rotational Influence on Convective Motions	126
6	Imaging Supergranular Flows in the Near-Surface Shear Layer	129
6.1	Introduction	129
6.1.1	Supergranulation	129
6.1.2	Near-Surface Shear Layer	133
6.2	Matching to Independent Observations	134
6.2.1	Direct Doppler Imaging	137
6.2.2	Advection of Magnetic Field	137

6.2.3	Rotation Rate	139
6.3	Propagation of Convective Patterns	142
6.3.1	Structure Similarity	148
6.3.2	Inferred Vertical Velocities	150
6.4	Discussion	152
6.4.1	Self-Consistent Flow Fields	152
6.4.2	Supergranular Flows and Surface-Driven Convection	153
6.4.3	Rotational Influence on Convection in the NSSL	155
7	Conclusions and Future Directions	158
7.1	Thesis Conclusions	158
7.1.1	Improvements to Ring-Diagram Helioseismology	158
7.1.2	Properties of Near-Surface Solar Convection	159
7.1.3	Influence of Rotation on Convection	160
7.2	Future Directions in High-Resolution Ring-Diagram Analysis	161
7.2.1	Spatial Extent of Analysis Regions	161
7.2.2	Horizontal Resolution	162
	Bibliography	164
	Appendix	
A	Center-to-Limb Velocity Systematic	173
A.1	Introduction	173
A.2	Temporally-Averaged Frequency Shift Data	174
A.3	Fitting Procedure	175
A.4	Results & Analysis	177
A.5	Discussion & Conclusions	178

B	Averaging Kernels for Three-Dimensional OLA Inversion	181
C	Code Listing	186

Tables

Table

2.1	List of analysis periods used in this thesis.	40
-----	---	----

Figures

Figure

1.1	Observations of solar granulation and magnetic activity	2
1.2	Fundamental observations from global helioseismology	2
1.3	Magnetic butterfly diagram	3
1.4	Solar simulations of global dynamics and magnetic structure	5
1.5	Direct Doppler imaging of supergranulation	6
1.6	Images of the Sun captured by SDO	10
1.7	Ring diagrams from a 16° tile.	13
1.8	Meridional flow measured with time-distance helioseismology	15
1.9	Large-scale subsurface flows from ring-diagram analysis	16
2.1	Steps involved in ring-diagram helioseismology	22
2.2	Cuts of tracked and projected Dopplergrams including high-pass and low-pass filtering	26
2.3	Tiling scheme diagram	29
2.4	Azimuthal average of 16° power spectrum	35
2.5	Unwrapped power spectrum in real data, model, and linear filter.	35
2.6	Three-dimensional sensitivity kernels	39
3.1	Ring-diagram power spectrum cuts with ring-fitting guides.	45
3.2	Azimuthal average of power spectrum showing mode blending.	47
3.3	Guess-table frequencies for various tile sizes.	50

3.4	Fit success rate as a function of mode lower turning mode.	53
3.5	Fitting code estimated uncertainty comparison.	54
3.6	Fitting success rate as a function of disk position.	56
3.7	Estimated Uncertainty as a function of disk position.	58
3.8	Velocity response for low-phase-speed modes as a function of disk position.	60
3.9	Velocity response for high-phase-speed modes as a function of disk position.	62
3.10	High-resolution frequency shift map for two measured wave modes	64
3.11	Frequency shift map for wave modes that display anomalous behavior	65
3.12	Comparison of 16° modeset after post-processing.	66
4.1	1D Boxcar Sensitivity Kernels	82
4.2	1D inversion solution versus regularization	83
4.3	1D boxcar inversion averaging kernels	84
4.4	1D solar model sensitivity kernels	85
4.5	One-dimensional averaging kernels for a solar-like inversion	86
4.6	Coefficients and errors for 1D solar inversion	87
4.7	Sensitivity kernel for spatially invariant 1D inversion	89
4.8	Averaging kernels for the spatially-invariant one-dimensional inversion.	93
4.9	Error and resolution of spatially invariant 1D inversion as a function of regularization	94
4.10	Depth profile of 1 Mm 3D averaging kernels versus regularization	99
4.11	Diagnostics for shallow 3D inversion	101
4.12	Power spectra of shallow 3D inversion results	102
4.13	Maps of inversion solution for a range of regularization values	103
4.14	Depth profile of 20 Mm 3D averaging kernels versus regularization	104
4.15	Diagnostics for deep 3D inversion	105
4.16	Power spectra of deep 3D inversion results	106
4.17	Maps of inversion solution for a range of regularization values	107

4.18	Three-dimensional averaging kernels	109
5.1	Inversion averaging kernels for a few depths.	117
5.2	Vector fields of inversion solution at two depths along with observed magnetic field strength at the photosphere.	119
5.3	Velocity spectrum of inversion results.	120
5.4	Root-mean-square horizontal velocity as a function of depth.	121
5.5	Spectral power density of the Rossby number	123
5.6	Rossby number as a function of depth below the photosphere	124
5.7	Comparison of velocity spectra from observations and numerical simulations.	125
6.1	Supergranules as seen in direct Doppler imaging.	130
6.2	Inverted horizontal flow fields and divergence maps for a few depths	135
6.3	Horizontal and vertical size of averaging kernels as a function of depth	136
6.4	Supergranulation seen with direct Doppler and ring-diagram analysis.	138
6.5	Comparison of average magnetic field strength to observed horizontal divergence	140
6.6	Correlation between divergence maps separated by one analysis period as a function of longi- tudinal shift	141
6.7	Cuts of the inferred vertical velocity field.	144
6.8	Distributions of inferred vertical velocity	145
6.9	Correlation of horizontal divergence as a function of time lag and depth separation	146
6.10	Cuts of the time-depth divergence correlation map.	147
6.11	Demonstration of divergence structure correlation	149
6.12	Structure correlation across time and depth	150
6.13	Vertical velocities inferred through the anelastic approximation and the divergence correlation	151
6.14	Measured propagation time from the photosphere to any depth in the NSSL.	156
A.1	Vector maps of large-scale systematics in ring fits.	175

A.2 Examples of multi-component fits for a few wave modes. 177

A.3 Radial systematic amplitude at 45° from disk center as a function of mode frequency and
phase speed. 179

Chapter 1

Introduction

1.1 Overview of Solar Dynamics

The surface of the Sun displays a complex assortment of continually evolving structures arising from highly turbulent magnetized convection. Images of the photosphere reveal a near-uniform covering of small convection cells broken up by the occasional magnetic active region (Figure 1.1). This pattern of convection is termed granulation and is responsible for transporting thermal energy from the interior out into the tenuous solar atmosphere. Granules have a typical size of 1 Mm and consist of broad, warm upflows surrounded by narrow, cool downflow lanes. Individual granules do not live for much longer than a few minutes before succumbing to violent restructuring by other cells within this near-surface convection. Coexisting with the granulation are a collection of longer-lived structures including larger scales of convection (supergranules), meandering zonal jets, and flows associated with the sporadic outbursts of magnetism. This collection of ever-evolving patterns hints at the complexities hidden within the solar interior.

Despite the chaotic appearance, some aspects of the Sun are observed to be surprisingly ordered. The surface rotation rate shows a smooth variation with latitude from a 27-day rotation period at the equator to about a 35-day rotation period near the poles. Measurements of the interior rotation rate from global helioseismology have revealed a set of distinct radial shells (Figure 1.2) (e.g., Brown et al. 1989; Thompson et al. 1996). The radiative zone, from the center to $0.72R_{\odot}$, rotates as a solid body with a rotational period of around 27 days and connects to the differentially rotating convection zone ($r = 0.72R_{\odot}$ to $1.00R_{\odot}$) through a narrow region of strong rotational shear known as the tachocline. The upper 5% of the Sun by radius




Figure 1.1: Left: TiO image of granulation from Big Bear Solar Observatory. The field extends 19 Mm (1.6°) in each direction. The irradiance pattern is dominated by a characteristic scale of solar convection called granulation. Right: Complex magnetic activity seen with the Swedish Solar Telescope. Strong magnetic fields allow for stronger cooling at the solar surface, causing the regions of magnetic activity to look cooler/darker.




Figure 1.2: Left: Power spectrum of photospheric solar oscillations from SOHO/MDI. The ridges of power seen in this ν - l (frequency vs. harmonic degree) diagram are created by the resonant waves of the solar interior. The lowest-frequency ridge corresponds to the surface gravity wave (the f -mode). Each higher ridge corresponds to pressure waves of increasing radial order (p_1 , p_2 , etc.). Subtle variations in these ridges are measured and used to infer subsurface properties. Right: Rotation rate (angular velocity Ω) of the solar interior as a function of depth and latitude measured with global helioseismology (Howe 2009). These measurements reveal a set of distinct radial shells in the solar interior differentiated by rotational properties.

(35 Mm) displays further rotational shear as the interior rotation rate drops by around 3%. This upper boundary of the solar convection zone is known as the near-surface shear layer (NSSL). While the NSSL is the most accessible region of the solar interior to subsurface seismic observations, much about this layer remains a mystery.

Solar magnetism and dynamo theory play prominent parts in our understanding of global solar dynamics. The sporadic appearance of magnetic activity on the surface of the Sun is modulated by the global solar magnetic cycle. Active regions are spawned in belts of magnetic activity that evolve with the 11-year solar magnetic cycle (Figure 1.3). These belts appear at mid-latitudes in each hemisphere and propagate towards the equator over the course of the solar cycle, spawning outbursts of strong surface magnetic field along the way. At the end of a cycle, the belts vanish, only to be replaced again with a new set at high latitude for the next solar cycle.




Figure 1.3: A Magnetic Butterfly Diagram showing longitudinally averaged magnetic field obtained from the Vector Spectromagnetograph (VSM) instrument at Kitt Peak and the Michelson Doppler Imager (MDI) onboard the Solar and Heliospheric Observatory (SOHO). The colors indicate the polarity of surface line-of-sight magnetic field. This demonstrates the nature of the solar cycle as well as some of the large-scale magnetic patterns that exist within the solar cycle. Strong magnetic fields tend to appear within bands of magnetic activity that migrate in latitude over the course of the 11-year solar cycle. Adapted from Hathaway (2010)

The broad range of observed characteristics of the Sun are the result of a highly nonlinear magnetized fluid system. Many of the chaotic actions of the Sun (such as convection) influence the large-scale steady dynamics and vice versa. The patterns of convection seen at the surface likely operate in only the upper reaches of the convection zone and are required to transport thermal energy outwards. Deeper convective motions are likely much larger in size than the patterns seen at the surface and have comparably longer evolution times. These large-scale motions feel the influence of the Sun's rotation as they evolve and con-

sequently provide anisotropic momentum and heat transport. The Reynolds stresses associated with these rotationally influenced turbulent motions help maintain the global differential rotation and indirectly the meridional circulation (Miesch 2005). Thus, there is an intimate connection between the transient convective structures in the solar interior and the long-lived global-scale flows.

These subsurface convective patterns are hidden from view at the photosphere and can only be revealed through helioseismic analysis. This has proven to be exceedingly difficult due to the complexities of the observational procedure, but progress is being made. Lacking precise measurements of deep convection, numerical simulations have provided a way of investigating the balance between convection and global-scale flows. Figure 1.4 displays one such simulation that exhibits rotationally influenced convection interplaying with the global differential rotation and meridional circulation. The benefit of using numerical simulations to discern the interior dynamics of the Sun relies on having a set of precise observational constraints as a guide. Measurements of the global differential rotation profile have for many years provided such a constraint, allowing significant progress to be made on the theoretical front. Advancements in the techniques of helioseismology are providing more extensive and precise measurements of the solar interior that can further enhance our understanding of this highly complex system. To aid in the process of unraveling the intertwined mysteries of rotation, convection, and magnetism, I will present new observations of convection in the solar interior. These observations have direct implications for the dynamical balance that holds within in the Sun and will further our understanding of solar convection.

1.2 Near-Surface Convection

Despite its proximity to observations of the solar surface, the NSSL remains one of the more perplexing large-scale features of the solar interior. While numerical simulations of global solar dynamics have been able to reproduce a wide range of observed properties of the Sun, none so far have been able to reliably produce the strong rotational shear characteristic of the NSSL. A leading theory as to the origin of the NSSL involves the interaction of near-surface convective flows and the global rotation rate. It has been suggested that the distinction between this layer and the bulk of the convection zone comes from the rotational influence on convection. The deep convective motions below the NSSL feel the influence of rotation, while the convective




Figure 1.4: Results from a global simulation using the Anelastic Spherical Harmonic (ASH) code. Left: Panel (a) shows the temporally-averaged angular velocity as a function of latitude and depth, and panel (b) shows the average meridional circulation. This numerical simulation provided guidance on the global patterns of differential rotation and meridional circulation that might be expected in the Sun. Right: Vertical velocity at four different radii within the simulation. Dark colors indicate downflows and light colors indicate upflows. Near the surface, the cell-like structure of convection is clear. At depth, the effects of rotation can be seen with the vertical aligned convection cells (termed banana cells). Adapted from Miesch et al. (2008).

flows in the NSSL are fast enough to traverse the layer without significant deflection due to the Coriolis force (Gilman 1977; Hathaway 1982; Aurnou et al. 2007). The Coriolis deflection of convective motions induces Reynolds stresses that drive the anisotropic transport of angular momentum necessary to maintain the NSSL. While this theory may prove useful for understanding the origin of the NSSL, its confirmation relies on observations of near-surface convection which have proven exceedingly difficult to make.

The observations presented in this thesis weigh on this subject in two ways. First, new measurements of the convective speeds throughout the NSSL provide an estimate of the Rossby number, which is used to describe the degree of rotational influence on flows (see Chapter 5). I find that the NSSL is a transition region for the Rossby number, suggesting that the degree of rotational influence on convective flows does in fact change between the top and bottom of the layer. Further, through the analysis of resolved convective structures in the NSSL, I find that the radial travel time of convective motions in the NSSL roughly matches the rotation rate of the Sun (see Chapter 6). This again suggests that the interaction between rotation and convective motions plays a key role in the dynamics of the NSSL.

1.2.1 Supergranules




Figure 1.5: Northern hemisphere of direct Doppler imaging of supergranulation. The red and blue colors indicate line-of-sight velocity towards and away from the observer. The image was made by averaging 30 minutes of consecutive HMI Dopplergrams and subtracting off a background map made with 24 hours of consecutive Dopplergrams. The prominent spatial scale of 35 Mm (3°) is due to supergranulation. Supergranules have weak vertical flows compared to their horizontal flows, resulting in little observed signal near disk center.

In studying the convective motions of the NSSL, we may begin by considering the largest scale of convection seen in photospheric observations: supergranulation. Supergranulation refers to a particular pattern of convective motions seen in photospheric observations that co-exists with granulation. Photospheric Doppler velocity imaging has been a reliable tool for measuring this prominent scale of motion from the discovery of supergranulation (Hart 1956; Leighton et al. 1962) through modern high-precision analysis (Hathaway et al. 2000; 2002). This pattern consists of cell-like outflow sites surrounded by narrow downflow lanes. Figure 1.5 demonstrates the horizontal flow pattern of supergranulation in direct Doppler images averaged over 30 minutes. The horizontal size of a supergranule can range between 20 Mm and 60 Mm with a typical size around 35 Mm (Hathaway et al. 2000; 2002; De Rosa and Toomre 2004; Del Moro et al. 2004; Hirzberger et al. 2008; Rieutord et al. 2008). Like granulation, the pattern of supergranulation is constantly evolving. While it is difficult to identify the exact start or end time of a particular supergranule, measurements of their temporal coherency provide an estimate for their lifetime of around 1.7 days (e.g., Hirzberger

et al. 2008). The flows associated with supergranulation are primarily horizontal, with Dopplergrams revealing horizontal speeds of 300 m s^{-1} (Hathaway et al. 2002) and vertical speeds of 30 m s^{-1} (Hathaway et al. 2002; Rieutord et al. 2010). A more extensive review of the observed properties of supergranules can be found in Rieutord and Rincon (2015) and Gizon and Birch (2005).

Despite the wealth of information about the surface characteristics of supergranulation, the physical origin of this prominent scale of convection is unknown. This is due in part to significant uncertainty in the vertical profile of supergranules. The subsurface structure of supergranulation has been the subject of a multitude of studies using local helioseismology and no agreement has been reached. A few of these studies report a reversal in the horizontal divergence pattern, indicating a cell-like return flow (Duvall 1998; Zhao and Kosovichev 2003). The depth of this return flow is highly disparate between studies, ranging from very shallow (1 Mm, Rieutord et al. 2010) to rather deep (15 Mm, Zhao and Kosovichev 2003). The disagreement on the depth of a return flow is caused by the difficulty of the measurement. Interpreting a reversal of the flow pattern in depth must be done with care as many of the local helioseismic methods are susceptible to contamination from leakage of the signal from surface flows into deductions about deeper layers (Braun et al. 2004). Due to the wide range of answers for the depth of supergranulation, the precise subsurface structure of supergranulation is still unclear.

Further obscuring the source of supergranulation is the difficulty in identifying their origin and predicting their properties (particularly their distinct size) through theoretical means. Several models have been proposed that select a single length scale amidst the highly turbulent convective background (see Rieutord and Rincon 2015, for a review), but none so far have been supported by strong observational evidence. One of the first theories on the origin of supergranulation is that He^{++} recombination in the upper convection zone acts as a driver of this characteristic scale of convection (e.g., Simon and Leighton 1964). Schwarzschild (1975) further supposed that supergranulation could be driven at depth by an opacity break, He^+ recombination, and H^+ recombination. These theories share a common idea that supergranules are driven at their deepest point by an excess of thermal energy. Competing theories claim that supergranulation can be generated at the surface either by an instability of the upper boundary of the convection zone (Murphy 1977; Rincon and Rieutord 2003), the interaction and merger of downflow plumes from granulation (Rieutord

and Zahn 1995; Rast 2003), or possibly by a characteristic scale set by the advection of magnetic field in the photosphere (Crouch et al. 2007). Due to a lack of observational support, none of these theories have been confirmed. The advancement of high performance supercomputers has led to increasingly sophisticated numerical models of near-surface convection (e.g., Stein and Nordlund 1998; Rempel et al. 2009; Hotta et al. 2014), yet a simple and consistent explanation of supergranulation remains elusive. Despite the struggles of observing the subsurface structure of supergranules, precise helioseismic measurements will likely play a key role in unraveling the mysteries of this prominent scale of motion. The new observations presented in this thesis reveal much about the structure of supergranulation and hint at a surface-driven origin. These results will be presented in Chapter 6

1.2.2 Convective Amplitudes

Extending the measurements of near-surface convection into the deeper interior, the speed of convective flows throughout the NSSL has become a topic of interest. Using time-distance analysis, Hanasoge et al. (2012) presented the first measurements of horizontal convective amplitudes at significant depth (30 Mm and deeper). In this, the flows were found to be exceptionally weak compared to predictions made by essentially all numerical convection simulations (Vögler et al. 2005; Miesch et al. 2008; Rempel et al. 2009; Trampedach and Stein 2011; Hotta et al. 2014). The discrepancy between the observations of Hanasoge et al. (2012) and theory is a difference of nearly two orders of magnitude. This has caused significant confusion in recent years, as simulations typically rely on much faster convective amplitudes to generate the Reynolds stresses needed to maintain the differential rotation profile which is constrained by global helioseismic observations. However, in this thesis I will present new measurements using ring-diagram analysis of the convective amplitudes throughout the NSSL that generally agree with numerical simulations. Chapter 5 presents these results along with a discussion of the discrepancy between the two sets of observations.

1.3 Observing the Solar Interior

Figures 1.6a and b show full-disk snapshots of the Sun in continuum intensity (a) and 304 Å (He^{++}) (b). These images were captured with the Helioseismic and Magnetic Imager (HMI) and the Atmospheric

Imaging Assembly (AIA) aboard the Solar Dynamics Observatory (SDO) which has provided unprecedented coverage of solar activity in recent years. Despite the wealth of information contained in these images, there is a fundamental limit to how deep into the Sun one can make direct photometric or spectroscopic observations. This limit is set by the transition from an optically thin atmosphere to an optically thick atmosphere, and the radius at which this transition occurs varies with the wavelength with which one chooses to observe the Sun. These types of observations can provide invaluable information on the structure of the corona and chromosphere, but these observations cannot penetrate the photosphere. Non-photometric observations such as line-of-sight magnetic field strength (Figure 1.6c) and line-of-sight Doppler velocity (Figure 1.6d) can be made to enhance our understanding of the solar surface, but these measurements are also relegated to photospheric layers and higher. Without a method of directly observing the solar interior, we must rely on more indirect measures.

Solar oscillations with a period of around 5 min (3 mHz) were first detected by Leighton et al. (1962) using Doppler velocity measurements of the photosphere. These oscillations are the observable signal of resonant acoustic waves trapped beneath the photosphere, and they are generated by near-surface turbulent convection (Goldreich and Keeley 1977; Libbrecht et al. 1986; Goode et al. 1992). These acoustic waves are a window to the solar interior, as they relay information about the structure of the interior out to the surface of the Sun where we observe them. As seen in Figure 1.2, the Sun has a rich spectrum of resonant waves with well-measured properties. Different wave modes sample different portions of the Sun, and the wavenumber and frequency of a given wave mode determines the lower turning point of the wave (the maximum depth to which the wave can propagate). Low wavenumber and high frequency waves penetrate deepest into the Sun. As the waves travel through the interior, they can be modified by various effects such as thermal variations and local flows. The alterations to these waves can be directly observed as they reflect from the photosphere back into the interior. Helioseismology relies on observing the vast assortment of waves visible in the photosphere, analyzing their properties to deduce how they have been modified, and then inferring the subsurface structure of the Sun that caused the modifications. Measurements from many distinct wave modes with different lower turning points can be combined to infer variations in the subsurface flow field with depth.




Figure 1.6: Four full-disk images captured by instruments onboard the Solar Dynamics Observatory (SDO) on 31 August 2012. (a) Continuum intensity (artificially colored). (b) 304 Å wavelength band. (c) Line-of-sight magnetogram, colors indicate regions of strong magnetic field (blue/green is positive, red/yellow is negative). (d) Line-of-sight Dopplergram. The overall gradient in this image is due to the rotation of the Sun. A region of strong magnetic activity (with a sunspot) can be seen in the continuum intensity as a dark region, in the 304 Å band as a bright region, and in the magnetogram as a large complex region of mixed polarity. The active region is less obvious in the Dopplergram image.

The ability to accurately infer the subsurface characteristics of the Sun relies on both careful modeling of the acoustic wave field and precise numerical tools to measure small changes in the observations. The fidelity of helioseismic results depends critically on the accuracy of these analysis components. Over time, a wide variety of helioseismic techniques have been constructed to transform photospheric observations of the solar wave field into inferences of subsurface structure. The specifics of these methods vary considerable but can generally be grouped into two categories: global and local methods.

1.3.1 Global Helioseismology

Global helioseismology focuses on the largest and longest-lived wave modes of the Sun and is used to estimate the properties of the solar interior as a function of radius and latitude (e.g., Deubner and Gough

1984; Gough and Toomre 1991; Christensen-Dalsgaard 2002; Thompson et al. 2003). The observed full-disk wavefield is analyzed with spherical-harmonic transforms and a variety of properties of the resonant wave modes are measured from the power spectrum of the spherical harmonic coefficients. Global analysis can measure the deepest-traveling modes and therefore allows measurements of the solar interior from the photosphere to the nuclear-burning core. Perhaps the most notable result from global helioseismology is the accurate determination of the Sun’s core temperature, which played a critical part in the discovery of neutrino oscillations (for a review, see Christensen-Dalsgaard 2002). As mentioned in the previous section, global helioseismology has also provided accurate measures of the internal rotation rate of the Sun as a function of depth and latitude (Howe et al. 2000; Thompson et al. 2003). These results have revealed the existence of the tachocline, solid body rotation of the radiative interior, and the shear in the NSSL. Latitudinal variations in the rotation profile due to torsional oscillations have been studied with this method as well (Woodard and Libbrecht 1993; Howe et al. 2000; Toomre et al. 2000). Originally detected in direct Doppler imaging (Howard and Labonte 1980), torsional oscillations manifest as bands of fast zonal flow in each hemisphere that migrate equatorward over the solar cycle.

The eigenfunctions of global modes, which determine the volume of the solar interior over which a given measurement pertains, have an energy density which is longitudinally invariant and symmetric about the equator. Thus, properties of the interior can only be measured as a function of depth and latitude (symmetrized about the equator). Small-scale flow patterns, like that of convection, are averaged out in global helioseismic analysis. Further, the meridional component of the subsurface flows do not significantly influence the frequencies of the wave modes. Therefore, traditional global mode analyses which rely on frequency measurements are only sensitive to the zonal component of the axisymmetric flow field.

1.3.2 Local Helioseismology

An alternative form of analysis is that of local helioseismology, which enables measurements of small-scale flows that do not respect the axisymmetric and hemispherically-symmetric constraints of global analysis. With local helioseismology, transient convective flows can be mapped out in longitude, latitude, and depth. While global helioseismology is able to measure properties of the solar interior at great depths, local

helioseismology is typically restricted to the upper reaches of the Sun. However, the ability of local helioseismology to probe the structure of small-scale flows makes it an excellent tool for studying the dynamics of the near-surface shear layer.

There are various implementations of local helioseismology that have been developed over the years. All methods generally use the same source data of direct Doppler imaging of the Sun, but differ in how the data is analyzed, what aspect of subsurface flows are considered, and what assumptions are made along the way. While the choice of local-helioseismic method often depends on the scientific goals at hand, the different methods have significant overlap in what they are capable of measuring. The two most common types of local helioseismology are ring-diagram analysis (Hill 1988; Basu et al. 1999; Haber et al. 2002) and time-distance analysis (Duvall et al. 1993; 1997). Both methods are able to map the horizontal components of subsurface flows across the solar disk and through a range of depths. The exact depth range, horizontal resolution, and overall uncertainty of these flow measurements depend on how the analysis is carried out and is not pre-determined by the general method itself. The observations of near-surface convection presented later in this thesis were produced with ring-diagram analysis, and so I will provide a brief overview of the method in the next section.

Ring-diagram and time-distance analysis are far from the only possible methods of doing local helioseismology. Fourier-Hankel Decomposition considers wave motion into and out of localized points on the solar surface and has been used to measure absorption, scattering, and mode mixing of acoustic power by sunspots and plage (Braun et al. 1987; 1992; Braun 1995). Acoustic holography uses the observed wavefield at the photosphere to measure subsurface flows and to isolate localized regions of acoustic emission (Lindsey and Braun 1997). One of the most interesting applications of acoustic holography is far-side imaging, where active regions on the far-side of the Sun can be detected using the near-side wavefield (Lindsey and Braun 2000; Braun and Lindsey 2001; Gonzalez-Hernandez et al. 2007).

1.3.2.1 Ring-Diagram Analysis

Ring-Diagram analysis infers subsurface horizontal flows in localized patches of the solar disk by measuring the properties of three-dimensional (two spatial dimensions plus time) Fourier power spectra created

from each patch (Hill 1988; Schou et al. 1998; Haber et al. 2002). This is similar to global helioseismology where distinct wave modes are analyzed in the spherical-harmonic spectrum. Here, small regions of the solar surface are extracted over some time span to create a local analysis region that can be treated as a local Cartesian domain analyzed in Fourier space. The modal structure of the observed wave field can be seen in cuts of a three-dimensional power spectrum at constant frequency. Figure 1.7 shows three such cuts for a single power spectrum. The details of how this kind of power spectrum is generated are provided in Chapter 2. At a constant frequency, the different radial orders of solar wave modes appear as concentric rings (the source of the name ring-diagram analysis). The outermost ring corresponds to the surface gravity mode (f -mode) and each subsequent inner ring corresponds to pressure waves with increasing radial order (p_1 , p_2 , etc.).




Figure 1.7: Constant-frequency cuts from the power spectrum of a 16° analysis region. In each, power from modes at each radial order show up as concentric rings. The outermost ring in each panel is the surface gravity wave (f -mode). Darker colors indicate higher power on a logarithmic scale. The rings are shifted slightly in the $+k_x$ direction due to a subsurface flow of around 20 m s^{-1} .

Sub-surface flows underneath a ring-diagram analysis region induce Doppler shifts in the wave modes captured in the power spectrum, as in Figure 1.7. Seen at a constant frequency, these mode rings become elongated in the direction corresponding to the direction of subsurface flow. Measuring the shift of spectral power for each mode individually provides information about the subsurface flows at different depths. Spatial information about flows across the solar surface and below can be recovered by collecting such measurements from analysis regions extracted from different disk positions.

In this thesis, I use ring-diagram analysis to measure convective flows in the NSSL. One of the historic limitations of ring-diagram analysis is a trade-off between depth and resolution. The size of the analysis region used to create a power spectrum and frequency-shift measurements determines not only the amount of spatial averaging affecting the measurement (and thus the horizontal resolution), but also the maximum depth that can be measured reliably. Larger tiles provide more robust measurements of the low-wavenumber modes that travel deepest into the solar interior, but do so with coarse horizontal resolution. A general rule of thumb for this trade-off is that the size of the analysis region needed to reach a certain depth is around 6 times that depth. In the following chapters, I will discuss a new method for beating this trade-off and producing high-resolution flow fields at depth.

1.3.3 Scientific Results from Local Helioseismology

Local helioseismology has played an important role in our understanding of solar dynamics by permitting measurements of the structure of the interior. The tools of local helioseismology have been applied to a vast range of topics; from small-scale convective dynamics to large-scale global flow patterns. To get an idea of the range of topics that local helioseismology is capable of investigating, I will provide a brief overview here of some key discoveries made possible with these methods. An in-depth look into the explorations of local helioseismology can be found in Gizon and Birch (2005).

1.3.3.1 Large-Scale Flows

Local helioseismology has been regularly used to measure large-scale, axisymmetric flow patterns by averaging measurements across longitude to reduce noise. The large-scale variations in the differential rotation profile due to torsional oscillations can be seen using this method. Torsional oscillations have been measured extensively with both time-distance analysis (Giles et al. 1998; Beck et al. 2002; Zhao and Kosovichev 2004) and ring-diagram analysis (Basu et al. 1999; Basu and Antia 2000; Haber et al. 2000; 2002). Local helioseismic observations of these bands show that they are only a few m s^{-1} faster than the normal differential rotation rate, have a width in latitude of around 10° , and extend from the surface of the Sun to a depth of around 50 Mm (Kosovichev and Duvall 1997; Howe et al. 2000; Toomre et al. 2000; Antia




Figure 1.8: Meridional flow measured with time-distance helioseismology showing poleward motion near the surface and reversals deeper down. Positive velocity (yellow/red) is directed northward and negative velocity (cyan/blue) is directed southward. Figure adapted from Zhao et al. (2012).

and Basu 2000; Howe et al. 2006).

The surface meridional flow was first seen with direct Doppler imaging in Hathaway (1996). Using time-distance analysis, Giles et al. (1997) extended this measurement below the photosphere and found poleward flow throughout the upper 4% of the interior. Using mass conservation to extrapolate the flow field to the bottom of the convection zone, it was found that a large, single-cell circulation of meridional flow with a return flow near the bottom of the convection zone matched the observed data most simply. When used in a flux transport dynamo model, the shape and speed of this circulation could help reproduce the time-scale of the solar magnetic cycle (Dikpati and Charbonneau 1999). The basic reason that the observed meridional flow provides such a pairing is that the time it takes to complete one full circulation through the convection zone (based on the observations of Giles et al. 1997) roughly matches the timescale of the solar cycle. The structure of the global meridional circulation has been a continual target of local helioseismology (Basu et al. 1999; González Hernández et al. 1999; 2000; Haber et al. 2000; Zhao and Kosovichev 2004; Duvall and Hanasoge 2009; Zhao et al. 2013). Recent measurements of the meridional flow have suggested that there is in fact a reversal of the circulation well before the bottom of the convection zone (Hathaway 2011; Zhao et al. 2012; Rajaguru and Antia 2015). Figure 1.8 demonstrates the complicated multi-celled structure of the meridional flow seen in Zhao et al. (2012) using time-distance analysis.




Figure 1.9: Synoptic maps of subsurface flows at a depth of 0.9 Mm using ring-diagram analysis. Underlying the flow vectors is a map of the surface magnetic field with polarities indicated in red and green. The flows display large-scale, sweeping patterns as well as a significant amount of convergence towards the regions of high magnetic activity. Figure adapted from Hindman et al. (2004).

1.3.3.2 Small-Scale Flows

The local environment around magnetic active regions has also been the topic of extensive studies using local helioseismology. Figure 1.9 shows large-scale mapping of subsurface flows using ring-diagram analysis (vectors) on top of a map of photospheric magnetic field strength (colors). In this, the global flow pattern of meridional circulation and differential rotation are broken up by the presence of active regions. Near-surface flows have been found to display a weak convergence towards active regions as far as 30° away (Gizon et al. 2001; Haber et al. 2001) along with Coriolis-induced rotation (e.g., Hindman et al. 2009). The subsurface structure of flows around active regions has been studied with both time-distance analysis (Zhao and Kosovichev 2004) and ring-diagram analysis (Haber 2004). These same studies found that the horizontal inflow surrounding active regions reverses direction past a depth of around 10 Mm.

Small-scale measurements made with local helioseismology are often plagued by high noise levels. This problem can be mitigated by averaging over time and/or space, but this can be problematic when studying the transient, small-scale flows of near-surface convection. The flows associated with supergranulation can be directly observed in the upper megameter or so without significant averaging (see Chapter 6), but deeper flows are commonly below the noise level, as the waves that propagate deeper are more difficult to measure. Studies of supergranule flows often consider the structure of an ‘average supergranule’, where the

flow patterns observed for different supergranules are averaged together relative to the observed cell center at the photosphere (Duvall and Birch 2010; Švanda 2012; Duvall and Hanasoge 2013; Duvall et al. 2014; DeGrave and Jackiewicz 2015; Langfellner et al. 2015b). This procedure has led to the possible discovery of a supergranular return flow somewhere in the upper 15 Mm of the Sun.

1.4 Thesis Overview

This thesis presents significant contributions to the study of the dynamics of the solar interior. These contributions can be classified into two broad categories: improvements to ring-diagram analysis and new observations of near-surface convection. While the new observational methods have been designed to obtain measurements of supergranulation in the NSSL, they can be applied to a broader range of scientific investigations. Thus, the details of this new implementation of ring-diagram helioseismology is presented distinctly from the resulting scientific observations.

Chapters 2 through 4 are dedicated to revamping the methods of ring-diagram analysis with an emphasis on accuracy, depth sampling, and resolution. While any local helioseismic method incurs a tradeoff between resolution and depth, I present new methods that push this boundary to a point where new scientific discoveries can be made. This work focuses on resolving flows in the NSSL, but has the potential to be extended much deeper and to higher resolution.

Chapter 2 provides an overview of the ring-diagram analysis steps. Chapter 3 presents a novel method of measuring the frequency shifts from ring diagrams, and delves into both a comparative characterization of the method and the implications for subsurface flow results. Chapter 4 provides a detailed description of the inverse problem inherent in helioseismic analyses, along with a presentation of a new high-resolution inversion procedure.

In Chapters 5 and 6, I use these new methods to answer scientific questions about solar dynamics and present new observations of near-surface convection. In the view of solar dynamics, there are many components that build up the full picture. A few of these components have been determined with enough certainty to provide rigid constraints on models (rotation profile, magnetic bands, solar cycle). The characteristics of subsurface convection are much less certain, and provide a key linkage between the dynamics

and other aspects. In this thesis, I provide the first measurements of convective flows in the NSSL that have the resolution to infer the scales responsible for convective transport. Analysis of these flows has led to a deeper understanding of the role the upper boundary of the convection zone plays in global solar dynamics.

Chapter 5 presents measurements of convective flow speeds throughout the NSSL along with a measure of the degree of rotational influence at each depth. These measurements provide constraints on the bulk characteristics of near-surface convection and are easily relatable to theoretical models of solar dynamics. Chapter 6 presents evidence for the surface-driven nature of supergranulation as well as radially-coherent downflow plumes that traverse the NSSL. The precise nature of the convection seen in the NSSL further suggests a connection between the depth of the NSSL and the rotational influence felt by downflows.

Chapter 2

Helioseismic Ring-Diagram Methods

The scientific results presented later in this thesis were generated by the local-helioseismic technique of ring-diagram analysis, which produces subsurface flow measurements from a set of Dopplergram observations of the solar surface. The processing steps involved in transforming the latter into the former can be accomplished in a multitude of ways, and the details of the specific steps used must be understood to properly interpret the final results. A primary goal of this thesis is to analyze each step of the ring-diagram process with an eye for improvements and optimizations so that new scientific discoveries can be made. In particular, the goal of measuring transient, small-scale convective flows throughout the NSSL provides stringent requirements for the accuracy, resolution, and depth range of the analysis. The ability to measure these patterns depends crucially on how well the ring-diagram procedure can produce spatially localized estimates of the subsurface flow.

Inherent to the processing of ring-diagram analysis is the assumption of linearity in each step. In short, we rely on the ability to express the final set of velocity observations $\mathbf{w}$ made with ring-diagram analysis as a linear integral over the true subsurface flow field $\mathbf{v}$:

$$\mathbf{w}(\mathbf{x}, z) = \int \int d\mathbf{x}' dz' Q(\mathbf{x}, \mathbf{x}', z, z') \mathbf{v}(\mathbf{x}', z'). \quad (2.1)$$

The vector $\mathbf{x}$ contains the two horizontal directions (x, y) , and the Cartesian coordinate directions $\hat{\mathbf{e}}_x, \hat{\mathbf{e}}_y, \hat{\mathbf{e}}_z$ map locally to the longitudinal, latitudinal, and radial directions, respectively. Here, the function Q (called the averaging kernel) is the weighting function that relates the measured flow field $\mathbf{w}$ with the true flow field $\mathbf{v}$. While the spatial integral is taken over the entire volume of the Sun, we will strive to make the function

Q as compact as possible. Our ability to interpret the flow field $\mathbf{w}$ as a representation of $\mathbf{v}$ depends on the form of Q , which can only be derived by assuming each intermediate data product in ring-diagram analysis is linearly related to the subsurface flow field. While the processing steps themselves may not be linear (as with the ring-fitting code in Chapter 3), it is imperative to find a valid linear approximation.

This chapter will serve as both an overview of the steps carried out in ring-diagram analysis and a detailed description of the first few steps. I will not present any scientific findings here, but I will justify the modifications made to various processing steps with the scientific goals laid out in Chapter 1. To achieve the goal of measuring convective flows in the NSSL, I have modified some of the standard processing steps (such as the measurement grid shown in this chapter) and completely revamped others (such as the mode-fitting method and linear inversion presented in the following chapters).

2.1 Overview of Standard Processing Steps

The entire processing pipeline of ring-diagram helioseismology has been constructed incrementally since its inception and has become a standardized process (Bogart et al. 2011a). Figure 2.1 shows the series of steps necessary to produce subsurface flow measurements using ring-diagram analysis. A telescope observes the Sun and produces a set of Dopplergrams, which measure the line-of-sight velocity on the surface at each point on the solar disk. A set of consecutive Dopplergram images constitute the entire data source for a single analysis period. While we are free to choose the duration over which the Dopplergram set is collected, the work in this thesis will treat a single analysis period as one spanning 25.6 hours of solar observations. In the set of Dopplergrams, the surface of the Sun is split up into distinct sub-regions called tiles that are independently analyzed. The characteristics of the solar wave modes within each region are analyzed and measured with a non-linear fitting code operating on the power spectrum of the tile. While many properties of the wave modes can be extracted, the measurement of interest is the amount by which the frequency of the mode Doppler shifts relative to a central rest frequency. This data exhibits a large-scale systematic bias that is removed before further processing. The frequency-shift measurements made in each analysis region are then recombined in a linear inversion to produce a three-dimensional subsurface flow field that spans a large portion of the solar surface and below. This process is repeated independently for each

25.6-hour analysis period.

In this chapter, I will describe both the numerical procedure required to achieve the first few steps of ring-diagram analysis (Figure 2.1a-c) as well as a rough outline of how the linear relationship in Equation 2.1 can be justified. While the details of the final two steps (Figure 2.1d-e) are covered in the next two chapters, I will briefly discuss them here in the context of the desired linear relationship so that the function Q can be examined later. In this chapter, I will follow the linear relationship up to the point of relating the measured frequency shifts to the subsurface flow field. Relating the final measured frequency shifts to a subsurface flow can be done by following the effect of a subsurface flow on each of the intermediate data products of ring-diagram analysis incrementally. The full calculation to derive this particular part of the relationship is presented by Birch et al. (2004) and Birch et al. (2007). I will not attempt to reproduce the full derivation here, as it is lengthy and adds little to the current discussion of ring-diagram methods. Instead, I will discuss the assumptions that allow such a connection to be made, as well as the main conclusions of the derivation.

The methods presented in this chapter are largely unchanged from what is considered standard ring-diagram processing (Bogart et al. 2011a), while those presented in the next two chapters are novel methods created for the purpose of achieving the scientific goals of this thesis. I will also discuss the benefit of greatly increasing the number of measurements made over the solar disk in order to obtain the high-resolution flow maps needed for measuring small-scale convective patterns.




Figure 2.1: The steps involved in ring-diagram helioseismology. We can consider ring-diagram analysis as a sequence of processing steps that transform one data product into another. Here, the data products are in boxes and the processes are arrows. The data with which we start the process is a set of full-disk Dopplergrams (a). Dopplergram tiles (b) are created from this data set by tracking and projecting a mosaic of square regions. These tiles are then Fourier transformed and made into three-dimensional power spectra (c). The power spectra are then measured with a non-linear fitting code to produce frequency-shift measurements (d). Up to this point, each tracked and projected tile has been processed independently. The measured frequency shifts exhibit large-scale systematics as a function of disk position which can be measured and removed (e). The final step combines all of the frequency shifts in a single inversion to produce an estimate of the subsurface flow field (f). The colors of the boxes and arrows indicate the chapter of this thesis in which the topic appears in.

2.2 Photospheric Observations of the Solar Wavefield

2.2.1 Dopplergram Data

Full-disk Dopplergrams form the sole observable used in this implementation of ring-diagram analysis. The Dopplergrams are measurements of the line-of-sight velocity made with the Helioseismic and Magnetic Imager (HMI) onboard the Solar Dynamics Observatory (SDO). The HMI instrument captures full-disk images in 6 narrow pass-bands around the Fe I absorption line at 6173.3 Å. Line-of-sight Doppler measurements at each pixel can be constructed through a linear combination of these pass-band images that are sensitive to shifts in the Fe I line position (Couvidat et al. 2012). The precision of this Doppler shift measurement at each pixel is 13 m s^{-1} and the dynamic range is $\pm 6.5 \text{ km s}^{-1}$. HMI records a 4096x4096 pixel full-disk Dopplergram once every 45 seconds. This pixel count leads to a resolution of 500 km per pixel at disk center. The height of formation of this absorption line is estimated to be around 200 km above the photosphere (Fleck et al. 2011).

To simplify the mathematical analysis of the observed wavefield, the line-of-sight Doppler signal is projected from the full-disk observations into local Cartesian coordinates. To construct a three-dimensional data cube with coordinates (x, y, t) , we first collect a sequence of consecutive full-disk Dopplergrams that span the time of interest. In this thesis, I use 2048 sequential Dopplergrams to represent one analysis period. With HMI's 45-second cadence, this leads to a duration of 25.6 hours. Ring-diagram analysis implicitly averages the subsurface flow over the analysis period, so it is important to keep in mind that any flows we wish to accurately measure must persist for at least that long. Shorter-lived features will be diminished proportional to the ratio of their lifetime to the analysis duration. Averaging over a shorter span of time would reduce this effect, but also increase the noise in the solution. The specific analysis period used here was chosen as a trade-off between noise and temporal averaging.

With a time assigned to each full-disk Dopplergram, the Cartesian data cube is then constructed by extracting a square projected region from each Dopplergram. The region may be centered around any chosen point on the disk and extends horizontally to a specified size. Each Dopplergram is projected using an equidistant azimuthal (Postel) projection (with bicubic interpolation) centered around the square region's

center. This projection preserves both distances along great circles that intersect the center of the region as well as the azimuthal angle of points relative to the center. Since we will be measuring plane waves traveling in different directions within the region, this projection is appealing because it preserves propagation direction and wavelength for waves near the center.

The size of the region is typically chosen somewhere between 2° (24 Mm) and 32° (388 Mm) in heliographic angle. The size of the region determines the error statistics and depth to which flows may be measured. We will find in later chapters that 16° tiles are able to measure flows down to a depth of around 30 Mm, which satisfies the goal of measuring flows throughout the NSSL. Larger tiles are able to measure deeper into the Sun, but require more computational resources. Further, ring-diagram analysis relies on a plane-wave approximation for the observed wavefield which assumes no significant sphericity. Larger tiles ($\gtrsim 30^\circ$) break this assumption and introduce systematic errors into the analysis. Initially, the horizontal resolution of the flow measurements is limited by the size of the tile. We will find in later chapters that it is possible to beat this resolution limit by clever inversion procedures.

2.2.2 Tile Tracking

Since each Dopplergram is projected independently into the final data cube associated with the chosen tile, the projection centers are allowed to change with time. Shifting the region's center linearly as a function of time is a Galilean transform, and is useful for either introducing a known velocity into the region or countering an existing known velocity. This shifting of central position with time is called tracking and plays a part in ensuring the linear relationship between the ring-diagram measurements and subsurface flows. The surface rotation rate of the Sun is well-known and not of any particular interest in many ring-diagram studies. We thus track in longitude at the surface rotation rate to subtract out the large-scale advection of waves by the Sun's rotation. Globally, this rotation velocity is around 2 km s^{-1} and is called the Carrington rotation rate. The Sun's differential rotation introduces yet another large-scale velocity signal that can be removed by using a tracking rate that depends on latitude. The latitudinal profile of differential rotation has been measured fairly precisely (Snodgrass 1984, and many others) and can be used as a tracking rate. The effects of tracking at the Carrington rate or one that accounts for differential rotation are considered

in the following chapter. Tracking in latitude is also possible, but has fewer uses. While a region may pass through a range of longitudes over the tracking duration, in practice it is assigned a central location equal to the position of the tile center exactly half-way through the analysis duration.

The final Cartesian data cube consisting of tracked and projected Dopplergrams is called a tile. Figure 2.2 shows cuts through such a tile. The tile contains many different signals, some useful for our purposes and some not. Figures 2.2a and 2.2d demonstrate that the direct Doppler field is a complicated pattern that is hard to describe as a simple wavefield. When filtered in time, two distinct signals emerge. Figures 2.2b and 2.2e show the high-pass-filtered Doppler signal with a frequency cut-off of 1.5 mHz. In this, the low-frequency components of the signal have been removed. As a function of time and space, this part of the Doppler signal shows wave packets propagating short distances before disappearing. It is this wavefield we are interested in measuring, as it contains information on the subsurface flows. Figures 2.2c and 2.2f show the low-frequency contribution to the signal using the same cut-off frequency as before. In space, granules can be seen as packed cells with compact upflow sites. This view of granulation is far less detailed than that seen in Figure 1.1 due to the lower resolution of HMI. Here, there are thousands of individual granules in the field of view, and each granule is only around 2-3 pixels across. In time and space, these granules can be seen advecting and merging with each other. While the frequency filter used in Figure 2.2 seems to isolate this signal from the wavefield fairly well, granules have a broad frequency spectrum and cannot be separated so cleanly. When measuring the properties of the wavefield in the following chapter, I do not use filters such as this and instead account for background granular power with a model.

2.2.3 Tile Apodization

The properties of the wave modes observed in a given tile will be measured from the Fourier power spectrum of the tile. Before computing the three-dimensional Fourier transform of a tile, we apodize the tile in both space and time. This is done to remove anisotropies in the power spectrum due to the square shape of the tile, mitigate spectral ringing, and to control the inevitable spatial leakage.

In space, the apodization function is defined as a function of radial distance r from the geometric center of the tile, resulting in a circular apodization shape. This function is equal to one at $r = 0$ and




Figure 2.2: (Top row) Line-of-sight Dopplergram data as a function of position on the solar disk near disk center. The left column shows raw Dopplergram data, while the middle and right columns show the same data high-pass-filtered and low-pass filtered in time, respectively. (Bottom row) The same Dopplergram data as in the top row, but shown as a function of position and time. Note the different ranges on the vertical axis for each column. The high-pass-filtered data show the surface wavefield, while the low-pass filtered data barely resolves the pattern of granulation evolving over time. Bright shades indicate a line-of-sight velocity away from the observer, dark indicates towards.

remains unity out to $r = 0.875 R_t$, where R_t is the distance between the center of the tile and the edge of the tile along one of the cardinal directions. Between $r = 0.875 R_t$ and $0.9375 R_t$, the apodization function

tapers smoothly to zero with a quartic function of radius r . Past $r = 0.9375 R_t$, the apodization function remains zero. Thus the tapering of the apodization function occurs over $0.0625 R_t$, which for a 16° tile is 0.5° . The temporal apodization is similar in that it is equal to unity for most of the tile duration and smoothly tapers to zero at the beginning and end of the temporal duration with the same quartic function. The interval over which the temporal apodization is tapered is 0.03125 times the full duration, or 48 minutes of time for the analysis durations used here.

2.2.4 Tile Sets

A single tile provides measurements of the subsurface flow field from the region below the tile, but reconstructing a more complete picture of the solar interior requires measurements from many tiles. Spatial information about the subsurface flows can be gained by extracting tiles from different locations on the solar disk. A collection of tiles from the same set of Dopplergrams (the same analysis period) is called a tile set or mosaic. It is most useful to have a tile set where the tile centers are positioned on a regular grid in longitude and latitude. The extent of the grid in both directions determines the size of the analysis region, and the grid spacing (separation between adjacent tile centers) influences the horizontal resolution with which one can measure flows.

As mentioned earlier, the initial horizontal resolution of measurements made with ring-diagram tiles is set by the tile size, but further mathematical analysis in the inversion step (Figure 2.1) allows one to push past this limit. The specifics of how the inversion can create high-resolution results are covered in detail in Chapter 4. The horizontal resolution in this case is limited by the spacing between adjacent tiles in the tile set and also influenced by the error statistics of the mode measurements. The overlap between adjacent tiles (however significant) is accounted for in the inversion.

While the scientific goal of measuring flows throughout the NSSL is achieved by using tiles of a large enough size (for depth information), the goal of resolving supergranular flows can be achieved by using an appropriate tiling scheme. In practice, the highest useful resolution at which the inversion is capable of producing flow fields is roughly three times coarser than the grid spacing of the tile set. Thus in order for the inversion to resolve supergranules (two data points across a typical diameter of 30 Mm), the tile set grid

spacing must be at most 0.4° (5 Mm). In this thesis, I use a grid spacing that is slightly finer than this (0.25° , 3 Mm).

Figure 2.3 demonstrates the relative sizes of a tile, an analysis region, and the grid spacing between tiles. In panel (a), a full-disk Dopplergram is shown with the size of a 90° square analysis region centered on disk center (outer rectangle) and the size of a 16° circle. While the tiles are square, the data within each tile is apodized with a circular function. Panel (b) shows an enlargement of panel (a), focusing on the 16° apodization circle and projected onto a flat plane. The grid of dots in panel (b) indicate the spacing between adjacent tile centers in for the gridding scheme used in this thesis. It is important to note that with this scheme, adjacent tiles overlap by a large fraction and therefore contain much of the same Dopplergram data. This is in contrast to the standard practice in ring-diagram helioseismology where the grid spacing is set to be half the tile size (Haber et al. 2000; Hindman et al. 2000). This type of grid is called a *dense-pack* grid and is performed in the HMI Ring-Diagram Pipeline (Bogart et al. 2011a). Since the grid employed in this thesis is much denser, it is called an *ultra-dense pack*.

Using a finer grid than absolutely necessary to measure supergranulation is done for two reasons. The first reason is to simply provide more data points across the typical size of a supergranule. While two data points per wavelength of a signal is enough to resolve it, supergranules are not simple sine waves. The higher spatial resolution will allow a more precise mapping of supergranular flows. The second reason for erring on the side of a finer tile-set grid is to reduce the effect of aliasing by small-scale flows. The photospheric spectrum of convective velocities (Hathaway et al. 2002) contains a broad range of scales and peaks at the scale of granulation ($\approx 0.08^\circ$, 1 Mm). Flows such as these that are not resolved on the tile-set grid become aliased and are no longer separable from the resolved flows. The small-scale flows are largely uncorrelated from the larger scales, so the aliased signal shows up as uncorrelated noise in the measurements. This source of noise is not explicitly accounted for in any of the formal uncertainty analyses. The consequence of under-resolving the dominant flows in the solar interior is a reduction in the effective signal-to-noise of the inferred flow fields. This is why it is advantageous to use as fine a grid of tiles as possible. Without knowledge of the true spectrum of convective flows at each depth, it is impossible to determine the exact amount of aliased power as a function of grid spacing. However, using known parts of the near-surface




Figure 2.3: (a) Example full-disk Dopplergram with lines indicating the full extent of an analysis region (outer rectangle) and the size of a single 16° tile (inner circle). (b) A blow-up of the Dopplergram and tile from (a) along with dots indicating the tile centers for a 0.25° grid spacing. The circle once again indicates the size of a 16° tile. Filling the entire analysis region with tiles spaced by this amount results in 130321 total tiles.

spectrum, I have estimated that the grid used in this thesis (0.25° spacing) introduces roughly ten times less aliased power than the standard dense-pack grid (7.5° spacing). This estimate is made by considering the observed high-resolution spectrum of flows and considering the effects of under-resolving the signal and thus exchanging resolved signal for aliased power.

2.3 Representation in Spectral Space

The Doppler shifts of wave modes are most easily measured in Fourier space. Therefore to extract this information, we compute the power of the Fourier transform of each tile in the mosaic. Despite the spherical nature of the Sun, we forego the use of a spherical harmonic transform in favor of the simplicity of the Fourier transform. The small spatial extent of the tiles compared to the radius of the Sun allows us to treat each tile as a Cartesian box. The three-dimensional, Fourier power spectrum $P(\mathbf{k}, \omega)$ created from a given tile contains the information about subsurface flows we wish to extract. The vector $\mathbf{k}$ indicates the two-dimensional wavenumber (k_x, k_y) , and ω is the frequency. As seen in the previous chapter, cuts at constant frequency of such a power spectrum reveal the modal structure of solar acoustic waves (Figure 1.7). Analyzing the properties of these modes will provide measurements of the subsurface flows.

2.3.1 Flow-Induced Perturbation to the Wavefield

The first step in building up the linear relationship between ring-diagram results and the subsurface flow field is to consider the effect a subsurface flow has on the observations of the solar wavefield captured in a tile. The observed waves permeate the solar interior and have predictable frequencies and line widths in the absence of subsurface flows. Assuming the waves themselves are not a significant contributor to the large-scale dynamics of the solar interior, one can consider linear perturbations to these waves due to the presence of a subsurface flow.

The full calculation to derive the perturbation of the observed wavefield in the presence of a subsurface flow is carried out in Birch et al. (2004). In this, acoustic-gravity waves traveling through an adiabatic background state in hydrostatic equilibrium are considered. With the assumption that the region in question is small compared to the radius of the Sun, sphericity is ignored and solutions are cast in Cartesian

coordinates. As mentioned in the previous section, we have chosen to analyze tiles that are small enough to justify this assumption.

The primary ingredients of this calculation are a description of the background stratification around which to linearize the momentum and energy equations, a description of the wave modes being considered, and a model of wave excitation. The background stratification is taken from model S (Christensen-Dalsgaard et al. 1996), which provides the density, sound speed, and pressure as a function of depth. The frequencies of the wave modes at each wavenumber and radial order are determined by these profiles, which are able to reproduce the ridge structure seen in solar oscillation measurements. The damping rates of these waves (and thus their line-widths in spectral space) have been modeled separately and also tuned to match observations (Birch et al. 2004). The source of wave energy is modeled as stochastic excitation due to near-surface convection (i.e., granulation). Each source is taken as a vertical momentum impulse at a depth of 100 km below the photosphere. The parameters of this source model were also tuned to match observations of photospheric oscillations (Birch et al. 2004). The Green's functions of the system are derived through perturbation analysis and allow one to express the response of the observed wavefield to the source term.

The end result of this calculation is an expression of the modification of the observed wavefield caused by subsurface flows. This is the first step in defining the relationship in Equation 2.1 that will linearly relate the measured subsurface flow field to the true subsurface flow field. A key consequence of this calculation is the fact that the observed wavefield within a single tile is predominantly affected by the subsurface flows directly below the tile and contained within the horizontal reaches of the apodization circle. This is to say that the sensitivity of the observed perturbation is localized to the apodization disk of the tile.

2.3.2 Perturbations to the Power Spectrum

Just as with the real-space representation of the photospheric wavefield observations, the three-dimensional power spectrum can be related linearly to the subsurface flow field. This step of the calculation can be found in Birch et al. (2007), which applies the preliminary results of Birch et al. (2004) to the specific case of ring-diagram analysis. In this, we find that the power spectrum can be related linearly to the

subsurface flow field as such:

$$\delta P(\mathbf{k}, \omega) = \int \int d\mathbf{x}' dz' \mathbf{G}(\mathbf{k}, \omega, \mathbf{x}', z') \cdot \mathbf{v}(\mathbf{x}', z') \quad (2.2)$$

Here, δP is the perturbation to the power spectrum relative to a model spectrum $P_0(\mathbf{k}, \omega)$ at a given horizontal wavenumber vector $\mathbf{k}$ and frequency ω . The Green's function $\mathbf{G}(\mathbf{k}, \omega, \mathbf{x}', z')$ specifies how the perturbed power at a given $(\mathbf{k}, \omega)$ is influenced by each component of the subsurface velocity at any point in space $(\mathbf{x}', z')$. The form of $\mathbf{G}(\mathbf{k}, \omega, \mathbf{x}', z')$ is derived in Birch et al. (2007) and once again relies on the assumptions and models used in the previous derivation. This relation between the perturbed power and the subsurface flow field is similar to the expression for the final ring-diagram result (Equation 2.1). The following steps of ring-diagram analysis that result in measurements of the subsurface flow can be seen as a series of manipulations to $\mathbf{G}(\mathbf{k}, \omega, \mathbf{x}', z')$ that result in a desired form of the averaging kernel $Q(\mathbf{x}, \mathbf{x}', z, z')$.

By comparing a model of the unperturbed power $P_0(k, \omega)$ to a power spectrum created from real data, one could measure the perturbation to the power spectrum δP at a given $(\mathbf{k}, \omega)$ fairly easily. Then, given the function $\mathbf{G}(\mathbf{k}, \omega, \mathbf{x}', z')$, one could use the measured power perturbation $\delta P(\mathbf{k}, \omega)$ to infer the subsurface velocity by inverting Equation 2.2. While this would be a straightforward way of measuring subsurface flows, it is not standard practice. The variations in the power spectrum compared to a mean power spectrum are not dominated by flow-induced perturbations, but instead by realization noise, i.e., the random fluctuations in the observed wavefield produced by the stochastic wave source.

2.4 Measuring Shifts in Spectral Power

A simpler and more intuitive description of the influence of subsurface flows on the surface wavefield is that the waves experience a Doppler shift as they pass through regions of flow. The modal structure of the wavefield can be seen clearly in the power spectrum (Figure 2.4), indicating that it is more useful to consider the distinct wave modes of the solar interior as opposed to the Fourier components of the surface Dopplergrams. Wave propagation for a single mode is isotropic, creating the uniform rings of spectral power seen in Figure 1.7. These waves travel through the solar interior and experience reflection just below the photosphere, but tunnel sufficiently into the atmosphere where they are observed. The relative angle between

the subsurface velocity through which they travel and the wave propagation vector determines how much a particular wave will be Doppler shifted when observed at the surface. Waves travelling parallel to the flow will be effected to the greatest extent, while waves travelling perpendicular to the flow will remain unaffected. It is then more intuitive to measure the Doppler shift of a distinct wave mode of the Sun and relate that measurement to a subsurface velocity.

The typical way in which this Doppler shift is measured from a power spectrum is with a non-linear fitting code. A multi-parameter function that mimics the shape of the modal power and allows for a directional Doppler shift is fit to different modes. The Doppler shift of the mode is parameterized in such a way that the two horizontal components of an inferred velocity can be extracted from the fit. In the following chapter, I will discuss non-linear ring-fitting codes in more detail.

Unfortunately, passing the power spectra through a non-linear process like this severely limits how much further we can follow the mathematical analysis presented in this chapter. To get around this, we use a linear approximation to the non-linear process that has been found by Birch et al. (2007) to closely match the results obtained with a fitting code. The measured Doppler shift is taken as a weighted integral over the perturbed power at a constant wavenumber:

$$\mathbf{u}_i = \int_{\omega_o - \Delta\omega_i}^{\omega_o + \Delta\omega_i} \int_0^{2\pi} W_i(\omega) \delta P(k_i, \theta, \omega) \hat{\mathbf{k}} d\theta d\omega, \quad (2.3)$$

where $\hat{\mathbf{k}}$ is the unit vector pointing in the direction of the wavevector, i.e., in the θ direction. Here, the weighting function $\mathbf{W}_i$ is a function of wavenumber direction θ and frequency ω and is applied to the perturbed power at a constant wavenumber $k_i = |\mathbf{k}|$ appropriate for the wave mode. The angle θ is the angle that the wavevector $\mathbf{k}$ makes with respect to the k_x -axis. The index i specifies the distinct mode within the power spectrum that is being measured and the frequency window specified by $\Delta\omega_i$ is chosen to isolate the mode in question. The weighting function $\mathbf{W}_i$ is defined as

$$W_i \sim \frac{1}{P_0(k_i, \omega)^2} \frac{\partial P_0(k_i, \omega)}{\partial \omega}. \quad (2.4)$$

In order to apply this linear measurement process to a power spectrum, the spectrum must be transformed from Cartesian coordinates (k_x, k_y, ω) to polar coordinates (k, θ, ω) . The linear measurement procedure

acts on cuts of the polar-coordinate power spectrum at constant wavenumber k . This transformation is applied to each power spectrum as soon as it is generated. The number of pixels in the k direction is chosen to match the number of distinct wavenumber magnitudes along a cardinal direction. For a 16° tile using HMI data, this results in 192 k -bins. The resolution in θ is chosen to be a constant for all values of k for simplicity. My choice of 64 pixels in θ results in low wavenumbers ($k < 10$) being oversampled and high wavenumbers ($k > 10$) being undersampled. Higher θ resolution results in both higher precision measurements of flows and higher computational cost. For the primary tile size considered in this thesis, the original Cartesian dimensions are ($n_{k_x} = 384$, $n_{k_y} = 384$, $n_\omega = 1024$) and the final polar dimensions are ($n_k = 192$, $n_\theta = 64$, $n_\omega = 1024$). The polar-coordinate power spectrum is often called the ‘unwrapped’ spectrum. The modal structure of the power spectrum is obvious when displayed in an azimuthal average (Figure 2.4).

As opposed to using single-pixel measures of perturbed power as measurements of the subsurface velocity, this method of measuring frequency shifts of modal power is conceptually straightforward. The non-linear fitting codes are provided a model of how the spectral power from a given mode is expected to change as a function of subsurface velocity. In this way, the fitting code can determine the value of the subsurface velocity that best matches the power spectrum. How the linear power-spectrum weighting mimics this kind of measurement is less clear. Figure 2.5a shows a cut at constant wavenumber of a 16° power spectrum showing a single wave mode that exhibits a frequency shift due to a subsurface velocity. The direction of the frequency shift relative to the azimuthal direction within the power spectrum indicates a flow in the positive x -direction. Figure 2.5b shows an unperturbed model of spectral power for a single wave mode. Figure 2.5c shows the weighting function W_i for measuring the x -component of the frequency shift shown in Figure 2.5a. A single scalar measurement of the x -component of the frequency shift is produced by multiplying panels (a) and (c) in Figure 2.5 together at each pixel and then summing over all pixels. The weighting function looks for power that is shifted either above or below the central frequency of the mode depending on the direction θ . The weighting function for the y -component of the frequency shift is similar to Figure 2.5c, but shifted in θ by $\pi/2$.

Producing frequency shift measurements with a linear weighting function applied to the power spectrum is a simple and robust way of measuring subsurface flows. However, this method has its drawbacks




Figure 2.4: Azimuthal average of 16° power spectrum. The full three-dimensional Cartesian power spectrum has been interpolated to polar coordinates (k, θ, ν) and averaged along the azimuthal direction θ . Darker colors indicate higher power on a logarithmic scale. The distinct ridges of power are the different wave modes. The lowest frequency ridge is the surface gravity wave mode (f -mode). The ridges above this are the pressure-wave modes, with p_1 directly above the f -mode and p_2, p_3 , etc. continuing higher in frequency. The ridges at very low wavenumber are spaced close enough to be indistinguishable given the tile size and apodization.




Figure 2.5: (a) Spectral power for a wave mode at constant wavenumber ($k = 0.8 \text{ Mm}^{-1}$) mapped as a function of azimuthal direction θ and frequency ω . The mode shows a significant frequency shift as a function of direction due to a subsurface flow in the zonal ($\theta = 0$) direction. (b) Model of unperturbed spectral power for the mode shown in panel (a). In (a) and (b), darker colors indicate higher spectral power. (c) Linear weighting function $W_i(\theta, \omega)$ for the model in panel (b). Blue indicates positive weighting and red indicated negative weighting.

when dealing with real solar data. While the linear method is a good approximation for how the non-linear fitting codes measure frequency shifts, it cannot account for variations in the power spectrum due to non-flow related effects. The reliability of the linear measurement procedure relies on the applicability of the assumed model to the real data in all situations.

The model of spectral power has been found to be reasonably accurate under ideal situations (Birch et al. 2007), but there are a few effects seen in real data that can reduce the reliability. One such effect is that wave modes seen in the power spectrum can have significant overlap with neighboring modes. At high horizontal phase speeds (ω/k), the modes become closely spaced compared to their observed line widths and eventually appear as a single, inseparable mass of spectral power (see Figure 2.4). These high phase speed modes travel deepest into the Sun and thus contain information about deep flows. A linear weighting function is unable to separate the contributions from individual modes in these circumstances, limiting the depth to which one can reliably measure flows. Another effect that complicates the use of a linear weighting function is foreshortening of the analysis region as the tile center approaches the limb of the solar disk. A tile extracted close to the limb has an effective spatial resolution that is anisotropic. In the direction towards or away from the limb, the resolution is degraded due to foreshortening. In the perpendicular direction, the resolution remains roughly the same. This effect is seen in the power spectrum as an anisotropic diminishing of power at higher wavenumbers. While the power attenuation is symmetric about the point $(k_x = 0, k_y = 0)$, it breaks the assumption of the weighting function that the unperturbed (zero velocity) modal power is constant as a function of azimuthal direction θ .

A remedy for these detrimental effects is to use a model of spectral power that can directly match these qualities of the data. A non-linear fitting code with the appropriate model is capable of parameterizing these effects and reducing their influence on the frequency-shift measurement. In the following chapter, I will discuss the development of my own fitting code and demonstrate how it can account for these adverse effects.

2.4.1 Sensitivity Kernels

Implicit in the goal of measuring subsurface solar convective flows is the ability to isolate estimates of the subsurface velocity to localized positions in space. We have found that each frequency-shift measurement encodes information about the subsurface flows, but we have yet to determine what the spatial distributions of the flows are. Recovering spatially-isolated estimates of the subsurface flow requires precise knowledge about how the direct measurements (frequency shifts) are connected to the true flow field. The size of the tile used to create the power spectra implies a horizontal region that a particular frequency shift is sensitive to, but the vertical extent has yet to be determined. To determine the exact three-dimensional volume over which a particular frequency-shift measurement is sensitive, we compute a sensitivity kernel.

Equation 2.2 shows that a measurement of perturbed power is related to an integral over the subsurface flow field. The weighting of this integral specifies the volume of the solar interior to which the measured power is sensitive. In this sense, the function $\mathbf{G}(\mathbf{k}, \omega, \mathbf{x}', z')$ is the sensitivity kernel for a measurement of $\delta P(\mathbf{k}, \omega)$. If one were to instead use the linear weighting function to measure frequency shifts (Equation 2.3), a new sensitivity kernel could be derived that includes contributions from each δP included in the weighting. As mentioned in the previous section, a non-linear measurement procedure is preferred to the linear weighting to account for non-flow related effects seen in the data. In theory, the use of a non-linear measurement of the power spectrum might preclude the possibility of computing a linear sensitivity kernel. However, it has been estimated in Birch et al. (2007) that the linear weighting function (Equation 2.4) achieves a reasonable approximation of a fully non-linear fitting method. Thus, for the purposes of computing sensitivity kernels, I have assumed that the linear weighting function given in Equation 2.4 is sufficient.

The Green's function to the perturbed power can be assumed to be isotropic and translationally invariant (Birch et al. 2004). Combining Equations 2.2 and 2.3 and considering only a single horizontal flow component, we find that the linear frequency-shift measurement is related to the true subsurface flow as

$$u_i(\mathbf{x}) = \int \int d\mathbf{x}' dz' K_i(\mathbf{x}' - \mathbf{x}, z') v(\mathbf{x}', z'), \quad (2.5)$$

where the new sensitivity kernel is created from a weighted combination of the Green's functions:

$$K_i(\mathbf{x}' - \mathbf{x}, z') = \int_{\omega_0 - \Delta\omega_i}^{\omega_0 + \Delta\omega_i} \int_0^{2\pi} W_i(\omega) \hat{\mathbf{k}} \cdot \mathbf{G}(k_i, \theta, \omega, \mathbf{x}', z') d\theta d\omega \quad (2.6)$$

This relation holds for a single directional component of the measured shift; the perpendicular direction is computed similarly. A cross-term kernel can also be constructed by considering how the linear filter may be sensitive to flows perpendicular to the intended direction. As noted in Birch et al. (2007), the overall amplitudes of the cross-term kernels are generally small compared to the single-term kernels and thus they can be safely ignored.

Figure 2.6 shows depth profiles and horizontal cuts for a selection of these three-dimensional sensitivity kernels. These generally exhibit greatest sensitivity near the surface and oscillate as a function of depth. In depth, the number of peaks of sensitivity corresponds to one plus the radial order of the mode n . The horizontal cuts shown in the right panels of Figure 2.6 are taken from the depth of the final peak in sensitivity for each kernel. The horizontal planform roughly matches the apodization circle (dashed circle), but varies somewhat with the wavenumber of the mode. Kernels from high wavenumber modes have relatively sharp edges that match the apodization circle, while kernels from low wavenumber modes show a smoother spatial profile.

The broad horizontal extent and complex vertical structure of these sensitivity kernels severely limits our ability to directly interpret the frequency-shift measurements as measures of the subsurface flow field. Chapter 4 will discuss a method of combining large sets of frequency-shift measurements and their associated sensitivity kernels to produce estimates of the flow field that are much easier to interpret.

For a more in-depth discussion on what these sensitivity kernels look like, see Featherstone (2011). The set of sensitivity kernels used in this thesis are largely taken from Featherstone (2011), with a few additions to account for a larger mode set.

2.5 High-Resolution Data Sets and Their Properties

The scientific results presented in the latter chapters of this thesis are based on measurements from eleven different analysis periods. Each analysis period is created from a 25.6 hour sequence of HMI Dopplergrams. The Carrington rotations (Carr. Rot.), central-meridian longitudes (CM Lon.), and International Atomic Times (TAI) of the midpoint of each analysis period are listed in Table 2.1. The first seven periods (A1-A7) are taken from consecutive timespans to allow near-continuous monitoring of convective patterns




Figure 2.6: Left: Horizontally integrated sensitivity kernels plotted as a function of depth. The four panels from top to bottom are different measured wave modes. The vertical lines indicate the depth of the final peak of sensitivity in depth. Right: Cuts of each sensitivity kernel at the depth indicated by the vertical lines in the left panel. The dashed circles indicate the size of the apodization circle for a 16° tile. The higher wavenumber modes (a,b) tend to show a clearer edge near the apodization circle and often exhibit some ringing. The lower wavenumber modes (c,d) have smoother sensitivity as a function of position.

as they rotate across the solar disk. The final four periods (B1-B4) are spaced out by a few weeks from each other and from the A group. These have been chosen because they contain completely independent convective realizations from the first group of analysis periods. These are particularly useful for determining the level of systematic errors. When describing the analysis period used for a particular scientific result in later chapters, I will refer to the identifiers listed in this table.

Table 2.1: List of analysis periods used in this thesis.

Identifier	Carr. Rot.	CM Lon.	Central Time
A1	2099	195	2010.07.25_20:51:42_TAI
A2	2099	180	2010.07.27_00:04:27_TAI
A3	2099	165	2010.07.28_03:17:20_TAI
A4	2099	150	2010.07.29_06:30:17_TAI
A5	2099	135	2010.07.30_09:43:15_TAI
A6	2099	120	2010.07.31_12:56:22_TAI
A7	2099	105	2010.08.01_16:09:32_TAI
B1	2100	210	2010.08.20_23:04:50_TAI
B2	2101	270	2010.09.12_16:08:56_TAI
B3	2102	180	2010.10.16_18:32:46_TAI
B4	2103	180	2010.11.13_01:44:38_TAI

A set of 16° tiles have been extracted from each analysis period spanning 90° in both longitude and latitude with a spacing of 0.25° in-between. This results in 130,321 tiles for each region, and 1,433,531 tiles in all. Each tile has been converted into an unwrapped power spectrum and fed independently to the ring-fitting code that is detailed in the following chapter. The fitting code produces around 220 reliable frequency-shift measurements per tile, resulting in a final data set of 315 million frequency shifts. Each analysis period required around 150,000 cpu-hours to go from Dopplergrams to frequency-shift measurements, resulting in 1.7 million cpu-hours for the whole set of analysis periods. This computing was done on NASA's Discover supercomputer using a code I developed called ATLAS. ATLAS is a highly parallelized code intended for

doing high-resolution ring-diagram analysis and can run on thousands of cpus simultaneously.

The final inversion step operates simultaneously on every frequency shift measurement from a chosen analysis period. The end result of the inversion is one three-dimensional flow field realization for each of the eleven analysis periods. These flow fields are the final data product of the high-resolution ring-diagram analysis described in this chapter. The information contained in these flow fields form the basis of the scientific discoveries presented in Chapters 5 and 6. These three-dimensional flow fields are the final result of the entire high-resolution ring-diagram analysis. The scientific discoveries in Chapters 5 and 6 are based solely on these results.

Chapter 3

Multi-Ridge Fitting

Inferences of subsurface flow velocities using ring-diagram helioseismology depend on measuring the frequency shifts of oscillation modes seen in acoustic power spectra. Current methods for making these measurements utilize maximum-likelihood fitting techniques to match a model of modal power to the observed spectra (Basu et al. 1999; Haber et al. 2002). The model typically describes a single oscillation mode, and each mode is fit independently from the others. We present a new method that produces measurements with greater reliability and accuracy by fitting multiple modes simultaneously. We demonstrate how this method permits measurements of subsurface flows deeper into the Sun while providing higher uniformity in data coverage and velocity response. While the previous fitting method performs marginally better for some low-phase-speed modes, we find this new method to be particularly useful for high-phase-speed modes that penetrate most deeply below the photosphere.

This chapter is based on work previously published in Greer et al. (2014)¹ and is largely a restatement of that paper with some augmentations to be more relevant to the thesis as a whole. I was the primary author on that paper and did both the code development and data analysis presented in the paper. My co-authors provided valuable guidance as to the types of statistical tests that would be useful to a broad audience.

¹ Greer, B. J., Hindman, B. W., & Toomre, J. 2014, Multi-Ridge Fitting for Ring-Diagram Helioseismology, *Solar Physics*, 289, 2823

3.1 Introduction

Helioseismology determines the structure and dynamics of the solar interior through analysis of seismic waves observed at the surface. Ring-diagram helioseismology (Hill 1988; Basu et al. 1999; Haber et al. 2002) investigates subsurface horizontal flows by measuring the direction-dependent frequency shift of oscillation modes. The Doppler shift of the frequency of an oscillation mode due to a subsurface flow is expressed as

$$\delta\omega_n(\mathbf{k}) = \mathbf{k} \cdot \mathbf{u}_i, \quad (3.1)$$

where $\mathbf{k}$ and n are the horizontal wavenumber and radial order of the oscillation mode (indexed by i), and $\mathbf{u}_i$ is the associated spatial average of the horizontal velocity within the Sun,

$$\mathbf{u}_i = \langle \mathbf{v}(\mathbf{x}, z) \rangle = \int \int d\mathbf{x} \, dz \, K_i(\mathbf{x}, z) \, \mathbf{v}(\mathbf{x}, z). \quad (3.2)$$

Here, $\mathbf{v}(\mathbf{x}, z)$ is the true horizontal subsurface flow velocity at any point $(\mathbf{x}, z)$ in the Sun, and $K_i(\mathbf{x}, z)$ is the weighting function—or sensitivity kernel—associated with each distinct mode (i), which describes the spatial extent over which the true velocity is averaged to create a single frequency shift (Birch et al. 2007).

While the interpretation of the frequency shift is straightforward, the method of extracting it from the data is not. Traditionally, a model of the spectral power is fit to those oscillation modes visible in power spectra of the line-of-sight velocity observed in the photosphere. The model accounts for a frequency shift of the modal power as a function of horizontal direction, and this shift is directly related to the velocity as shown in Equation 3.1. The specifics of the fitting procedure determine how well the frequency shifts are measured, as well as what other qualities of the power spectra are taken into account. There are currently two commonly used fitting procedures in the Helioseismic and Magnetic Imager (HMI) Ring-Diagram Pipeline (Bogart et al. 2011a). The first method considered in this chapter is one introduced by Haber et al. (2002), which fits a frequency-shifted Lorentzian model to individual modes. Since this method analyzes single ridges of modal power sequentially, we refer to it as the Single-Ridge Fitting method (SRF). The second method used in the HMI Pipeline, which will not be considered here, also fits modes independently (Basu et al. 1999), but uses a significantly different power model that includes asymmetries in modal power, two background terms, and a host of other wavenumber- and direction-dependent modifications.

In this chapter, we present a new fitting method that utilizes a model similar to that used in the SRF method, but modified to permit multiple radial orders to be fit simultaneously. The development of this Multi-Ridge Fitting (MRF) method is an attempt to improve upon the performance of the SRF method in terms of reliability and accuracy of measured frequency shifts. We present a comparison of the frequency shifts from a common data set processed with both fitting methods. We focus on the performance for measurements of modes that reach deepest into the Sun and for measurements made near the solar limb. As metrics of reliability and accuracy we consider the fit success rate, the typical estimated uncertainty in the measurements, and how well each method recovers an introduced frequency shift from a known velocity.

Both fitting codes operate on the unwrapped three-dimensional power spectra described in the previous chapter. For better comparison to previous studies, three different tile sizes have been analyzed (16° , 4° , and 2°). While the performance of the two fitting codes for smaller tile sizes is instructive, the rest of this thesis only relies on results made with 16° tiles. For the majority of this chapter, the standard dense-pack tiling scheme has been used.

The MRF code uses these spectra as input, while the SRF code requires an additional processing step. For the SRF method, the power is normalized by the average power computed at each wavenumber k and azimuth angle θ (Haber et al. 2000). The purpose of this filtering is to eliminate large power variations as a function of azimuth which arise from a variety of sources including camera astigmatism and power foreshortening. The MRF method does not use filtered power spectra, as it allows for such variations of power in the fitting procedure itself.

In Section 3.2 we describe the SRF and MRF fitting methods and how they differ in procedure. In Section 3.3 we compare the performance of the two fitting methods using a common data set. In Section 3.4 we discuss the implications of the results in the context of improving accuracy and data coverage in ring-diagram helioseismology.




Figure 3.1: Sections at $k = 0.8 \text{ Mm}^{-1}$ of the power spectra for three different 16° tiles: (a) tile extracted from disk center tracked at the Snodgrass (1984) rate, (b) tile from disk center tracked at 500 m s^{-1} relative to the Snodgrass rate, and (c) tile from the central meridian and 75°N latitude tracked at 130 m s^{-1} . Red dashed lines indicate approximate bounds for the frequency windows used for independent mode fitting in the SRF method. Green dashed lines indicate the single frequency window used in the MRF method. Velocity-induced frequency shifts are evident in (b) and (c), and significant power foreshortening is visible in (c).

3.2 Fitting Methods

3.2.1 Single-Ridge Fitting (SRF)

In a cylindrical section at constant wavenumber $[k]$, distinct modes can clearly be seen (Figure 3.1). Each radial order intersects the surface of constant wavenumber to create a band of power with a central frequency. Sub-surface flows cause each band to trace out a slight sinusoidal undulation in frequency (ν) as a function of azimuth (θ). In order to separate each radial order seen in the data, the SRF method extracts a small range of frequencies at a constant wavenumber centered around a specific mode. The bounds in frequency are based on a guess table that provides the initial values of the parameters that will be fit to the data. When a mode is framed both above and below in frequency by other modes, the bounds in frequency

are taken as part-way between the central frequencies of the adjacent modes. For the highest and lowest modes in frequency that do not have adjacent modes on both sides, the outer boundaries are taken as a few linewidths away. SRF sequentially extracts frequency windows around each desired radial order and fits a six-parameter model—Equation 3.3—to each window using a maximum-likelihood technique (Anderson et al. 1990; Haber et al. 2002). The SRF method represents the power in each mode as a symmetric Lorentzian with an angle-dependent frequency:

$$P(\nu, \theta) = B + \frac{A\Gamma/2}{[\nu - \nu_0 + (2\pi)^{-1}\mathbf{k} \cdot \mathbf{u}]^2 + (\Gamma/2)^2}, \quad (3.3)$$

$$\mathbf{u} = u_x \cos\theta \hat{\mathbf{x}} + u_y \sin\theta \hat{\mathbf{y}}. \quad (3.4)$$

Here $P(\nu, \theta)$ is the power as a function of frequency and azimuth; A is the amplitude; Γ is the line-width; ν_0 is the central frequency; k is the wavenumber; u_x and u_y are the frequency shift components; and B is a constant background. Since the measured frequency shift is an average measure of the subsurface horizontal velocity, it has two horizontal components and can itself be written as a vector quantity. There are six parameters in this model ($A, \Gamma, \nu_0, u_x, u_y, B$), that need to be fit to each mode independently at every wavenumber and radial order. By marching through each discrete wavenumber for a given tile size and attempting to fit each mode listed in the guess table, the SRF method produces the optimal values of these six parameters for every mode. A measurement of the random error of each frequency shift component (u_x and u_y) is determined by the curvature of the maximum-likelihood function evaluated at the point of optimization (Anderson et al. 1990). This provides an estimate of how much each parameter can vary without significantly decreasing the quality of the fit. It is this that we identify as the estimated uncertainty on each measurement that we will be making. Each measured frequency shift has an associated estimated uncertainty that is measured from the same data.

There are a number of ways in which a fit can fail or be deemed invalid. The numerical optimization procedure can fail to converge on a solution, causing no valid data to be produced for a single mode. If there is a successful fit, the data can still be rejected if the parameters are outside of predetermined bounds.




Figure 3.2: Azimuthal averages of a spectrum generated from a single 16° tile extracted at disk center. Each curve corresponds to a different horizontal wavenumber: $k = 0.23 \text{ Mm}^{-1}$ (red, top), $k = 0.58 \text{ Mm}^{-1}$ (green, middle), and $k = 1.42 \text{ Mm}^{-1}$ (blue, bottom). Curves are offset in power for clarity. Modes at lower wavenumber and higher frequency are spaced closer together, causing significant overlap. Ridges at high wavenumber look flattened due to azimuthal averaging over a flow-induced frequency shift.

These bounds are chosen to ensure that the parameters obtained through fitting are physically relevant.

At moderate wavenumbers ($k \geq 0.5 \text{ Mm}^{-1}$) in larger tile sizes, individual modes are sufficiently separated in frequency to justify this approach of independent fitting (Figure 3.2). However, at lower wavenumbers, modes of neighboring radial order blend together to form a mound of spectral power without prominent individual peaks. Smaller tile sizes also have more blending of modes at all wavenumbers due to the smaller spatial apodization. In these situations, severe mode blending brings into question the idea of fitting each mode independently with a single Lorentzian model with a constant background term. The frequency window used around each mode also raises issues when a subsurface flow introduces a large frequency shift. At $k = 0.5 \text{ Mm}^{-1}$, a 300 m s^{-1} flow is sufficient to shift significant mode power past the frequency boundaries chosen to split up the data for fitting. Tiles produced in the standard HMI Ring-Diagram Pipeline are tracked at the Carrington rate (Bogart et al. 2011a), which at extreme latitudes differs from the local surface rotation rate by nearly 250 m s^{-1} . Not only does mode power shift outside the fitting window for a given mode, but power from neighboring modes that is not accounted for in the model enters the window and potentially throws off the fitting (Routh et al. 2011).

3.2.2 Multi-Ridge Fitting (MRF)

To mitigate the issues that come with independent fitting of modes, we have developed the Multi-Ridge Fitting (MRF) code. This new method fits multiple radial orders simultaneously at a given wavenumber over a wider window in frequency. The function used to describe the collection of modes is a sum of symmetric

Lorentzians, each similar to those used in the SRF method:

$$P(\nu, \theta) = B(\nu, \theta) + \sum_{n=0}^{N-1} \frac{A_n(\Gamma_n/2)}{[\nu - \nu_n + (2\pi)^{-1} \mathbf{k} \cdot \mathbf{u}_n]^2 + (\Gamma_n/2)^2} F(\theta; f_n, \theta_n), \quad (3.5)$$

$$F(\theta; f_n, \theta_n) = 1 + f_n \cos[2(\theta - \theta_n)], \quad (3.6)$$

where A_n , Γ_n , ν_n , $u_{x,n}$, and $u_{y,n}$ retain the same purpose as in the SRF method, and the sum is taken over the number of modes specified in the guess table at a given wavenumber. The new factor $[F(\theta; f_n, \theta_n)]$ in the numerator of each Lorentzian accounts for amplitude variation as a function of θ using the parameters f_n and θ_n . While the constant background term used in Equation 3.3 is roughly valid when considering the narrow range of frequencies used in SRF, the background term is modified here to represent background power over a wider range by using a modified Harvey law (Harvey 1985):

$$B(\nu, \theta) = \frac{B_0}{1 + (\nu/\nu_{\text{bg}})^b} F(\theta; f_{\text{bg}}, \theta_{\text{bg}}). \quad (3.7)$$

Here, B_0 is the amplitude; ν_{bg} is a roll-off frequency; and b is the power law index. Again, there is a power anisotropy term ($f_{\text{bg}}, \theta_{\text{bg}}$) in the amplitude. The low-frequency end of the fitting window is taken as 0.3 mHz at all wavenumbers to allow a large fraction of the background power to be utilized in constraining the model. The top of the frequency window is taken as two linewidths above the highest central frequency listed in the guess table at a given wavenumber. As a pre-conditioning step before performing the optimization of every parameter in Equation 3.5, the MRF method performs a simple three-parameter fit (A, Γ, ν_0) to each mode listed in the guess table. This step is similar to the fit performed in the SRF method, but does not take any variation along azimuth into account. This step provides an improved initial condition for the full optimization for each mode. Improved initial parameters for the background term are also obtained prior to the final optimization by fitting the parameters (B_0, ν_{bg}, b) to the spectrum between 0.3 mHz and 1.5 mHz. Uncertainty estimates are once again measured from the curvature of the maximum-likelihood function.

It is important to note the significant increase in the number of parameters from Equation 3.3 to Equation 3.5. Typical nonlinear optimization algorithms scale in time as m^2 , where m is the number of parameters being fit simultaneously. To fit all of the ridges at a single wavenumber, the SRF method performs N sequential optimizations with $m = 6$. The MRF method fits the same set of ridges simultaneously

with $m = 7N + 5$, leading to an expected optimization time $\approx 1.4N$ times slower for large N . Not only is this new procedure more computationally demanding, it is much more susceptible to the numerical issue of wandering through parameter space. To ensure a timely convergence, constraints are placed on some of the model parameters:

- Amplitudes and widths must be positive.
- Central frequencies must not cross those of adjacent modes.
- Widths must be within an order of magnitude of lookup table values.
- Both zonal and meridional frequency shifts must be smaller than $\pm 1\text{km s}^{-1}$.
- Fractional anisotropy must be between 0 and 1.

These constraints are enforced during the optimization process to limit the available parameter space.

3.3 Comparison

Each fitting method requires a set of initial values for each parameter, often referred to as a guess table. The guess table contains typical values of the central frequencies, widths, and amplitudes of each mode that are fit. To have a fair comparison between the two methods, they have been supplied the same guess table, and thus the same set of modes to attempt for each tile size (Figure 3.3).

While the MRF method accounts for power anisotropy in the model (Equation 3.5), the SRF method depends on additional processing of the input power spectra to flatten power anisotropy. The unfiltered power spectra are fit using both methods, while the filtered ones are only fit by the SRF method. We confirm that the SRF method performs slightly better on the filtered spectra (Haber, personal communication, 2013), so the analysis in this chapter compares unfiltered spectra fit with the MRF method to filtered spectra fit with the SRF method.




Figure 3.3: Guess-table frequencies for each tile size demonstrating which modes are attempted in both fitting methods. The gray dashed lines are lines of constant phase-speed (ω/k). From left to right, the values of phase-speed are 10, 20, 40, and 80 km s^{-1} . The discrete ridges of modes indicate the individual radial orders, with the f -mode at low frequency, followed by p_1 through p_9 .

3.3.1 High Phase-Speed Modes

Each distinct mode that is fit has an associated kernel that specifies the spatial sensitivity to subsurface flows (see Equation 3.2). While the horizontal profile of these kernels is largely determined by the tile apodization (Hindman et al. 2005; Birch et al. 2007), the vertical profile changes significantly across the mode set as it depends on the mode’s radial eigenfunction. The sensitivity for f -mode kernels has a single peak in depth, and each successively higher radial order adds an additional peak. The final peak of sensitivity for any mode occurs approximately where the horizontal phase-speed $[\omega/k]$ equals the local sound speed, causing each mode within a single radial order but different horizontal wavenumber to have a slightly different kernel. By matching phase-speeds to the sound speeds tabulated in solar Model S (Christensen-Dalsgaard et al. 1996), we can use phase-speed as a proxy for the depth to which a mode reaches into the Sun.

Current results that utilize the SRF method can reach down to 20 Mm ($0.97 R_\odot$), partway into the near-surface shear layer (Haber et al. 2000). In order to extend modern ring-diagram analysis deeper into the Sun, we must obtain frequency-shift measurements for higher phase-speed modes. An easy way to do this is to increase the tile size or the tracking duration, with the result of increasing the signal-to-noise across

the entire power spectrum of each tile. Increased tile sizes also suffer less blending and can sample lower wavenumbers. However, if we wish to preserve horizontal and temporal resolution, we must consider how improvements to the fitting method can produce reliable measurements for higher phase-speed modes of a constant tile size.

The first aspect to consider is how reliable each method is for obtaining measurements of high phase-speed modes. By dividing the number of successful measurements by the number attempted, we get a measure for each mode of the success rate on a per-mode basis. There are two primary reasons that either fitting method will fail to produce a successful measurement: poor data quality and poor choice of model. In the first case, as the signal-to-noise ratio of a particular mode decreases, the estimated uncertainty on each model parameter increases. At some point, the fitting methods will fail to locate the mode amongst the noise and the fit will be deemed a failure. In the second case, if there is no way for the fitting method to adjust the provided model in a way that appropriately matches the data, it will struggle to converge on an optimal and unique solution. Both fitting methods will judge the attempt a failure if it does not converge within a specified number of optimization iterations or converges to non-physical model parameters. Measurements of the uncertainty on the model parameters for this case are not necessarily an accurate estimation due to the irreparable disparity between model and data.

Plotted as a function of lower turning point depth, Figure 3.4 demonstrates how deep into the Sun each fitting method is capable of measuring. To isolate the effects of phase-speed, the success rate illustrated here was compiled using only tiles bounded within 30° of disk center. Since the spacing between adjacent tiles in our mosaic depends on the tile size considered, there is a different number of tiles within this boundary for each tile size (49 for 16° tiles, 805 for 4° , and 3249 for 2°). We average the success rate over ten independent realizations of all tiles found within this boundary, resulting in 490 attempted frequency shift measurements of each mode for 16° tiles, 8050 for 4° tiles, and 32,490 for 2° tiles. While the guess table has been tuned to provide optimal results for the SRF method, the scatter seen in the SRF success rate is primarily caused by imperfect guess parameters. The MRF method is in general less susceptible to this effect. A striking difference between the two methods is how the success rate falls off with increasing depth. Both methods are most successful for the shallowest depths sampled with a given tile size. The SRF method tends to have

a gradual decrease in the success rate as the depth increases, while the MRF method maintains a nearly constant success rate through the entire mode set. As the tile size decreases, the MRF success rate again stays largely constant while the SRF success rate decreases for all depths. The success rate for the SRF method also has a slight dependence on wavenumber, causing the variation of success rate between each radial order at every depth.

The measured uncertainty as determined by each fitting method also show a strong dependence on phase-speed (Figure 3.5). As a function of lower turning-point depth for each mode, the average estimated uncertainty is bounded from below by an envelope that exponentially increases with depth. Near-surface flows in the upper few megameters of the Sun (typical flow speed $\approx 150 \text{ m s}^{-1}$) can be determined with sufficient precision using a single mosaic of tiles tracked through one day, while flows at the bottom of the near surface shear layer ($\approx 35 \text{ Mm}$, typical flow speed $\approx 50 \text{ m s}^{-1}$; Schou et al. (1998)) require a significant amount of averaging in time and space to achieve a reasonable signal-to-noise ratio. Despite the considerable differences in success rate at extreme depths, there are only slight differences between the typical measured uncertainties of each method. The average uncertainty obtained with the MRF method is slightly larger for shallow modes than that obtained with the SRF method. For deeper penetrating modes and for smaller tile sizes, the MRF uncertainties become smaller than the SRF uncertainties.

Through these two metrics (success rate, estimated uncertainty), we see that the MRF method performs more consistently with depth while maintaining similar estimates of the average uncertainty.

3.3.2 Approaching the Solar Limb

The previous results regarding how each fitting method performs at high phase-speeds were compiled using tiles near the center of the solar disk ($r \leq 30^\circ$). Both methods perform optimally near disk center and produce lower-quality results closer to the limb. The primary cause of this is foreshortening, which causes the effective resolution in one direction to be considerably less than that in the perpendicular direction. In the power spectra, this appears as a reduction of power along the direction of foreshortening. Since the MRF method has additional terms to account for this type of power variation in both the mode power and background power, it is expected that it will outperform the SRF method near the limb.




Figure 3.4: Fit success rate as a function of lower turning point of each mode found in the guess table. Left column is SRF results. Right column is MRF results. Rows correspond to the three tile sizes. To minimize the effects of disk position, all tiles of a given size falling within 30° of disk center have been considered here. The success rate is averaged over ten independent realizations of these tiles, covering ten days of tracking. The MRF success rate is higher at nearly every depth for every tile size.




Figure 3.5: Left: Average estimated uncertainty of the SRF method as a function of lower turning point for each mode. Right: Ratio of MRF uncertainties to SRF uncertainties. Rows correspond to the three tile sizes. To minimize the effects of disk position, all tiles of a given size falling within 30° of disk center have been considered here. The success rate is averaged over ten independent realizations of these tiles, covering ten days of tracking. The MRF method produces uncertainty estimates that are largely similar to those from the SRF method, with some differences for modes that reach deepest.

The ability to push closer to the limb has many benefits. Regions of strong magnetic activity are continually evolving as they pass across the solar disk from the east limb to the west limb. To maximize the amount of time that these regions are accessible with ring-diagram analysis, we must be able to measure frequency shifts close to the solar limb. The consistency of measurements as a function of disk position is important for studying the evolution of flows across the disk. The possibility of high-latitude meridional counter-cells has been of great interest (Rightmire-Upton et al. 2012; Komm et al. 2013; González Hernández et al. 2010), increasing the need for analysis methods that perform consistently between disk center and high latitudes. Careful removal of large-scale systematics has a noticeable effect on the determination of subsurface flows (Zhao et al. 2012; Greer et al. 2013; Zhao et al. 2013). The accurate measurement of these systematics is also dependent on the spatial uniformity of the methods that obtain frequency shift measurements.

Once again, it is useful to first consider the success rate of each method as a function of disk position (Figure 3.6). A single mosaic of tiles tracked through one day covering the entire solar disk was analyzed to produce 223,750 attempted measurements for 16° tiles, 611,390 for 4° tiles, and 632,740 for 2° tiles. For the SRF method, not only does the success rate at every disk position depend on the radial order of a mode, but the success rate for each radial order also falls off at different distances from disk center. Each radial order has a distance from disk center where the success rate transitions from a nearly constant value near disk center (0°) to a steep drop towards the limb. For 16° -tile f -mode fits made with the SRF method, this distance is around 75° in heliographic angle, but for each higher order the distance decreases. By p_9 , this distance has dropped to only 45° . With fewer high phase-speed measurements being made away from disk center, the depth to which subsurface flows can be determined becomes position dependent. The smaller tile sizes display a similar trend, only with a lower success rate for each order at all disk positions.

In contrast, the MRF method has a higher success rate in nearly all cases. For 16° tiles, the distances at which each radial order begins to drop off are now nearly equal at around 75° . This results in the MRF method producing much more uniform results across the solar disk, both in spatial coverage and in depth. Smaller tile sizes show a decrease in the success rate closer to disk center, but it is still roughly constant for all modes.

The average uncertainty as a function of distance from disk center (Figure 3.7) shows similar trends




Figure 3.6: Success rate for each radial order averaged over all wavenumbers as a function of distance from disk center. The success rate as a function of distance from disk center for the SRF method is different for each radial order, while the success rate for the MRF method is largely constant. The f -mode has roughly the highest success rate, while each higher radial order has a successively lower success rate.

for both fitting methods. Both have nearly constant estimated uncertainties from disk center out to around 60° . Beyond this, the magnitude of the estimated uncertainties for all radial orders rise steeply in both fitting methods. The primary difference between the two methods is how severe this rise in uncertainty is near the limb. For 16° tiles fit with the SRF method, the average uncertainty for high-radial-order modes quickly passes 1 km s^{-1} outside of 60° from disk center. The MRF method shows a slower rise in estimated uncertainty for these modes, reaching 1 km s^{-1} only for the worst cases.

3.3.3 Frequency Shift Accuracy

As demonstrated in the previous sections, the MRF method produces a larger quantity of successful frequency-shift measurements with estimated uncertainties comparable to those produced with the SRF method. However, this does not guarantee that the flows determined by either method will be correct. To ascertain the accuracy of each fitting method, we must compare the measurements to a known average velocity. We accomplish this by introducing known frequency shifts into the data set through the tile tracking rate. Usually every tile is tracked longitudinally at a rate close to the surface differential rotation rate in order to minimize the effects of a large zonal flow. By altering the tracking rate for each tile such that it slides east or west at some Δv relative to the standard tracking rate, we introduce a known frequency shift as a Galilean transformation. We track the same set of tiles from the same temporal sequence at a variety of Δv values and use both fitting procedures to gather frequency-shift measurements. By comparing the difference in measured zonal frequency shift between a region tracked at $\Delta v = 0$ and the same region tracked at a non-zero Δv , we determine how accurately each fitting method can reproduce the known frequency shift. By considering only the differences in measured frequency shifts made at the same disk position, any systematics that are independent of tracking rate are removed from the analysis.

A set of tiles covering the entire solar disk are tracked for a single day over a range of Δv spanning 500 m s^{-1} in each direction. The zonal frequency shift measurements made at $\Delta v = 0$ are subtracted from those made at all other values of Δv independently for each mode of each tile to remove physical flows and systematics. The ideal velocity response is to have the measured frequency shift difference consistent with the introduced velocity difference. By subtracting the known velocity difference from the measured




Figure 3.7: Average estimated uncertainty as a function of distance from disk center for SRF on left, MRF on right. The average uncertainty for both methods stay roughly constant near disk center and increase dramatically outside 60° . The f -mode has the lowest average uncertainty, while each higher radial order has a successively higher average uncertainty.

difference, we obtain the frequency shift deviation as a function of Δv . At each disk position, we perform a linear regression of these deviations against Δv for a collection of modes of similar phase-speed,

$$[u(\Delta v) - u_0] - \Delta v = a_0 + a_1 \Delta v, \quad (3.8)$$

$$u_0 \equiv u(0). \quad (3.9)$$

Here, $u(\Delta v)$ is the measured frequency shift of a single mode obtained with either fitting procedure for a given tile tracked at Δv relative to the default rate, and u_0 is the measured frequency shift from the same tile tracked at the default rate. The left side of Equation 3.8 represents the deviation of measured frequency shift from the introduced velocity, while the right side expresses this deviation with a linear dependence on the introduced velocity. We bin together measurements made from modes within a range of phase-speeds at each disk position and fit the coefficients a_0 and a_1 . We use the coefficient a_1 as a diagnostic for how each method responds to velocities where $a_1 = 0$ is ideal. Mapped as a function of disk position (Figures 3.8, 3.9), this quantity shows the spatial uniformity of velocity response. The measured value of a_0 is consistent with zero in all cases, so it will not be discussed further. The mode set for each tile size is divided into two subsets based on the lower turning-point depth of each mode. The depth at which each mode set is split is dependent on tile size and is chosen to get roughly the same number of modes in each subset. Modes with lower turning points in the upper 16 Mm for 16° tiles, 4 Mm for 4° tiles, and 2 Mm for 2° tiles are shown in Figure 3.8, while modes with lower turning-point depths below this are shown in Figure 3.9.

The velocity response of modes with shallow lower turning points from 16° tiles is largely ideal across the entire disk for both the SRF and MRF methods (Figure 3.8), but with a slight overestimation at high latitudes. For smaller tile sizes, the SRF method produces position-dependent results, while the MRF method begins to show a uniform underestimation of frequency shifts. For 4° tiles, the SRF method underestimates frequency shifts far from the equator by around 10% and overestimates far from the central meridian by the same amount. The SRF pattern for 2° tiles is much less clear due to scatter in the measurements and lower success rate, but there is still a trend for overestimating in the east–west direction and underestimating in the north–south direction. The MRF results exhibit a much more spatially uniform pattern while still




Figure 3.8: Velocity response for modes with lower turning points in the upper 16 Mm for 16° tiles, 4 Mm for 4° tiles, and 2 Mm for 2° tiles. The inner dashed circle is at 30° from disk center and the outer dashed line is at 75°. The SRF velocity response (left) has significantly higher spatial variability over the solar disk than the MRF method (right).

showing a consistent underestimation of 5 % to 10 %. For every tile size, the SRF results tend to be most accurate within 30° of disk center (the inner dashed circles of Figures 3.8 and 3.9) while the MRF results show a consistent underestimation that depends on tile size. From 30° out to 75° from disk center (the outer dashed lines of Figures 3.8 and 3.9), the MRF results are more accurate and spatially consistent than the SRF results. Beyond 75° for smaller tile sizes, the MRF results begin a consistent underestimation while the SRF results maintain the spatially varying pattern seen inside 75° . For modes sensitive to deeper flows (Figure 3.9), spatial uniformity is less of an issue while both methods obtain underestimations (5 % to 10 %) of the introduced velocity.

As tiles are tracked at large values of Δv , they begin to average over the subsurface flows in a more extended physical region than the same tiles tracked at $\Delta v = 0$. Since this extra tracking distance is constant for all tile sizes, it is the smallest tile sizes that are affected the most. Sub-surface flows that vary on a length scale comparable to the tile size add to the scatter seen in Figures 3.8 and 3.9. Since the zonal component of these small-scale flows averages to zero over large regions of the disk, there is no significant bias introduced to the analysis. The small-scale pattern that emerges in the 2° map for MRF measurements (and 4° to a lesser extent) is due to supergranular flows. While the SRF maps should display the same exact pattern, it is interesting to note that the variance appears to obscure the pattern.

3.3.4 Post-Processing of Frequency Shift Measurements

As seen in the overview of ring-diagram processing in Figure 2.1, the next step in the helioseismic pipeline is to combine a large set of frequency shift measurements in an inversion. This inversion creates estimates of the subsurface flow at each point in space that are easier to interpret than the raw frequency shifts. As we will find in the next chapter, the particular inversion method used in this thesis requires the data coverage and estimated uncertainties of the frequency shifts to be spatially invariant. While a range of tile sizes have been considered in this chapter, the inversion procedure uses only a single tile size (16°). The inversion requires that the tiles at every disk location produce the exact same mode set with the exact same estimates uncertainties. The regions under consideration fall within 45° of disk center, which allows us to use the mean uncertainty of any given mode as a representative value over the whole region (Figure




Figure 3.9: Velocity response for modes with lower turning points between 16 and 48 Mm for 16° tiles, between 4 and 12 Mm for 4°, and between 2 and 6 Mm for 2°. The inner dashed circle is at 30° from disk center and the outer dashed line is at 75°. Frequency shift measurements from deeper modes show less spatial variability over the solar disk for both the SRF method (left) and the MRF method (right).

3.5). Figure 3.4 shows the fit success rate for each mode of 16° tiles using the MRF method. The few modes that have a success rate lower than 0.98 are removed from the data set altogether. The remaining modes still need post-processing in order to be useful to the inversion, as no mode has a success rate of exactly 1.

Post-processing is done to fill any gaps in the data and is done on a mode-by-mode basis. Frequency shift measurements made across the solar disk for the same mode are compared to each other in order to determine the quality of the mode. Mapping the x -component of the frequency shift as a function of disk position for a few modes, we can understand how reliable the fitting code really is. Figure 3.10 shows measured frequency shifts for two modes that the MRF code can fit reliably. While previous results from 16° tiles were spaced out by half the tile size, these measurements have been immensely oversampled for the benefit of the inversion procedure. Panel (a) shows a fairly smooth map of frequency shift measurements with very little uncorrelated noise. Panel (b) has a small amount of noise, but this does not appear to be pervasive over the whole region.

The post-processing applied to each mode individually has a few steps. First, separate Gaussian profiles are fit to the distributions of frequency shifts and estimated uncertainties. These Gaussian fits provide a robust measure of the mean and standard deviation of the two distributions. The mean of the distribution for the estimated uncertainty is used later as the representative value for all measurements from that mode. Second, bad data points within the set are located and flagged. A bad data point is identified by either a point that is lacking any measurement or has a frequency shift value or estimated uncertainty more than eight standard deviations away from the mean of the respective distributions. This criterion based on the distribution selects measurements that are far outliers from the rest of the measurements. The bad points are then given new values that are linearly interpolated from the nearest-neighbor locations. This interpolation step is repeated a few times to account for small groups of neighboring bad pixels. The result is a map of frequency shift measurements that has uniform coverage and lacks significant single-pixel outliers. High resolution maps such as those in Figure 3.10 provide a plethora of data points and allow a robust measure of the distributions.

Even if a mode has a high enough success rate and can be successfully post-processed to make a spatially continuous map, it can still be deemed invalid. By checking each of the 250 or so modes by eye, we




Figure 3.10: Measured frequency shifts as a function of disk location for two wave modes. For each day of analysis, 220 maps such as these are included in the inversion. The typical estimated uncertainty on any measurement in the left and right maps are 19 m s^{-1} and 64 m s^{-1} , respectively. The data shown in panel (a) appears smooth across the entire domain, while the data shown in panel (b) shows small-scale noise. This noise appears as single-pixel deviations from the overall distribution. High-phase-speed modes such as the one in panel (b) generally exhibit more noise, and are discarded if the noise level exceeds a set value. The dashed circle indicates the size of the tile (16°). It is important to note that the maps have prominent structure on scales smaller than the tile size.

find that some modes show strange spatial jumps when mapped as a function of disk position. Figure 3.11 shows two such modes that display sharp jumps in the measured frequency shift as a function of position. The frequency shift from any given tile is an average measure of the flow field below the tile, and thus we would generally expect that measurements made on a grid with a resolution much smaller than the tile size would appear smooth. Instead, preliminary analysis of these jumps in the data has shown that they are likely due to a failure of the fitting code to measure the intended frequency shift. Since the fitting code performs a non-linear optimization, it is possible for the fit to wander away from the appropriate frequency shift and latch onto some other feature in the power spectrum. It would appear as through some small-scale variation present in the Dopplergram data causes the fitting code to pick between different distributions of

measurements.




Figure 3.11: Measured frequency shifts for two high-phase-speed modes presented similarly to Figure 3.10. The arrows highlight regions that demonstrate nonphysically sharp transitions between areas of smooth data.

It is worth noting that the frequency shift measurements for these problematic modes all lie within the estimated uncertainties. Without the extreme oversampling done to create maps such as those in Figure 3.11, there would be no reason to suspect these measurements. Since the cause of this effect is uncertain and the validity of the frequency shift measurements is in question, any modes that show these strange patterns are thrown away. This culling leaves 220 useable modes (Figure 3.12). Since most of these modes are high-phase-speed modes, this reduces the depth to which we can measure subsurface flows. This high resolution mapping has not been repeated for measurements made with the SRF code, and so it is unknown whether the SRF code produces similar results.




Figure 3.12: Comparison of the 16° tile guess table modeset (crosses) and the modeset after post-processing. The loss of high-phase-speed modes is largely due to the non-physical patterns seen in Figure 3.11.

3.4 Discussion

Extracting frequency shifts that can be interpreted as subsurface velocities is done through fitting a model of acoustic power to a three-dimensional power spectrum. By fitting a model that includes various effects seen in the data to different segments of the three-dimensional spectrum, many unique fitting methods can be created. I have presented two such fitting procedures that differ both in the model used and the selection of data considered in a single optimization. These methods use identical expressions for the velocity-induced frequency shift for each oscillation mode, yet produce significantly different results for the frequency-shift measurements themselves.

3.4.1 Improvements in Depth for Inversions

The MRF method is able to obtain a greater number of frequency shift measurements for higher phase-speed modes in a given tile size than the SRF method. The increase in the success rate is also prominent for small tile sizes, suggesting that the treatment of mode blending plays an important role. Smaller tiles exhibit more leakage of modal power across wavenumbers than larger tiles due to spatial apodization. The spreading of power increases the effective width of modes at a single wavenumber. High-phase-speed modes

are already spaced closer in wavenumber than other modes, enhancing the mode blending for their region of each spectrum. The SRF model does not accommodate any overlap of power from neighboring modes when fitting a single mode, causing a disparity between model and data that worsens at higher phase-speeds and smaller tile sizes. The MRF method, in accounting for mode overlap, returns a higher number of successful measurements for these cases that can be used for studying deep flows and high-horizontal-resolution flows.

The post-processing of the MRF results mentioned in Section 3.3.4 involves removing a subset of the modes from the measurement set. Through high-resolution analysis, the frequency shift measurements from certain modes have been found to display non-physical flow patterns. This analysis has not been repeated for SRF frequency shift measurements, so it is impossible to compare the effective useful mode sets of the two methods. A direct comparison of the depths to which each method provides measurements requires knowing the exact mode sets.

3.4.2 Spatial Variability and General Implications

Anisotropy along the azimuthal direction of the power spectrum appears to be the primary cause of fitting problems near the solar limb. The SRF method tackles this issue by operating on spectra that have been “flattened” along azimuth. This step not only amplifies the modal power in certain directions in order to eliminate anisotropy, but also amplifies the noise. By accounting for the natural variation of power along each ridge as well as in the background power, the MRF method is able to push closer to the solar limb. As seen in Figure 3.6, the mode set measureable by the SRF method is position-dependent. The higher-order modes see a drop in success rate closer to disk center than lower-order modes. The depth to which one can measure is then shallower near the limb. The MRF method, in contrast, allows for consistent determinations of subsurface flows within 75° of disk center.

Large tile sizes tend to provide the most accurate velocity measurements for both fitting methods, with a slight tendency for both to overestimate at high latitudes (Figure 3.8). Both methods exhibit spatial variability in the velocity response for smaller tile sizes, but only for low-phase-speed modes. The MRF response is radially symmetric with the least accurate velocity response occurring near disk center, while the worst regions of the SRF response appear far from disk center and have opposite signs between the north–

south and east–west direction. While the magnitude of this variability is relatively small, there are significant implications when attempting to analyze small residual flows left over from the subtraction of large-scale flows. Tiles at high latitudes tracked at the Carrington rate see an overall flow speed of nearly 200 m s^{-1} due to differential rotation. The systematic underestimation seen in the SRF method introduces an anomalous retrograde flow of around 20 m s^{-1} , similar to the magnitude of the center-to-limb velocity systematic for low-phase-speed modes (Greer et al. 2013). It is unclear what causes the systematic inaccuracy in either fitting method, although the strong dependence on disk position suggests a possible coupling between velocity measurements and power anisotropy.

As mentioned in Section 3.2.2, there is a significant increase in the computational cost when switching from the SRF method to the MRF method. While processing a day’s worth of tiles with the SRF method is typically one of the fastest steps in the HMI Pipeline, the MRF method brings the process of fitting more in line with the time needed for tracking. It is important to consider when this additional computational burden is justified. The SRF method produces high-quality results for low-phase-speed modes and for large tiles near disk center. This leaves three distinct cases where the MRF method provides improvements: measuring higher-phase-speed modes, pushing closer to the limb, and using small tiles. The implications of this new procedure on determining subsurface flows follow these three technical improvements. For a given tile size, we are able to extend our analysis deeper into the Sun while maintaining a constant horizontal and temporal resolution. Analysis to these extended depths can be performed consistently across most of the solar disk, providing uniform coverage over a larger fraction of the Sun. The increased reliability of small tile sizes permits higher horizontal resolution analysis of subsurface flows.

In this chapter, I have shown how the MRF method is capable of measuring frequency shifts from ring-diagram power spectra more reliably and accurately than the previously accepted method. The low level of spatial variability is ideal for the methods of the next chapter, which assume complete translational invariance of the frequency shift uncertainties and data coverage. For the remainder of this thesis (with the exception of Appendix A), I will use the MRF method as the sole method of obtaining frequency shift measurements from ring-diagram power spectra.

Chapter 4

Helioseismic Inversions

4.1 Introduction

The final step required to produce subsurface flow maps in ring-diagram helioseismology is the inversion. This is a linear procedure that converts the highly-convolved frequency shift measurements from the previous chapter into high-resolution flow maps that are localized to a single depth. In the context of the linear relationships detailed in Chapter 2, this final step combines the sensitivity kernels from each frequency shift measurement and determines the form of the averaging kernel function Q that relates the final inversion solution w to the true subsurface flow field v :

$$w(\mathbf{x}, z) = \int \int d\mathbf{x}' dz' Q(\mathbf{x}, \mathbf{x}', z, z') v(\mathbf{x}', z'). \quad (4.1)$$

This process relies critically on the structure of the sensitivity kernels associated with each frequency shift measurement (see Chapter 2). The name “inversion” arises from the fact that we wish to solve a set of coupled integral equations such as Equation 2.5 by creating an inverse operator. The fundamental complication with the inversion problem is that it does not produce a single solution, but instead permits a range of solutions dependent on user-defined criteria. Within the range of possible solutions, the resolution, estimated uncertainty, and localization in depth can vary substantially. The scientific results that arise from the inversion solutions depend critically on these quantities, and so it is necessary to understand the inner workings of the inversion step.

Inversions play an important role in helioseismology, as the direct observations are linked to the desired subsurface flows through integral equations. The results of global helioseismology (see Chapter 1)

have relied on two-dimensional (latitude and depth) inversions to obtain isolated measures of the internal rotation rate from surface observations (e.g., Schou et al. 1994). While the nature of local helioseismic measurements lends itself to using three-dimensional inversions (due to the additional spatial information), the high computational cost of three-dimensional inversions is often not justified. Instead, one-dimensional inversions can be done independently at different locations on the solar disk with relative ease and patched together after the fact to create a three-dimensional map (e.g., Haber et al. 2001; 2002). Compared to what a three-dimensional inversion can accomplish, this scheme is severely limited in how well it can isolate flow measurements horizontally. Few three-dimensional inversion procedures exist due to their complexity and computational requirements. In recent years, separate three-dimensional inversions designed for ring-diagram and time-distance analyses have seen use in measuring small-scale flows associated with supergranules and active regions (Featherstone 2011; Jackiewicz et al. 2012).

As shown in Chapter 2, deriving the mathematical connection between observations and subsurface flows requires knowledge of the observational process and of any subsequent steps to manipulate the data. Thus, while each flavor of helioseismology utilizes inversions, the details for how this is carried out vary. In particular, the collection of the sensitivity kernels used to perform an inversion determines how well the inversion can isolate measurements of the flow.

The purpose of addressing the inversion step in great detail in this chapter is to enable subsurface flow observations with unparalleled resolution and accuracy. A recent study from Švanda (2015) showed that the analysis of near-surface convection is highly susceptible to small problems within the inversion technique and that an improperly tuned inversion can completely change the scientific conclusions made about convection. This result emphasizes that the details of the inversion procedure for local helioseismology should not be taken lightly, as they can have significant and sometimes unintended effects on the results.

There is a fundamental trade-off between the depth to which a ring-diagram frequency shift measurement is sensitive and the horizontal resolution of the measurement. Large tiles sample deeply but average over a large horizontal area. Small tiles average over smaller horizontal regions of the Sun, but are limited in depth. The Adaptive-Resolution Ring-Diagram Inversion (ARRDI) code (Featherstone 2011) is the first three-dimensional ring-diagram inversion to be developed and is able to combine multiple tile sizes to produce

flow maps that have the highest possible resolution allowed by the trade-off at each individual depth. With this method it has been possible to measure subsurface moat flows around sunspots among other subsurface flow structures. Three-dimensional ring-diagram inversions hold the promise of allowing direct observations of convective structures throughout the solar convection zone, but are limited by the depth-resolution trade-off. In this chapter, I present a new inversion technique that utilizes highly-overlapping ring-diagram tiles to defeat this tradeoff and allow observations of subsurface flows with unprecedented resolution.

The scientific results in the following chapters rely on this inversion procedure, and thus their interpretation hinges on how the inversion acts. This chapter will serve to break down the details of the inversion into manageable pieces such that the qualities of the inverted data are understandable, and a correct interpretation is possible. This chapter also serves as a reference for the details of how the results in later chapters were produced. The inversion procedure permits a wide range of mathematically valid solutions even though only one solution may be needed. The specific choice of solution from the larger set of possibilities is important, but the reasoning behind the choice is equally meaningful. In this chapter I will demonstrate how the inversion solution may vary outside of the single solution that is chosen for scientific analysis.

In section 4.2, I will provide a mathematical description of the inverse problem at hand and discuss multiple solution methods. In section 4.3, I will present a set of one-dimensional inversions that demonstrate some of the basic features of a helioseismic depth inversion such as estimated uncertainty and vertical resolution. These inversions use measurements that all have unique kernels. In section 4.4, I will use one-dimensional inversions to demonstrate the horizontal deconvolution aspect of the full helioseismic inversion and how it can be explained in Fourier space. These inversions use measurements that have translationally invariant sensitivity kernels and rely on the significant overlap between adjacent measurements. In section 4.5, I will present the full procedure for high-resolution, three-dimensional, ring-diagram inversions and explain the specifics of how it is used to obtain the results in the following chapters.

4.2 Mathematical Background

4.2.1 Frequency Shift Measurements

The frequency shift measurement made in the previous chapter is related to the subsurface velocity through a three-dimensional integral:

$$u_i(\mathbf{X}) = \int d\mathbf{x}' \int dz' K_i(\mathbf{X}, \mathbf{x}', z') v(\mathbf{x}', z') \quad (4.2)$$

where i denotes the distinct wave mode from which the frequency shift is measured. The set of modes measured from the power spectra are ordered based on radial order and wavenumber into a one-dimensional array. I will use vector notation ($\mathbf{x}$) to denote two-dimensional horizontal vectors and z to indicate vertical position. The vector $\mathbf{X}$ indicates the position on the solar disk that the measurement has been made from (the tile center). The sensitivity kernel in this equation has been written to allow for a kernel that depends on the location of the measurement on the disk, but here on forward I will assume translational invariance such that $K_i(\mathbf{X}, \mathbf{x}', z) = K_i(\mathbf{x}' - \mathbf{X}, z)$.

The structure of the sensitivity kernel K_i can be derived through perturbation analysis of the relevant wave equation (see Chapter 2). It is important to note that Equation 4.2 only describes a single directional component of the frequency shift $\mathbf{u}_i$ and subsurface flow $\mathbf{v}(\mathbf{x}, z)$. This single component can be either the zonal component (east–west) or meridional component (north–south). The integral equations for each of the two horizontal velocity components are identical and thus can be treated as separate inversions. In this chapter, I will present the inversion procedure as if I am solving for only one component.

Equation 4.2 demonstrates that while we may express the frequency shift measurements like a velocity, they are in fact measurements of a convolved velocity field. However, the convolution being performed does not always nullify the interpretation of frequency shifts as direct velocity measurements. Take for example the sensitivity kernel for f -mode frequency shifts (Figure 2.6a). This kernel has sensitivity that peaks just below the photosphere, has a simple circular planform in the horizontal directions, and most importantly, remains positive definite at all points in space. Sub-surface flow fields that vary on scales much larger than this kernel will not be altered significantly when convolved with this kernel, and so we can interpret the relevant frequency shift measurements as (approximately) direct measurement of the flow field. This interpretation

of frequency shifts as direct velocity measurements has been used extensively in helioseismology (Hindman et al. 2006; Hirzberger et al. 2008; Hindman et al. 2009), but limited primarily to the upper megameter or so of the solar interior. The difficulty in using this interpretation is knowing *a priori* that the flow field in question has the appropriate structure to remain unperturbed by the convolution process.

Sensitivity kernels from modes that reach deeper into the Sun have more complex structure, and are not always positive at all points in space (Figure 2.6d). The convolution of a flow field with a non-compact or otherwise complicated kernel results in a non-intuitive representation of the flow that is difficult to interpret. Frequency shift measurements of subsurface flows that do not respect the structure of the sensitivity kernels require further processing before an interpretation can be made. As the goal of this processing is to produce a flow field that is inferred from the frequency-shift measurements, we can call it a velocity field and interpret it as such. The process required to disentangle the integral in Equation 4.2 and estimate this inferred velocity is the focus of this chapter. Throughout this thesis, I will denote the true subsurface velocity inside the Sun as $\mathbf{v}(\mathbf{x}, z)$ (which has three components and is defined at all points in space), the frequency shift measurements made using ring-diagram analysis as $\mathbf{u}(\mathbf{X})$ (which has two horizontal components for each measured wave mode and disk location), and the final velocity estimate derived from frequency shifts as $\mathbf{w}(\mathbf{x}, z)$ (which has two horizontal components at any given target position in space).

4.2.2 Methods of Solution

Solving Equation 4.2 for v given a set of measurements u_i and kernels K_i is a linear inverse problem. The measurements contain information about v on all scales that the kernels sample, but in order to solve the inverse problem we must first discretize the equations. Each frequency shift measurement has an associated convolution relation as in Equation 4.2, giving us a set of integral equations. Accounting for the finite set of frequency shift measurements and the discretization of the sensitivity kernels, we can rewrite Equation 4.2 as:

$$u_i(\mathbf{X}_j) = \sum_n K_i(\mathbf{x}_n - \mathbf{X}_j, z_n) v(\mathbf{x}_n, z_n) \quad (4.3)$$

where $K_i(\mathbf{x}_n - \mathbf{X}_j, z_n)$ is the value of the sensitivity kernel at each discretized point in space $(\mathbf{x}_n, z_n)$ for a measurement of mode i made at disk position $(\mathbf{X}_j)$. The set of locations in three-dimensional space have

been compacted into a single one-dimensional array that I will index with the letter n . Each value of n provides a unique position in three-dimensional space $(\mathbf{x}_n, z_n)$. The set of measurements across all modes i and disk locations $\mathbf{X}_j$ can also be indexed with a single number m , leading to a simplified notation for Equation 4.3:

$$u_m = \sum_n K_{mn} v_n \quad (4.4)$$

If we include many frequency measurements and their associated sensitivity kernels, we can create a matrix problem containing the original set of integral equations:

$$\mathbf{u} = \mathbf{K}\mathbf{v}. \quad (4.5)$$

Here we have expressed the set of sensitivity kernels that relate the true subsurface flow to the frequency measurements as a matrix $\mathbf{K}$. Each row of $\mathbf{K}$ contains the appropriate sensitivity kernel for the associated frequency measurement in $\mathbf{u}$ with three-dimensional spatial information packed in the same way as the vector $\mathbf{v}$. $\mathbf{K}$ is thus a $N_{meas} \times N_{space}$ matrix, where N_{meas} is the number of measurements and N_{space} is the number of discrete points in space being targeted for an inversion solution.

The goal of the inversion is to come up with some estimate $\mathbf{w}$ of the true velocity field $\mathbf{v}$ such that $\mathbf{w} \approx \mathbf{v}$ for some appropriate interpretation of $\approx$. Since we are performing a linear inversion, we will look for solutions in the form of

$$\mathbf{w} = \mathbf{A}\mathbf{u}. \quad (4.6)$$

Here, the set of solution points $\mathbf{w}$ are related to the set of measurements through an inversion matrix $\mathbf{A}$. Before we specify how this inversion matrix is computed, we should consider how this type of linear solution is related to the true subsurface velocity field. Each solution point in $\mathbf{w}$ can be related to the full set of frequency shift measurements through a set of linear coefficients taken from the $\mathbf{A}$ matrix:

$$w(\mathbf{x}_t, z_t) = \sum_{i,j} A_i(\mathbf{X}_j, \mathbf{x}_t, z_t) u_i(\mathbf{X}_j). \quad (4.7)$$

Here, the position $(\mathbf{x}_t, z_t)$ is the target location within the solar interior that we want w to represent. It is the goal of the inversion to find a set of inversion coefficients a_i that achieve this goal. Translational invariance dictates that only the relative and not absolute locations of the solution target and the frequency

shift measurements matter, allowing us to say

$$A_i(\mathbf{X}_j, \mathbf{x}_t, z_t) = A_i(\mathbf{x}_t - \mathbf{X}_j, z_t). \quad (4.8)$$

Combining the definition of the frequency shift measurements from Equation 4.2 with this form of the inversion solution, we can relate the inversion solution at a single point to the subsurface flow field with:

$$w(\mathbf{x}_t, z_t) = \int d\mathbf{x}' \int dz' Q(\mathbf{x}' - \mathbf{x}_t, z', z_t) v(\mathbf{x}', z'), \quad (4.9)$$

where the function $Q(\mathbf{x}' - \mathbf{x}_t, z', z_t)$ (the averaging kernel) is defined as

$$Q(\mathbf{x} - \mathbf{x}_t, z, z_t) = \sum_i \sum_j A_i(\mathbf{x}_t - \mathbf{X}_j, z_t) K_i(\mathbf{x} - \mathbf{X}_j, z), \quad (4.10)$$

or in discrete matrix form,

$$\mathbf{Q} = \mathbf{A}\mathbf{K} \quad (4.11)$$

Thus, the averaging kernels for each target position are constructed by a linear combination of sensitivity kernels, where the linear weighting is determined by the specifics of the inversion procedure.

Just as the sensitivity kernel $K_i(\mathbf{x}_n - \mathbf{X}_j, z_n)$ provides the relationship between a given frequency shift measurement and the subsurface flow field, the averaging kernel function provides the same relationship for the inversion solution. The averaging kernel allows us to characterize the region of the solar interior corresponding to a solution point. Thus, the interpretation of the inversion solution as an estimate of the subsurface flow at a single point depends critically on the structure of the concomitant averaging kernel.

To understand how a useful inversion solution in the form of Equation 4.7 comes about, it is useful to investigate a few incorrect solutions first. A naïve solution derived from Equation 4.5 might look as follows:

$$\mathbf{w} = \mathbf{K}^{-1}\mathbf{u}, \quad (4.12)$$

where all we need to do is compute the inverse kernel matrix and apply it to the measurements ($\mathbf{A} = \mathbf{K}^{-1}$ in Equation 4.6). This can only be done when the number of measurements m is equal to the number of solution points n ($\mathbf{K}$ is square), and when the kernel matrix is not singular. To avoid the issue of fine-tuning the number of measurements based on the discretization of the kernels, we can choose a more robust method.

Pre-multiplying both sides of Equation 4.5 by K^T , we obtain

$$K^T u = K^T K w \quad (4.13)$$

which can be solved for w without worrying about the relationship between n and m :

$$w = (K^T K)^{-1} K^T u. \quad (4.14)$$

This form of the solution is similar to that in Equation 4.12, but uses the pseudo-inverse of K instead of the direct inverse. The use of the pseudo-inverse $(K^T K)^{-1} K^T$ allows us to form an answer even when we have an over- or under-constrained system and the kernel matrix K is degenerate. This form of solution is commonly used in linear least-squares applications.

4.2.2.1 Regularized Least Squares

Due to strong similarities between the various sensitivity kernels going into the kernel matrix, the condition number of K is very high. A high condition number in this context implies that a small change in the measurements can yield a very large change in the inversion result. Since the frequency shift measurements are susceptible to noise, this creates a solution with greatly amplified scatter. While Equation 4.14 is a valid method of solving the inverse problem, we may wish for more direct control over the solution. To achieve this control, we define an optimization problem that results in a solution with a similar form to Equation 4.14.

The simplest way to construct the problem is to say we are finding a solution w that optimally matches the observed flow field:

$$\Psi = \sum_i \sum_j \frac{1}{\sigma_{ij}^2} \left| u_i(\mathbf{X}_j) - \sum_n K_i(\mathbf{x}_n - \mathbf{X}_j, z_n) w(\mathbf{x}_n, z_n) \right|^2. \quad (4.15)$$

In this, we seek to find values for the entire discretized solution field simultaneously. The weighting factor σ_{ij}^2 is the uncertainty in the measurement $u_i(\mathbf{X}_j)$. The inclusion of the weighting puts more emphasis on low-uncertainty measurements. Therefore, the inversion tries to match the observations only to within the observational uncertainties. This type of weighting ends up rescaling the measurement vector u and kernel matrix K , and does not further change the form of the solution, i.e. $u_i/\sigma_{ij} = \tilde{u}_i$ and $K_i/\sigma_{ij} = \tilde{K}_i$.

Minimizing Ψ with respect to each element of w is equivalent to solving Equation 4.14. Written as a minimization problem, it is easier to see that the inversion solution represents the flow that—when convolved with the sensitivity kernels—matches the observed data. To get around the issue of noise amplification due to high condition number, we add a weighted regularization term to Equation 4.15. We are free to pick any form of regularization that depends on the inverted quantity w , but a common choice in this situation is to regularize against the derivative of the solution:

$$\Psi = \sum_i \sum_j \frac{1}{\sigma_{ij}} \left| u_i(\mathbf{X}_j) - \sum_n K_i(\mathbf{x}_n - \mathbf{X}_j, z_n) w(\mathbf{x}_n, z_n) \right|^2 + \lambda \sum_n |\mathcal{D}w(\mathbf{x}_n, z_n)|^2. \quad (4.16)$$

Here, $\mathcal{D}$ is the spatial derivative operator. Minimizing Ψ in this equation with respect to each element of w produces the following matrix problem:

$$w = A\tilde{u} = (\tilde{K}^T \tilde{K} + \lambda R)^{-1} \tilde{K}^T \tilde{u}, \quad (4.17)$$

where each solution point has an associated averaging kernel given by

$$Q = A\tilde{K} = (\tilde{K}^T \tilde{K} + \lambda R)^{-1} \tilde{K}^T \tilde{K} \quad (4.18)$$

The regularization matrix R contains the derivative operator. Through tuning λ , solutions can be achieved that are much smoother than before.

In helioseismology, this method of solving the inverse problem is called Regularized Least Squares (RLS) for obvious reasons. The benefit of this method is that it results in smooth flow fields (enforced by regularization) and the minimization quantity makes intuitive sense. While smooth flow fields make for an attractive solution, the enforcement of such a criteria is not always justified. Further, the smoothness is achieved by combining measurements made across the solar disk. This effect is best seen by studying the resulting averaging kernels, which often lack compactness when using an RLS method. Spatial variations in the measured frequency shifts will not always equate to spatial variations in the solution flow field, and can instead result in averaging kernels that vary in shape across the disk. This creates a solution that is extremely difficult to interpret, as each point in the flow field can represent the flow over a different shaped and large volume of the Sun.

4.2.2.2 Optimized Local Averaging

Proper interpretation of the inversion solution requires knowledge of the relevant averaging kernels. However, this is not sufficient for having a useful solution. As opposed to the RLS inversion, where the averaging kernel is allowed to take any form at any location, we wish to formulate an inversion that permits us direct control over the averaging kernel.

The definition of the averaging kernel (Equation 4.1) tells us that the size and the shape of Q determine how the solution relates to the true subsurface flows. Thus, the solution w is only as close to the truth v as the averaging kernel Q allows. The optimal situation is then that the averaging kernel matrix Q is equal to the identity matrix I , or equivalently that the averaging kernels are all delta functions. This would allow for the easiest interpretation, where each solution point represents the subsurface flow at exactly one point in space, and no significant averaging takes place. Due to the limited number of wave modes measured and the finite spatial resolution of the observations, true delta function averaging kernels are impossible. Instead, we can create an optimization problem that attempts to match the averaging kernels to a target function T :

$$\Psi = \int d\mathbf{x} \int dz |T(\mathbf{x} - \mathbf{x}_t, z, z_t) - Q(\mathbf{x}, z)|^2. \quad (4.19)$$

The choice of $A_i(\mathbf{X}_j - \mathbf{x}_t, z_t)$ (see Equation 4.10) that minimizes Ψ in this case provides the closest match of the averaging kernel to the target function. For the case of the target function being a delta function, this averaging kernel also provides the most localized measure of the subsurface flow. This general formulation of the inverse problem is called the Optimally Localized Averages (OLA) method (Christensen-Dalsgaard et al. 1990; Schou et al. 1998). The best choice for the target function $T(\mathbf{x} - \mathbf{x}_t, z, z_t)$ is a delta function centered on the point $(\mathbf{x}_t, z_t)$, but a narrow three-dimensional Gaussian function is sometimes used for numerical reasons. The averaging kernel is a linear combination of sensitivity kernels (Equation 4.10), allowing us to minimize Ψ with respect to the linear combination coefficients. In contrast to RLS, where an entire flow field is solved for in one matrix problem, OLA seeks to find the flow at a single target location by finding the set of inversion coefficients $A_i(\mathbf{X}_j)$ that provide a useful averaging kernel. A set of inversion solution points is built up by running multiple inversions targeting different positions in the solar interior. Fortunately, the inversion coefficients for one target location can often be applied to other locations (at the same target

depth), reducing the need for many independent inversions. Further, the same set of inversion coefficients used to combine the sensitivity kernels into the final averaging kernel are used to combine the frequency-shift measurements into the final inversion velocity solution w . Thus, building up a three-dimensional inversion flow field $w(\mathbf{x}, z)$ involves running a single OLA inversion for each target depth z_t , then applying the inversion coefficients $A_i(\mathbf{X}_j - \mathbf{x}_t, z_t)$ to each target position $\mathbf{x}_t$.

As with the RLS inversion, we are free to introduce a regularization term into Equation 4.19. Since the primary goal of introducing regularization before was to limit noise amplification, we add a regularization term to Equation 4.19 proportional to the propagated uncertainty on the solution. Rewriting Equation 4.19 explicitly as a function of the inversion coefficients with regularization, we obtain

$$\Psi = \int d\mathbf{x}' \int dz \left| T(\mathbf{x}' - \mathbf{x}_t, z, z_t) - \sum_i \sum_j A_i(\mathbf{X}_j - \mathbf{x}_t, z_t) K_i(\mathbf{x}' - \mathbf{X}_j, z) \right|^2 + \lambda \mathcal{R} \quad (4.20)$$

$$\mathcal{R} = \sum_{i,l} \sum_{j,k} A_i(\mathbf{X}_j - \mathbf{x}_t) A_l(\mathbf{X}_k - \mathbf{x}_t) C_{il}(\mathbf{X}_j, \mathbf{X}_k). \quad (4.21)$$

Here, $C_{il}(\mathbf{X}_j, \mathbf{X}_k)$ is the error covariance between measurements $u_i(\mathbf{X}_j)$ and $u_l(\mathbf{X}_k)$. In the case of uncorrelated measurements, $C_{il}(\mathbf{X}_j, \mathbf{X}_k) = \delta_{il} \delta_{jk} \sigma_{ij}^2$. Minimizing Ψ with respect to the inversion coefficients $A_i(\mathbf{X}_j)$ for a given target function allows us to express the solution as yet another matrix problem:

$$\mathbf{A} = (\mathbf{K}^T \mathbf{K} + \lambda \mathbf{C})^{-1} \mathbf{K}^T \mathbf{T}, \quad (4.22)$$

$$\mathbf{Q} = \mathbf{A} \mathbf{K} = (\mathbf{K}^T \mathbf{K} + \lambda \mathbf{C})^{-1} \mathbf{K}^T \mathbf{T} \mathbf{K}, \quad (4.23)$$

where $\mathbf{a}$ is the vector containing each inversion coefficient. The vector $\mathbf{T}$ contains the discretized target function, packed using the same standard as the sensitivity kernels.

It is interesting to note the similarity between this solution and that of the RLS inversion (Equation 4.18). The benefits of using the OLA method over the RLS method are that it allows for more direct control of the averaging kernel and does not make any assumptions about the characteristics of the velocity field. The two formulations of the inversion procedure (RLS, OLA) have been compared against each other in helioseismic applications and found to produce similar results (Basu et al. 1998).

One must pay careful attention to the normalization of the inversion solution, as it has not yet been specified. It is most useful to have averaging kernels that integrate to unity, as they allow for the simplest

interpretation. In an OLA inversion, an attempt can be made to produce properly normalized averaging kernels by matching to a properly normalized target function. However, an OLA inversion does not (and often cannot) produce an exact match to the target function. There are a few ways in which averaging kernel normalization can be achieved. The first is to append a second regularization term to the OLA equation that penalizes solutions that are not properly normalized. This creates the same problem as before, where the solution will tend to be normalized, but there is no guarantee. Another way to enforce normalization would be to remove a degree of freedom from the inversion. Instead of solving for m coefficients, one could solve for $m - 1$ of the set in the inversion and then set the final coefficient such that the sum of all coefficients is unity. This will enforce normalization, but is a little cumbersome. The method I use for normalization is to perform the inversion as-is and normalize the A -coefficients by their norm after the fact. This additional step is easier to understand and allows one to see how far from proper normalization the inversion strayed as a diagnostic tool. All of the inversion results presented in this chapter and others have been normalized in such a way.

4.2.2.3 Inversion Tuning

The properties of the inversion solution and averaging kernels are functions of the mode set, data uncertainty, data availability, and sensitivity kernel structure. The regularization parameter λ acts as a trade-off parameter between optimally localized averaging kernels and low uncertainties. Free choice of λ results in an infinite number of possible inversion solutions for a given set of measurements and sensitivity kernels. Since there is no wrong choice for λ , there are in a sense no wrong inversion solutions (assuming they have been computed properly). Changing λ changes both the inversion flow field and the relevant averaging kernels in such a way that they always provide a mutually-consistent answer. In this sense, the inversion solution is only as useful as the data provided to it and the averaging kernels that come out of it. The previous chapter focused on making sure the frequency-shift measurements are accurate, and this chapter focuses on making sure the inversion provides a useful interpretation.

In the following sections of this chapter, I will demonstrate a few of the most important qualities of helioseismic OLA inversions. The full three-dimensional problem has many subtleties, so I will first present

one-dimensional inversions that exhibit these qualities in idealized environments.

The three-dimensional inversion I use in this thesis to transform frequency measurements into subsurface flow fields is fairly complex. In depth, the full three-dimensional inversion uses measurements with substantially different error statistics and sensitivity kernels that span different volumes of space. In the two horizontal directions, the three-dimensional inversion uses measurements with identical error statistics and translationally invariant sensitivity kernels. The full inversion has to include all three directions simultaneously, but we can build up an intuition for how it acts on the data by considering one-dimensional inversions. In the following sections of this chapter, I will present one-dimensional inversions that mimic how the full three-dimensional inversion acts either in depth or horizontally. By understanding the basics of these idealized cases, the results of the full helioseismic inversion are easier to interpret.

4.3 One-Dimensional Depth Inversions

4.3.1 Two Box Kernels

First, let's consider an inversion where we have only made two uncorrelated measurements. We'll use the same language as for the full three-dimensional ring-diagram inversion. The two frequency splitting measurements are u_1 and u_2 , with measured uncertainties σ_1 and σ_2 , and their respective sensitivity kernels are $K_1(z)$ and $K_2(z)$. For simplicity, we pick boxcar functions for the two kernels that start at the surface $z = 0$ and end at different depths (z_1 or z_2 with $z_2 > z_1$). This mimics how the different wave modes used in the real inversion have acoustic cavities that reach from the photosphere down to a depth where the local sound speed equals the phase speed of the mode. The normalization of each kernel is picked such that the integral over the domain is unity (Figure 4.1). These amplitudes are $1/z_1$ and $1/z_2$.

For a target function, we pick a delta function located at z_0 which lies between z_1 and z_2 . It is not always necessary to use a delta function, as the main results that follow only depend on a sufficiently narrow and compact target function such that the integral of either kernel times the target is either zero or a constant, depending on the mode. Writing out the quantity to minimize in an OLA inversion (Equation




Figure 4.1: One-dimensional boxcar sensitivity kernels for mode 1 (red) and mode 2 (blue). The dashed line indicates the location of the target function.

4.20), we find it has a simple form:

$$\Psi = \int (\delta(z - z_0) - A_1 K_1(z) - A_2 K_2(z))^2 dz + \lambda [(\sigma_1 A_1)^2 + (\sigma_2 A_2)^2], \quad (4.24)$$

where the coefficients A_1 and A_2 are the quantities being solved for. Minimizing Ψ with respect to A_1 and A_2 independently leads to the solution:

$$\begin{aligned} A_1 &= -\frac{z_1/z_2}{(\lambda\sigma_1^2 z_1 + 1)(\lambda\sigma_2^2 z_2 + 1) - z_1/z_2}, \\ A_2 &= \frac{\lambda\sigma_1^2 z_1 + 1}{(\lambda\sigma_1^2 z_1 + 1)(\lambda\sigma_2^2 z_2 + 1) - z_1/z_2}. \end{aligned} \quad (4.25)$$

The solution velocity w , uncertainty estimate σ_w , and averaging kernel Q are constructed as in the previous section. As mentioned in the previous section, the inversion coefficients have been normalized after the inversion is computed.

Picking values for the kernel depths and uncertainties ($z_1 = 0.5$, $z_2 = 1$, $\sigma_1 = 1$, $\sigma_2 = 2$), we can look at how the coefficients depend on regularization. The relative values for σ_1 and σ_2 are chosen to mimic how the measured uncertainties on frequency shift measurements vary with the phase speed of the mode (Figure 3.5). Since the uncertainties only appear next to λ in Equations 4.25a and b, their ratio influences the behavior of the solution and their overall magnitude just scales the relevant values of regularization. Figure 4.2a shows the two coefficients A_1 and A_2 as a function of regularization. For small values of regularization, the coefficients converge to values that produce an averaging kernel that optimally matches the target function

(given the constraints of the kernels). The shallow mode K_1 is used to subtract off sensitivity from the deeper mode in order to get an isolated measure of the velocity around the target depth. This results in a boxcar averaging kernel spanning $z \in (z_1, z_2)$. High values of regularization lead to just one mode contributing to the final solution. In this case, isolation of the averaging kernel in depth has been abandoned in favor of reducing the contributions from different measurements to reduce the final uncertainty.




Figure 4.2: (a) Inversion coefficients A_1 (red) and A_2 (blue) as a function of the regularization coefficient λ . (b) Uncertainty on the inversion solution σ_w as a function of λ . The colored vertical lines indicate values of λ that are used in Figure 4.3.

Figure 4.2b shows the final uncertainty on the solution σ_w as a function of λ . Again, at high regularization the solution tends towards one that minimizes the uncertainty. The solution shows an uncertainty that converges to a constant at either extreme of regularization. For low values, the solution uncertainty is amplified compared to the individual values of σ from each measurement (dashed lines).

Figure 4.3 shows the averaging kernel $Q(z)$ for a few values of λ spanning the region of change in Figure 4.2. Since there are no rules as to what value of λ to pick, we use this figure along with Figure 4.1b to decide what solution allows the most useful interpretation. The lowest values of λ allow for the most isolated averaging kernels, but amplify the final uncertainty. However, in order to claim that the inversion solution represents the subsurface flow at the depth of the target function, we must sacrifice uncertainty for some degree of isolation. Luckily, in the case of two boxcar modes, the final uncertainty converges to a constant when the averaging kernel becomes most compact. In the next case, we will see that this is not always the




Figure 4.3: Averaging kernels for one-dimensional boxcar inversion for a few values of λ . Low values (blue) average over most of the domain, while high values (red) create a more isolated averaging kernel.

case.

In this idealized one-dimensional inversion, we find that it is the overlap between sensitivity kernels that allows one to produce compact averaging kernels. The range of sensitivity kernels included in the inversion increases the possibilities for creating isolated measurements at a variety of depths. Through mapping out how both the final uncertainty and the averaging kernel respond to regularization, we are able to pick a value for λ that provides a useful solution.

4.3.2 Solar Wave Mode Kernels

The solution to a one-dimensional inverse problem is highly dependent on the structure of the sensitivity kernels. In the previous example, two boxcar functions were used to demonstrate how the inversion solution depends on regularization and estimated uncertainty. To extend this simple example to something closer to the full inverse problem in high-resolution ring-diagram analysis, we now look at a one-dimensional inversion using realistic sensitivity kernels. These kernels are derived in Birch et al. (2007) for the purpose of three-dimensional ring-diagram inversions (see Chapter 2). The one-dimensional sensitivity kernels for this example problem (Figure 4.4) were generated by horizontally integrating the full three-dimensional kernels. I use 220 modes with unique sensitivity kernels. A selection of sensitivity kernels spanning the domain of interest are plotted in Figure 2.6.




Figure 4.4: A sample of the 220 unique one-dimensional sensitivity kernels used in the solar wave mode inversion. The different modes are denoted by their radial order n and their wavenumber k . The dashed lines indicate the locations of the shallow (1 Mm), medium (5 Mm), and deep (20 Mm) targets.

Like the boxcar functions from the previous example, these kernels generally have sensitivity from the surface down to some depth that depends on the mode in question; however their sensitivity over that range is not uniform. This creates a similar situation where deeper penetrating modes will generally be weighted differently than shallow modes. These kernels are more complicated not only because their amplitudes oscillate in depth, but they are not always positive valued throughout the domain. For the estimated uncertainty on each mode, we use the measurements from real 16° tiles (Figure 3.5). These estimates form the diagonal of the covariance matrix, and the off-diagonal elements are assumed to be zero.

Again, we set up a minimization problem in the form of Equation 4.20. Unlike the problem with only two modes, the solution to Equation 4.20 cannot be separated easily by mode. Instead, we solve the corresponding matrix problem provided in Equation 4.22 as-is. While the form of the solution is not as readily deconstructed as in the previous example, we find it is qualitatively similar to before. The target function is projected on to the basis set of kernels, and the regularization acts to reduce the projection coefficients in the event of an underconstrained problem.

Using a delta function in depth as the target function, we can solve the matrix problem above for many values of regularization λ and target depth z_t to see how the inversion coefficients relate. We pick three characteristic values for the target depth: a shallow target depth (1 Mm) where every kernel has some

overlap with the target function, a medium target depth (5 Mm) where only half of the kernels overlap the target function, and a deep target depth (20 Mm) where very few kernels overlap the target function.




Figure 4.5: One-dimensional averaging kernels for a solar-like depth inversion with a target depth of 1 Mm (a), 5 Mm (b), and 20 Mm (c). The different colors show averaging kernels produces with different amounts of regularization. As the regularization parameter λ increases, the vertical localization of the averaging kernel suffers. For constant λ , the width of the averaging kernels gets larger for deeper target depths.

Figure 4.5 shows the averaging kernels for a each target depth at a few selected values of the regularization parameter λ . The widths of the averaging kernels for all target depths get wider for larger values of λ . This reduces the final uncertainty, as the inversion is allowed to average together a larger number of independent measurements. It is interesting to note that for deeper target depths, the main peak of sensitivity is not always symmetric about the target depth. This is due to the fact that measurements with shallow acoustic cavities typically have lower uncertainties, and thus are preferred when creating an averaging kernel. Due to the irregular nature of the sensitivity kernels, the averaging kernels consistently show positive and negative sidelobes away from the primary peak of sensitivity. These sidelobes complicate the interpretation

of the inversion solution, as they can bring in information from far-away depths. While sidelobes are largely unavoidable, it is beneficial to look for averaging kernels with minimal sidelobes in depth.




Figure 4.6: (a) Inversion coefficients for a shallow mode (red) and a deep mode (green) as a function of λ for the shallow target (solid lines), medium target (dotted lines), and deep target (dashed lines). (b) Estimated uncertainty on the inversion result as a function of λ for the three target depths. The horizontal dashed lines indicate the minimum and maximum of the uncertainties on all modes individually. Note the range with which the final uncertainty can be either amplified or reduced.

Figure 4.6 shows both a selection of inversion coefficients for each target depth (a) and the final propagated uncertainty for each target depth (b) as a function of λ . In this, the solution for the shallow target depth is shown with solid lines, that of the medium target depth with dotted lines, and that of the deep target depth with dashed lines. The (a) panel simply demonstrates that the inversion coefficients are not always monotonic functions of λ . As the regularization term lessens and higher-uncertainty measurements are allowed to feed into the solution, lower-uncertainty measurements shift roles from being the sole contributors to the solution and instead to compensate for the new contributions. In contrast to the boxcar kernel example, Figure 4.6b shows that the final uncertainty does not converge as readily at low values of λ , and instead rises continually over the range of regularization shown here. But as before, high values of λ allow the final uncertainty to converge to a constant value. Due to the plethora of independent measurements, the final uncertainty is allowed to dip below the minimum uncertainty for any one measurement. In these cases, the inversion has essentially averaged many measurements together instead of playing them off of each other to create a useful solution. The transition between these two regimes is slightly different for each target

depth, suggesting the need to tune the value of λ independently for each depth.

Between Figure 4.5 and Figure 4.6, we can begin to pick a value of λ that is best suited for each target depth. Once again, we can start with the criterion that the averaging kernel needs to allow a useful interpretation. For each depth, values of λ less than 10^{-2} show reasonable isolation in depth and a peak in sensitivity that matches the target depth. Values lower than 10^{-4} start to show significant sidelobes in depth, which is to be avoided. While we are free to pick a different value of λ for each depth, we find that there is little need in this scenario as a single value of around $\lambda = 10^{-1}$ would likely work well at all depths. The final uncertainties for this range of λ generally show an amplification of uncertainty relative to the original measurements. This is to be expected, as the inversion has produced estimates of the subsurface flow field that average over a smaller volume of space than any of the original measurements.

4.3.3 Summary of Depth Inversion Behavior

The full three-dimensional helioseismic inversion is able to create estimates of the subsurface velocity at isolated locations. In depth, this isolation is possible through the use of measurements from different wave modes that have different acoustic cavities within the Sun. These one-dimensional inversions demonstrate that it is possible to achieve varying degrees of vertical compactness, allowing for a tradeoff with the final propagated uncertainty.

The regularization parameter λ is the mediator between compact averaging kernels and low uncertainty. Tuning an appropriate value for λ involves mapping out how both the averaging kernels and the final uncertainty react, then picking a value that allows for a useful inversion solution while minimizing the uncertainty. While this process needs to be performed for each target depth, the same values for λ can be used on different data sets as long as the measurement uncertainties do not change significantly.

4.4 Horizontal Inversions with Spatial Invariance

The measurements and sensitivity kernels used in the full three-dimensional inversion have vastly different properties when analyzed in depth or as a function of disk position. In the previous section, we find that the depth-dependence of the measurement uncertainties and sensitivity kernels produce complicated,

asymmetric averaging kernels. In this section, I will consider the horizontal directions of the full inversion by once again using a simplified one-dimensional example. Near disk center, the data coverage and measured uncertainties are essentially constant (Figure 3.6 and Figure 3.7). Thus, the ability for the three-dimensional inversion to create isolated averaging kernels in depth is constant as a function of horizontal position. From this, we can construct a new inversion that assumes a translationally invariant sensitivity kernel and uniform measurement uncertainty. The translational invariance permits simple analysis of the inversion process in spectral space.

We begin by defining the frequency shift measurement made at a location x as:

$$u(x) = \int K(x' - x)v(x')dx' \quad (4.26)$$

For now, we will assume the measurements have been made continuously in space, and later we can consider finite sampling effects. Equation 4.26 is a straight-forward convolution. It is thus the goal of the one-dimensional horizontal inversion to perform a deconvolution on the measurements. We can further mimic the horizontal component of the full three-dimensional inversion by using a rounded boxcar sensitivity kernel (Figure 4.7). The exact horizontal planforms of the sensitivity kernels used in the full inversion vary slightly between each other, but generally have this shape.




Figure 4.7: (a) Sensitivity kernel for spatially invariant one-dimensional inversion. The kernel extends one unit of distance and is tapered by a cubic spline at each end. Measurements made at different positions along the domain have identical kernels, and the measurement position determines the central position of the kernel. (b) Fourier power spectrum of the sensitivity kernel in (a). The horizontal axis is scaled by the wavenumber associated with the size of the sensitivity kernel in real space. The zeroes in the power spectrum are due to the fact that the kernel is compact and occur roughly at integer multiples of this characteristic scale.

4.4.1 Solution in Spectral Space

Writing the OLA minimization equation to include translational symmetry:

$$\Psi = \int |T(x) - Q(x)|^2 dx + \lambda \int dx \int dx' A(x) A(x') C(x, x') dx, \quad (4.27)$$

$$Q(x) = \int A(x') K(x - x') dx'. \quad (4.28)$$

The goal is then to find a function for $A(x)$ that minimizes Ψ . If we were to discretize the problem at this point, we would construct yet another matrix problem that solves for a set of inversion coefficients $A(x_j)$. Fortunately, we can instead take advantage of the spatial invariance of the problem at hand and find an analytic solution. Using Parseval's theorem and the convolution theorem to express the two terms on the right-hand-side of Equation 4.27 in Fourier space, we obtain the following minimization equation:

$$\Psi = \int |T(k) - A(k)K(k)|^2 dk + \lambda \sigma \int |A(k)|^2 dk. \quad (4.29)$$

In this, I have ignored the off-diagonal components of the covariance for simplicity. The definition of the averaging kernel is simplified to

$$Q(k) = A(k)K(k). \quad (4.30)$$

Here, we denote the Fourier transform of a function $f(x)$ by replacing the variable with its Fourier pair $f(k)$. By expressing the OLA minimization equation in Fourier space, we can solve for each Fourier component independently. This removes the need to construct a matrix problem to solve for each inversion coefficient simultaneously, simplifying the inversion tremendously. Setting λ to zero reveals a simple solution,

$$A(k) = T(k)/K(k). \quad (4.31)$$

In the case of a delta function target, the Fourier transformed inversion coefficient function is simply the reciprocal of the Fourier transformed sensitivity kernel, $A = K^{-1}$.

The need for regularization becomes apparent when considering the power spectrum of the sensitivity kernel (Figure 4.7b), which has zeroes at periodic scales. Using Equation 4.31 to solve the inversion for zero regularization causes a significant amplification of power at these scales, which tends to amplify noise.

Regularization is used to mitigate the effect of having significant dips in the denominator of Equation 4.31.

Including regularization, the function for $a(k)$ that minimizes Equation 4.29 becomes

$$A(k) = \frac{K(k)T(k)}{|K(k)|^2 + \lambda\sigma^2}. \quad (4.32)$$

The regularization term used to create this solution assumes that the frequency shift measurements are all statistically independent so that the off-diagonal terms of the covariance matrix can be ignored. This assumption was justified for the depth inversions, as the measurements came from disparate wave modes. For measurements made of the same wave modes, this is not necessarily true. The covariance between measurements made from adjacent overlapping tiles is approximately proportional to the tile overlap fraction squared:

$$C(x, x') = \sigma^2 \left(\frac{\int K(x)K(x-x')dx}{\int K(x)^2 dx} \right)^2. \quad (4.33)$$

Through use of Parseval's theorem and the convolution theorem, the solution in Fourier space ends up having a $|K(k)|^2$ included in the regularization term. There is an additional constant due to the normalization of Equation 4.33 that can be wrapped up into the parameter λ . The solution when a full covariance is included has the same form of Equation 4.31 with a multiplicative factor that depends on λ :

$$A(k) = \frac{K(k)T(k)}{|K(k)|^2 + \lambda|K(k)|^2} = \frac{K(k)T(k)}{|K(k)|^2(1 + \lambda)}. \quad (4.34)$$

Once the solution is normalized properly, this constant is removed, leaving a solution that once again does not depend on λ . It is interesting that when this form of the covariance is included, the inversion has no option but to create a solution that is effectively zero regularization. This result is a consequence of the covariance depending on the overlap of the kernels. This effect will be used later in the full three-dimensional inversion to create a regularization term that treats the horizontal and vertical components differently.

Since the inclusion of the full covariance between measurements creates a degenerate solution that does not depend on λ , I have chosen to use only the diagonal component of the covariance for the regularization of the present example inversion. However, In order to accurately represent the final solution the propagated uncertainty does use the full covariance matrix. Thus the inversion no longer uses the true uncertainty on the solution for regularization, but instead just a weighting against each inversion coefficient.

4.4.2 Regularization Tuning

Figure 4.8 shows averaging kernels for a few values of λ along with the original sensitivity kernel (dashed line) for comparison. As the regularization parameter λ is decreased, the averaging kernel becomes dramatically narrower. There is notable ringing in the averaging kernel that appears at fixed intervals. This ringing is inherent to the inversion procedure and is caused by the finite extent of the sensitivity kernel. This effect can also be seen in the power spectrum (Figure 4.8b) as severe dips in the spectral sensitivity. Measuring subsurface flows with a kernel like the one used here creates a measurement set that is insensitive to certain wavenumbers. Actual frequency shift measurements do not show as strong of dips in the power spectrum, as these wavenumbers have likely been filled in with random noise. The subsurface flow information at these wavenumbers is lost, so it is the role of regularization to make sure the inversion does not attempt to amplify them too much. The missing information could be brought back by using two or more tile sizes that have zeroes in the power spectrum at disparate wavenumbers. In practice, the spatial and spectral ringing cannot be completely removed, but their effects can be mitigated with the significant computational cost of using multiple tile sizes.

The purpose of tuning the regularization for the one-dimensional depth inversion was to find a useful tradeoff between solution uncertainty and localization of the averaging kernel. The tradeoff in this horizontal inversion is largely the same, but the localization of the averaging kernel can be analyzed in a more straightforward way. Figure 4.9 shows the averaging kernel FWHM (a) and the final solution uncertainty (b) as a function of regularization parameter λ . The dashed lines in Figure 4.9a indicate the spatial extent of the sensitivity kernel (upper) and the grid spacing between adjacent frequency measurements on the solar disk (lower). This grid spacing is also the resolution at which the sensitivity kernels have been discretized, but is not necessarily the final effective resolution of the inversion. The ratio between these values has been chosen to mimic that in the full three-dimensional inversion. As seen with the averaging kernels themselves in Figure 4.8, the FWHM of the averaging kernel gets smaller as λ is decreased. In general, these show a similar trend as in the one-dimensional depth inversion, where lower regularization leads to a more localized solution at the cost of increased uncertainty.




Figure 4.8: (a) Averaging kernels for the spatially-invariant one-dimensional inversion for a few values of the regularization parameter λ . The dashed line indicates the shape and size of the original sensitivity kernel. As λ decreases, the averaging kernels become more spatially localized. Ringing occurs for every value of λ that creates an averaging kernel with a FWHM narrower than the sensitivity kernel. The spacing of the ringing is determined by the size of the sensitivity kernel. (b) Power spectrum of the averaging kernels shown in panel (a) as a function of normalized wavenumber. A wavenumber of one indicates the spatial scale of the sensitivity kernel. The dashed line is the power spectrum of the sensitivity kernel, and the dotted line is the power spectrum of the target function. The effects of Equation 4.32 can be seen here, where the inversion acts to amplify power until it matches the target function.

The choice of regularization parameter λ once again comes down to localization versus uncertainty. Unlike the depth inversion from the previous section, a lack of spatial localization in this case does not always mean the solution cannot be easily interpreted. In depth, an improperly tuned inversion can create an asymmetric averaging kernel that can preferentially bring in near-surface contributions. For the translationally invariant case presented here, a low- λ inversion simply creates a broader averaging kernel that remains centered on the appropriate point in space. Tuning λ for spatial localization in this case is a straightforward




Figure 4.9: (a) The full width at half maximum (FWHM) of the averaging kernel as a function of λ . The upper dashed line indicates the horizontal size of the sensitivity kernel, and the lower dashed line is the separation between adjacent measurements. The measurement spacing is the theoretical limit of inversion resolution and is only achievable for very low values of λ . (b) Estimated uncertainty as a function of λ . The solid line is the uncertainty calculated with the proper covariance between measurements built in, and the dotted line is the uncertainty with no covariance between different measurements. The inversion regularizes against the latter, but the proper covariance is necessary to obtain a realistic estimate of solution uncertainty.

issue of resolution. The particular value of λ that is chosen for an inversion will depend on the quality of the data and the scientific questions that need to be answered.

4.5 Three-Dimensional Helioseismic Inversions with MORDI

The previous sections of this chapter have built up the framework around which we can construct a full three-dimensional inversion for ring-diagram frequency shift measurements. The set of distinct wave modes measured at each location on the disk determines how well one can create estimates of subsurface flows at an isolated depth, and the spacing between adjacent measurements on the disk influence the final horizontal resolution of the solution. Using the OLA inversion method, estimates of the subsurface flow can be made at highly localized points in space both horizontally and vertically. The three-dimensional OLA equation that is solved is

$$\Psi = \int d\mathbf{x}' \int dz' |T(\mathbf{x}', z') - Q(\mathbf{x}', z')|^2 + \lambda \mathcal{R}_1 + \mu \mathcal{R}_2 \quad (4.35)$$

$$\mathcal{R}_1 = \sum_i \sum_{j,k} A_i(\mathbf{X}_j) A_i(\mathbf{X}_k) C(\mathbf{X}_j, \mathbf{X}_k) \quad (4.36)$$

$$\mathcal{R}_2 = \sum_i \sum_j |\sigma_i A_i(\mathbf{X}_j)|^2 \quad (4.37)$$

where the averaging kernel is given by

$$Q(\mathbf{x}, z) = \sum_i \sum_j A_i(\mathbf{X}_j) K_i(\mathbf{x} - \mathbf{X}_j, z). \quad (4.38)$$

In these, the subscripts j and k refer to the locations of tiles on the solar disk. Further, I again assume that the covariance between two measurements is zero for disparate wave modes and proportional to the overlap fraction of the apodization function $f(\mathbf{x})$ for identical wave modes:

$$C(\mathbf{X}_j, \mathbf{X}_k) = C(\mathbf{X}_j - \mathbf{X}_k) = \left| \frac{\int d\mathbf{x} f(\mathbf{x} - \mathbf{X}_j) f(\mathbf{x} - \mathbf{X}_k)}{\int d\mathbf{x} f(\mathbf{x})^2} \right|^2. \quad (4.39)$$

Here, we use a fully three-dimensional target function and set of sensitivity kernels. The inversion coefficients are expressed as two-dimensional maps $a_i(\mathbf{X})$ where i indicated the wave mode measured from the ring-diagram power spectra. For numerical stability, the target function is chosen as a Gaussian function with widths in each direction slightly smaller than the grid spacing. The reason behind the two different regularization terms will be explained in detail later in this section.

Thanks to the translational invariance of the frequency shift uncertainties and sensitivity kernels, we can once again take advantage of minimizing the OLA equation in Fourier space. Just as in Section 4.4, this decouples the problem into an independent equation to be solved at each wavenumber vector $\mathbf{k}$. However, in the one-dimensional horizontal inversion, this resulted in a simple algebraic expression for the inversion coefficients at each wavenumber. In this three-dimensional problem, we are left with solving for the inversion coefficients of each distinct wave mode simultaneously at each wavenumber vector. This is very similar to performing a set of one-dimensional depth inversions at each horizontal grid point in real space independently, except in this case the inversion operates on each wavenumber component of the coefficients, kernels, and target function independently. Computationally, the two scenarios are nearly identical. Solving a three-dimensional helioseismic inversion by representing only the horizontal directions in Fourier space has been done in Jackiewicz et al. (2012) and has been called a Multi-Channel Deconvolution. This is an extremely efficient method of performing large three-dimensional inversions of helioseismic data, as it allows one to increase the number of horizontal solution points with ease. In this thesis, I have applied this method to high-resolution data sets in order to obtain measurements of small-scale convection in the NSSL.

The number of inversion coefficients that must be solved for simultaneously determine the size of the matrix problem that is created (Section 4.2). For the mode set and grid sizes considered in this thesis, solving a full three-dimensional inversion in real space requires inverting a $3 \cdot 10^7 \times 3 \cdot 10^7$ element matrix and requires approximately 6.5 petabytes of memory. This size of problem is not tractable on even modern supercomputers, and in the end only offers a single analysis period of inverted flow fields. The ability to deconstruct the horizontal domain into independent inversions significantly reduces the computational burden, as each wavenumber can be solved for serially instead of simultaneously. The inversion at each wavenumber essentially becomes a one-dimensional depth inversion with around 220 modes. This opens up the possibility for vast amounts of parallelism in the inversion code, limited only by the number of grid points present in the horizontal directions.

The Massively-Overlapped Ring-Diagram Inversion (MORDI) code performs three-dimensional, ring-diagram, OLA inversions following this procedure. I have developed this code to take advantage of the high-quality frequency shift measurements made possible by the advances detailed in the previous chapter. MORDI takes as input a set of frequency shift measurements from various modes spanning some rectangular region in longitude and latitude on the solar surface, a set of three-dimensional sensitivity kernels that match the frequency-shift mode set, and a set of control parameters for quantities like the regularization and target depth. The maps of measured frequency shifts and the associated sensitivity kernels are Fourier transformed horizontally and left in real space vertically. As seen in Section 4.4, this drastically reduces the computational burden of finding the optimal solution to the OLA equation. The solution to the inversion is computed independently and in parallel at each horizontal wavenumber. MORDI reconstructs and outputs a real-space two-dimensional flow field for every target depth specified and can optionally output the full three-dimensional averaging kernel for the associated target depth. The purpose of this section is to demonstrate how the full three-dimensional inversion can be treated like the simplified inversions in terms of regularization tuning and averaging kernel localization.

Since the inversion relies on Fourier transforms to solve the optimization equation, we must ensure that the frequency shift maps and sensitivity kernels respect the assumptions of the transform. The frequency shift measurements are made for a variety of wave modes over a rectangular grid of points in longitude and

latitude. A single wave mode defines a map of frequency-shift measurements that is Fourier transformed in both directions (longitude and latitude) during the inversion. Each of these maps must be apodized and zero-padded to prevent ringing in Fourier space. Hard edges at the ends of the maps can produce ‘echos’ in the final inversion solution throughout the domain. Smoothly tapering the edges of the frequency-shift maps to zero significantly reduces this effect. In practice, I have found that the inversion produces cleanest results when each map is padded by at least two tile sizes on all ends and the data is tapered to zero over the extent of at least one tile size. While it is always safer to add more padding than the bare minimum, this increases the computational burden of running the inversion. For the 90° square regions used in this thesis made with 16° tiles, this results in an inversion grid that is 154° on a side (617 pixels). Since the apodization reduces the usable area of the inversion, the final solution will not have the same extent as the frequency-shift maps. For the current inversions, the final usable area of the inversion is around 58° on a side.

4.5.1 Multi-Dimensional Localization

To allow more independent control over the horizontal and vertical aspects of the inversion, I have included two regularization terms. The first regularization term $\mathcal{R}_1$ (controlled by the parameter λ) is the standard regularization that is proportional to the final propagated uncertainty and includes the full covariance matrix. The second regularization term $\mathcal{R}_2$ (controlled by the parameter μ) is equal to the final propagated uncertainty in the case of no off-diagonal components of the covariance matrix. As found in the previous section on inversions with translational invariance, using this first form of regularization by itself leads to solutions that are identical to those made with no regularization. This is due to the fact that the covariance matrix had the same form as the kernel overlap matrix.

In the three-dimensional inversion, this degeneracy is not exact. As seen in Figure 2.6, the horizontal planforms of the sensitivity kernels are similar to the apodization circle but not exactly the same and the difference grows with depth. Had they been identical, the parameter λ would only effect the balance of uncertainty to vertical localization of the averaging kernels and by itself would leave the horizontal localization as high as possible, regardless of final uncertainty. In this case, it would be the role of μ to control the tradeoff between uncertainty and the overall localization in all dimensions. In this way, the two regulariza-

tion parameters can be used to adjust the horizontal and vertical aspects of the averaging kernels almost independently.

4.5.2 Regularization Tuning

Since the horizontal overlap between sensitivity kernels is not exactly proportional to the tile overlap, this separation between the two regularization terms is not as clean. Still, they are capable of influencing the vertical and horizontal localization to different degrees. The inclusion of an extra regularization term results in a two-dimensional parameter space of regularization (λ, μ) that needs to be explored in order to find a useful inversion solution. Since very little changes between inversions computed with different regularization values, the MORDI code can be told to loop over a range of both λ and μ while everything is held in memory. This significantly reduces the computational cost of running multiple independent inversions and provides an easy way of searching through the λ - μ parameter space to select appropriate values.

There are two categories of diagnostics used for tuning the three-dimensional, ring-diagram inversion. The first category is about interpretation and includes diagnostics like the averaging kernel localization and the final propagated uncertainty. The second category is about the actual flow field solution and its characteristics. The following few figures show these diagnostic quantities as a function of the two regularization parameters for two different target depths (1 Mm, 20 Mm). Instead of presenting the horizontal forms of the myriad of averaging kernels, I will use the horizontal FWHM as a simpler measure just as in Section 4.4.

Figure 4.10 shows the vertical structure of the averaging kernels for the 1 Mm depth inversion for different values of λ (different colors) and μ (different sub-figures). While all of the averaging kernels resulting from this set of inversions seem to create an averaging kernel centered on the target depth, there is significant variance as to how well the inversion has isolated the sensitivity to this depth. As expected from the previous arguments about multi-dimensional localization, the two regularization parameters λ and μ seem to influence the vertical localization of the solution equally.

The various panels of Figure 4.11 show diagnostic quantities for a target depth of 1 Mm as a function of μ and for different values of λ (line colors). Panel (a) shows the horizontal FWHM of the resulting averaging




Figure 4.10: Horizontally-integrated sensitivity for averaging kernels with a target depth of 1 Mm (dashed line). The different panels show averaging kernels for different values of μ and the colored lines in each panel indicate different values of λ . Both λ and μ effect the vertical localization of the averaging kernel. Low regularization values result in more localized averaging kernels, but also increase the amplitude of sidelobes.

kernels. Just as in Figure 4.9a from the previous section on translationally invariant inversions, the dashed lines indicate the tile size (upper) and the measurement grid spacing (lower). As expected, the horizontal localization is not significantly modified by λ due to the similarities between the sensitivity kernel overlap and the covariance matrix. Panel (b) shows the final propagated uncertainty (accounting for measurement covariance). As in Figures 4.6b and 4.9b, the dashed lines indicate the minimum and maximum values of estimated frequency-shift uncertainty used in the inversion. Panel (c) shows the rms zonal velocity of the flow field produced by the inversion. In general, this quantity cannot be used to reliably tune the inversion as the correct value is unknown. For this particular target depth, we may use this value to compare to other measures of photospheric convective velocities to make sure the inversion solution makes sense (see Chapter 5). Panel (d) shows the rms-to-error ratio, computed as simply the value in panel (c) divided by the value in panel (d) for each value of λ and μ . This gives a useful estimate of how significant any one point in the solution is. Values greater than one allow a straightforward interpretation of the flow field without any post-inversion averaging. The cross in each panel indicates the value of λ and μ used for later inversion results. In all, these diagnostics do not appear to depend strongly on λ . This is because a target depth of 1 Mm does not require much vertical localization, which the parameter λ has almost independent control over.

Figure 4.12 shows power spectra of the zonal velocity map obtained for different regularization values. These power spectra are originally computed as a function of two-dimensional wavenumber ($\mathbf{k}$) and have been integrated at each wavenumber magnitude ($k = |\mathbf{k}|$). While the rms zonal velocity (Figure 4.11c) can be useful for making sure the choice of regularization provides a reasonable answer, the power spectra allow one to consider which spatial scales are present in the solution and contribute to the rms value. The largest effect that regularization has on the power spectra is the amplification of high-wavenumber power. This is expected, as the goal of the inversion is to produce high-resolution maps which necessarily have more power at high wavenumbers. These power spectra are similar to the power spectra of the averaging kernel in Figure 4.8b. While the lower-wavenumber power on scales comparable to or larger than the tile size remains constant for most values of λ and μ , it is possible to amplify this power when both regularization terms are set very low. At this depth, the spectral signature of supergranulation ($\ell = 120$) can be seen in a few of the




Figure 4.11: Diagnostic plots for full 3D inversion targeting a depth of 1 Mm. The colored lines in each sub-figure show data for different values of λ . The crosses indicate the values of λ and μ used in later chapters. (a) Horizontal FWHM of the averaging kernel as a function of μ . The dashed lines in (a) are the original tile size (16°) and the tiling scheme resolution (0.25°). (b) Final propagated uncertainty (including covariance) of the inversion solution. The dashed lines in (b) are the minimum and maximum uncertainties of the frequency-shift measurements. (c) rms zonal velocity, showing a general increase for lower values of regularization. (d) Ratio of rms to uncertainty as a proxy for the signal-to-noise ratio.

curves, particularly for $\mu = 10^{-3}$.

To see visually what the flow fields look like for different values of λ and μ , Figure 4.13 shows maps of the horizontal divergence for the 1 Mm target depth for four different values of λ (bottom to top) and μ (left to right). The color table of each map has been normalized to two times the rms value of the map. The most obvious effect seen in this figure is that of horizontal resolution changing with regularization. As seen in the power spectra of these maps in Figure 4.12, the large-scale features are present in all cases and it is the presence of the small-scale features that changes with λ and μ . Despite this persistence of large-scale patterns, when scaled against the entire signal, the large-scale patterns seem to disappear behind the




Figure 4.12: Power spectra of the zonal velocity at 1 Mm depth for various values of μ (the different sub-figures) and λ (colored lines). The values of λ corresponding to each of the colors are indicated in the previous two figures.

stronger small-scale flows.

Figure 4.14 shows the vertical profile of a range of averaging kernels for the inversion targeted at 20 Mm depth. These averaging kernels have much greater variance over the displayed range of regularization values than the 1 Mm depth inversion due to the greater number of wave modes that are used to isolate sensitivity deeper into the Sun. Once again, we find that the two regularization parameters have roughly equal influence over the degree of isolation in depth. Unlike the 1 Mm target depth, we find here that improperly tuned inversions can create averaging kernels that are not at all useful for creating flow fields at isolated depths. The averaging kernels for high values of λ and μ not only have sensitivity that spans nearly the entire domain in depth, but contain significant oscillations in sensitivity that hinder the interpretation of the solution. For inversions like this one, tuning the regularization often involves finding the maximum values of λ and μ that permit a useful interpretation of the solution.

Figures 4.15, 4.16, and 4.17 respectively show inversion diagnostics, solution power spectra, and divergence maps for a target depth of 20 Mm. These figures show that the inversion acts largely the same way as it does for the shallow target depth, but is influenced to a larger degree by the regularization parameter λ . Since λ has little effect on the horizontal localization, it is the primary controller of vertical localization. Unlike the shallow target depth of 1 Mm, this target depth allows the inversion to utilize every measured wave mode to create an averaging kernel with isolated sensitivity at depth. It is interesting to note that




Figure 4.13: Maps of the inversion solution horizontal divergence for the shallow (1 Mm) target depth for various values of λ (bottom to top) and μ (left to right). Each panel covers the same 50° by 50° area on the solar disk and the color table has been scaled by the standard deviation of the map. Red indicates positive horizontal divergence and blue indicates negative. Comparing these maps with one another, it becomes clear how regularization affects the spatial resolution of the inversion solution. The asterisk indicates the solution that uses the regularization values indicated in Figure 4.11 by the cross.

the divergence patterns seen in the high-regularization solutions at 20 Mm (Figure 4.17) look similar to the high-regularization solutions at the surface (Figure 4.13). This is due to the fact that the averaging kernels in these cases span a large range of depths and therefore have significant overlap. Despite having different target depths, the solutions will show nearly the same patterns as they represent averages over




Figure 4.14: Horizontally-integrated sensitivity for averaging kernels with a target depth of 20 Mm (dashed line). The different panels show averaging kernels for different values of μ and the colored lines in each panel indicate different values of λ . Both λ and μ effect the vertical localization of the averaging kernel. High regularization values result in averaging kernels that show little localization near the target depth. Low regularization values result in more localized averaging kernels, but also increase the amplitude of sidelobes.

similar volumes.




Figure 4.15: Diagnostic plots for full 3D inversion targeting a depth of 20 Mm. The colored lines in each sub-figure show data for different values of λ . The crosses indicate the values of λ and μ used in later chapters. (a) Horizontal FWHM of the averaging kernel as a function of μ . The dashed lines in (a) are the original tile size (16°) and the tiling scheme resolution ($0^\circ.25$). (b) Final propagated uncertainty (including covariance) of the inversion solution. The dashed lines in (b) are the minimum and maximum uncertainties of the frequency shift measurements. (c) rms zonal velocity, showing a general increase for lower values of regularization. (d) Ratio of rms to uncertainty as a proxy for the signal-to-noise ratio.

To tune the regularization parameters for the full three-dimensional inversion, figures such as these have been created at each depth and studied to determine what values provide the most useful solutions. While one could pick new values for λ and μ at each depth, it seems more reasonable to either pick a single value for each to cover every depth or set each to be a smooth function of depth. In practice, a power law function works well for the data sets considered here:

$$\lambda(z) = 10^{az+b}. \quad (4.40)$$

Here, a and b are coefficients that determine how the regularization parameter changes with depth. A similar




Figure 4.16: Power spectra of the zonal velocity at 20 Mm depth for various values of μ (the different sub-figures) and λ (colored lines) as in Figure 4.12.

relation can be made for μ . The variation with depth (given by a) is often small compared to the value of b .

The crosses in Figures 4.11 and 4.15 indicate the rough values of λ and μ used in later chapters. While some inversions have been tuned slightly differently, these values are approximately $\lambda = \mu = 10^{-3}$ at all depths (no depth variation). While one is free to choose the regularization values based on any combination of the diagnostics, I have chosen them such that the rms-to-error ratio is approximately two at all depths. Other diagnostic quantities (like the averaging kernel profiles) have been considered as well, but only to ensure that the inversion solution can be interpreted easily.

4.5.3 Selected Averaging Kernels

With the values for the regularization parameters λ and μ from the previous section, a set of averaging kernels have been computed for target depths spanning 0 to 30 Mm depth in 1 Mm increments. Many of these averaging kernels show nearly identical qualities as those computed for adjacent target depths. A selection of four characteristic averaging kernels are shown in Figure 4.18 and a more complete set of averaging kernels are shown in more detail in Appendix B. The dashed circles in the right panels indicate the size of the apodization circle for a 16° tile. Horizontally, the averaging kernels show a much greater amount of localization than the original sensitivity kernels. The spatial ringing shown in the translationally invariant one-dimensional example is present in these averaging kernels but is less apparent.

A primary consequence of the inversion is obtaining localized sensitivity at the given target depth.




Figure 4.17: Maps of the horizontal divergence for the deep (20 Mm) target depth three-dimensional inversion for various values of λ (bottom to top) and μ (left to right). Each panel covers the same 50° by 50° area on the solar disk and the color table has been scaled by the standard deviation of the map. Red indicates positive horizontal divergence and blue indicates negative. Comparing these maps with one another, it becomes clear how regularization affects the spatial resolution of the inversion solution. The star indicates the solution that uses the regularization values also indicated in Figure 4.15.

The panels on the left side of Figure 4.18 show the horizontally integrated sensitivity of the various averaging kernels as a function of depth. The width of the primary peak of sensitivity increases monotonically as a function of target depth and has a measured FWHM of around one third the target depth. At shallower depths than the target depth, the averaging kernels tend to show ringing with a nearly equal amount

of positive and negative sensitivity. At deeper depths, the ringing persists but has a significantly lower amplitude.

From the averaging kernels presented in Figure 4.18, I find it appropriate to interpret the inversion solution for a given target depth as the subsurface flow field at that depth. The primary peak of sensitivity sits almost exactly at the target depth and the ringing in both the horizontal and vertical directions is minimal. The effective resolution of the inversion solution is characterized by the vertical and horizontal extent of the averaging kernels and thus must be taken into account when interpreting the flow fields.

4.5.4 Conclusions

The inversion technique presented in this chapter is capable of turning highly-overlapped, ring-diagram, frequency-shift measurements into three-dimensional maps of subsurface flows. The characteristics of the data and associated averaging kernels allow a better understanding of how the inversion works through analyzing one-dimensional inversions. Inversions in depth show us that localization in depth depends highly on the set of wave modes measured during ring fitting. Horizontal inversions with translational invariance show us that flows smaller than the analysis region can be resolved through highly overlapped measurements. Both types of inversions permit straightforward tuning of the regularization parameters in order to obtain useful results.

The full three-dimensional inversion demonstrates many of the aspects of the simplified one-dimensional inversions. The process of picking appropriate regularization parameters is inherent to the inversion procedure and necessary for obtaining scientifically useful results. The ability to create finely-tuned high-resolution flow maps has enabled a new class of observations of solar convection. The following chapters rely on observations of subsurface flows made with MORDI.




Figure 4.18: Left: Horizontally integrated averaging kernels plotted as a function of depth. The four panels from top to bottom show averaging kernels targeting different depths. The vertical lines indicate the approximate depth of maximal sensitivity. Right: Cuts of each averaging kernel at the depth indicated by the vertical lines in the left panel. The dashed circles indicate the size of the apodization circle for a 16° tile. This figure has been formatted the same as Figure 2.6 for easy comparison.

Chapter 5

Convective Amplitudes and Rotational Influence

The interaction of convection with rotation establishes Reynolds stresses that are a crucial ingredient in the dynamics of the Sun's convection zone. The degree of rotational influence on convective flows is often characterized by a nondimensional measure called the Rossby number (Ro). The Rossby number is usually defined as the ratio of a typical convective time scale to the rotational period. In order to assay the Rossby number, and therefore the degree of rotational constraint, one must measure the speed of the convective flows with sufficient resolution to sample those scales of motion responsible for Reynolds stresses. Such measurements are now possible within the near-surface shear layer due to the introduction of high-resolution ring-diagram analysis described in Chapters 2 through 4. Using this new implementation of ring-diagram helioseismology to measure previously unresolved scales, we ascertain the strength and spatial scale of convective flows throughout the near-surface shear layer (NSSL). In this chapter, I present a robust analysis of these convective flows and the concomitant Rossby number.

The ring-diagram technique employs highly overlapped analysis regions and an efficient method of three-dimensional inversion to measure convective motions with a resolution that ranges from 10 Mm at the surface to 70 Mm at the NSSL at a depth of 30 Mm. We find the rms horizontal flow speed to peak at 427 m s^{-1} at the photosphere and fall to a minimum of 124 m s^{-1} between 20 Mm and 30 Mm. We also find that just below the photosphere, the convective flows do not feel the influence of rotation ($\text{Ro} > 1$), while flows beyond a depth of 10 Mm are strongly influenced ($\text{Ro} < 1$). These results indicate that the near-surface shear layer is a transition region for the Rossby number and the degree of rotational impact on convective motions.

This chapter is based on work previously published in Greer et al. (2015)¹ with additional analysis and interpretation focusing on the Rossby number. I was the primary author on that paper and performed the data collection, helioseismic inversions, and subsequent data analysis presented in the paper. My co-authors provided valuable guidance on the language and style of the paper. While the work presented in Greer et al. (2015) uses the same analysis procedure detailed in the previous chapters, the results were based on a small subset of the full data set currently available (see Table 2.1 for a full listing). Since the additional analysis pertaining to the Rossby number was performed after the full data set became available, it is based on the full set. When appropriate, I will indicate the portion of the full data set from which a particular result was created.

5.1 Introduction

The Sun’s differential rotation and meridional circulation are maintained by turbulent transport of angular momentum through the action of Reynolds stresses. The spatial and temporal correlations between flow components that generate such stresses are thought to be imposed through Coriolis deflection of convective motions. Due to this deep connection between convection and rotation, the helioseismic measurement of the rotation rate as a function of latitude and depth (e.g., Howe et al. 2000; Thompson et al. 2003) has provided an inflexible constraint on numerical simulations of the convection zone. An even more stringent constraint would be imposed by direct helioseismic determination of the convective flow velocities.

As the first significant step in this direction, Hanasoge et al. (2012), hereafter HDS12, have used time-distance helioseismology to infer the subsurface convective flow amplitude. Specifically, they found an exceptionally low rms velocity which had an *upper* limit of 1 m s^{-1} at a depth of 30 Mm. Such low speeds are in conflict with numerical simulations, which generate speeds that are several orders of magnitude larger (Vögler et al. 2005; Miesch et al. 2008; Rempel et al. 2009; Trampedach and Stein 2011; Weber et al. 2011; Hotta et al. 2014). This discrepancy has cast doubt in both directions, implying that either the dynamical balance achieved in many simulations is far from the one that holds in the Sun, or that helioseismic methods (unknowingly) underestimate velocities below the first few megameters in depth. In

¹ Greer, B. J., Hindman, B. W., Featherstone, N. A., & Toomre, J. 2015, ApJ, 803, L17

this context, we attempt to provide another source of comparison using a new implementation of ring-diagram helioseismology to directly image convective flows throughout the upper 30 Mm. With this technique we find convective amplitudes much larger than those deduced by HDS12, and quite similar to the amplitudes predicted by numerical simulations.

These measurements also allow us to investigate the role of rotation on the observed convective motions. The upper boundary layer of the solar convection zone is a region of strong rotational shear known as the near-surface shear layer. The region spans the upper $0.05 R_{\odot}$ (35 Mm) and exhibits a 3% decrease in the global rotation rate compared to the interior. While global helioseismology has provided precise measurements of the rotation rate in this region (Howe et al. 2000; Thompson et al. 2003), the physical origin of the layer is unknown. It has been proposed that the distinction between this layer and the bulk of the convection zone is that convective flows in the upper 35 Mm are fast enough to traverse the layer without considerable deflection by Coriolis forces (Gilman 1977; Hathaway 1982; Aurnou et al. 2007; Augustson et al. 2011), unlike the flows that exist deeper in the Sun. Hence, the flows within this layer are only weakly organized by rotation and the Reynolds stresses resulting from that organization are reduced. Since Reynolds stresses are a key component in the balance of torques in the Sun, the bulk rotation rate in the NSSL decreases in order to maintain equilibrium. While this theory may play a key part in understanding the physical origin of the NSSL, it currently lacks observational support.

The degree of rotational influence on convective flows is typically characterized by a nondimensional measure called the Rossby number, Ro , which is meant to express the ratio of advection to Coriolis forces in the momentum equation for the fluid. The Rossby number is usually defined such that it depends only on the rotation rate Ω and on a typical velocity U and length scale L :

$$Ro \equiv \frac{U}{\Omega L} \sim \frac{|\mathbf{v} \cdot \nabla \mathbf{v}|}{|2\boldsymbol{\Omega} \times \mathbf{v}|}. \quad (5.1)$$

A Rossby number greater than unity ($Ro > 1$) is meant to indicate that the flows do not feel a significant influence from rotation, while a Rossby number less than unity ($Ro < 1$) indicates strong rotational constraint. Practical use of the Rossby number defined this way is dependent on picking a *single* representative velocity scale U and length scale L from the flow field. While the rms velocity is commonly used for the velocity

amplitude U , the rms value depends on the motions occurring over the entire range of physical scales. This is problematic, as the convective spectrum at both the photosphere (see Hathaway et al. 2000) and at depth (Hanasoge et al. 2012; Greer et al. 2015) has significant power over a broad range of spatial scales. Further, the rms velocity obtained by any measurement scheme is resolution dependent. Since small-scale motions in the Sun generally have faster amplitudes than large-scale motions, the dominant visual scale seen in a map of the flow is often the resolution cut-off. Thus, different measurement procedures with different resolutions produce different Rossby number estimates.

In this paper I develop an alternate definition of the Rossby number that is (a) more representative of the ratio of forces presented in Equation 5.1 and (b) is a spectral representation that allows direct comparison of the contributions from each resolved spatial scale. Integrated over a band of length scales, this Rossby number serves the same function as the traditional definition and provides an easily interpreted measure of rotational influence for convective motions. Using the helioseismic analysis of HMI data, we measure horizontal flows throughout the NSSL and calculate the Rossby number associated with those flows as a function of depth. I find that the NSSL is a zone of transition for rotational influence. Immediately below the surface, the flows are fast, the Rossby number is large, and the flows lack rotational constraint. At the bottom of the NSSL, the flows have become sufficiently slow and large scale that the Rossby number is small and rotation has a strong influence on the dynamics.

In Section 5.2, I provide a synopsis of the analysis methods discussed in detail in Chapters 2 through 4. In Section 5.3, I present measurements of the horizontal flow in the near-surface solar interior in the context of convective amplitudes and horizontal scales of motion. In Section 5.4, I cast these measurements as an estimate of the Rossby number to infer the degree of rotational constraint on the observed convective motions. Finally, in Section 5.5 I compare these findings to both observations and simulations and discuss the implications for rotational constraint in the NSSL.

5.2 Methods

5.2.1 Measurement Procedure

Ring-diagram helioseismology deduces subsurface flows within the Sun through the measurement of Doppler shifts of the Sun’s acoustic wave modes. Here, we observe these waves using full-disk, line-of-sight Dopplergrams produced by the Helioseismic and Magnetic Imager (HMI) aboard the Solar Dynamics Observatory (SDO). We analyze 2048 consecutive Dopplergrams with a cadence of 45 seconds for each analysis period and collect eleven such analysis periods in all (see Table 2.1). Each analysis period has a duration of 25.6 hrs. Portions of each full-disk Dopplergram are projected onto a mosaic of smaller analysis regions called tiles. Each square tile is 194 Mm in longitude and latitude (16° in heliographic angle) and the separation between tile centers in the mosaic is 3 Mm (0.25°) in latitude and longitude. All together, each mosaic spans 90° in longitude and latitude and is comprised of 130,321 individual tiles. From image to image in the time series, the central longitude of the projection for each tile is shifted with time at the latitudinally-dependent surface differential rotation rate as measured by Snodgrass (1984). This tracking removes the large-scale velocity signal of solar differential rotation.

Any analysis of velocity measurements is limited in depth by the tile size (larger tiles sample deeper flows) and limited in horizontal resolution by the spacing between adjacent tiles. In previous studies the mosaic of tiles has typically had a much sparser spacing, with tile centers separated by half the tile size. Thus, standard ring-diagram techniques have achieved only coarse resolution (e.g., 100 Mm). Here we decrease the spacing between adjacent tiles while keeping the tile size constant. Consequently, we are able to substantially refine the resolution through deconvolution without sacrificing depth information.

The work originally published in Greer et al. (2015) relied on the central portion of a single analysis period, since that was the only data set available at the time. This original analysis region is the central 60° in each direction of one of the full analysis periods described in Table 2.1. While the measurements of the convective amplitudes presented here (and in Greer et al. 2015) are based on this small region, the results pertaining to the Rossby number use results from the full spatial extent of all the listed analysis periods.

After tracking each tile independently through the entire sequence of Dopplergrams, we apply a circular

apodization function and create a three-dimensional power spectrum (two horizontal spatial wavenumbers and temporal frequency). For each wave mode in the Sun, a subsurface horizontal flow induces a Doppler shift that is measurable in the spectrum as a frequency splitting. We use the Multi-Ridge Fitting (MRF) code (Greer et al. 2014, Chapter 3) to measure both the Doppler shift and its associated uncertainty for typically 220 unique modes of different radial order n and horizontal wavenumber k . The measured Doppler shifts exhibit a large-scale systematic bias of currently-unknown origin that depends on disk position and wave mode (Zhao et al. 2012; Greer et al. 2013; Kholikov et al. 2014). This systematic error has an amplitude of 10 to 20 m s^{-1} and contributes only to global scale flows ($\ell < 10$). Results relying on the full set of analysis regions have been corrected for this effect by subtracting the mean measured frequency shift (as measured over the eleven analysis periods) from each wave mode at each disk position. The results originally published in Greer et al. (2015) could not rely on such an average, and so relied on a different method of systematic removal. We have measured the frequency shift systematic using a large set of ring fits spanning 80 days and removed it using the procedure detailed in Greer et al. (2013) and Appendix A. This large set of ring fits was obtained using low-resolution ring-diagram analysis and interpolated on to the high-resolution data set. The longitudinal mean of each flow component is removed after this step. In each figure, I will state which data set was used so that the method of systematic removal can be inferred.

5.2.2 Inversions

The Doppler shifts measured by the MRF method are generated by horizontal flows in the solar interior, and adjacent measurements are highly correlated due to our dense tiling scheme. Each measured shift is a weighted average over the velocity field within a three-dimensional region. The weighting function for this average is called the sensitivity kernel and has a structure that depends on the wave mode used to obtain the Doppler shift (Birch et al. 2007). The broad horizontal and vertical extent of the sensitivity kernels complicate the interpretation of the Doppler shifts as a direct measure of subsurface flows, since most of the convective structures under consideration are smaller than the kernels. However, high-resolution can still be achieved by deconvolving the measurements while accounting for the high degree of correlation. We use the Massively-Overlapped Ring-Diagram Inversion (MORDI) procedure from Chapter 4 to achieve high

horizontal resolution and produce estimates of the horizontal velocity field at strongly localized points in 3-D space. This method combines every mode from all 130,321 tiles in a given analysis period and finds a balance between producing spatially isolated flow estimates and the final propagated uncertainty. The balance is tuned with regularization parameters, which we tune such that the signal-to-noise of our results is roughly constant at all depths (see Chapter 4). The covariance between any two measurements is assumed to be the product of the uncertainty measured with the MRF method for each mode times the fractional overlap area of the tiles from which the measurements are obtained.

Just as each Doppler shift measurement is related to the velocity field inside the Sun through a sensitivity kernel, each solution point in the inversion is related to the true velocity field through an averaging kernel. The purpose of the inversion is to produce averaging kernels that are much smaller and spatially localized than the sensitivity kernels. For every solution point, we compute the full 3-D averaging kernel in order to properly interpret our results. The horizontal size of the averaging kernel for a particular depth demonstrates the effective resolution of the inversion solution at that depth. The vertical structure of the averaging kernel provides an estimate of not only the vertical resolution, but the actual depth achieved for a given set of modes.

In Figure 5.1, we show cuts of three averaging kernels computed as part of the inversion. The three averaging kernels are for depths of 0.25 Mm, 5 Mm, and 30 Mm, spanning the depth range of our results. As shown in Figure 5.1a, the averaging kernels increase in horizontal size as the depth increases. Immediately below the surface, the inversion achieves a horizontal resolution of 10 Mm. The resolution scale increases with depth, degrading to 70 Mm at a depth of 30 Mm. The loss of horizontal resolution with depth is due to increasing uncertainty for Doppler shift measurements that reach deeply coupled with a dearth of such modes. At all depths, the averaging kernels demonstrate the ability for the procedure presented here to resolve flows much smaller than the tile size. Figure 5.1b shows the vertical structure of the same three kernels. These give a sense of how much each depth in the solution is correlated with any other depth. The averaging kernels for all other depths presented in this paper show a steady progression of the vertical and horizontal extents as a function of depth. A larger selection of averaging kernels can be found in Appendix B.




Figure 5.1: (a) Longitudinal cut through the peak of three different averaging kernels in the inversion, each targeting a different depth (Solid = 0.25 Mm, Dotted = 5 Mm, Dashed = 30 Mm). (b) Cuts in depth of the same averaging kernels. Both the horizontal and vertical extent of the averaging kernels increase monotonically as a function of depth.

5.3 Convective Amplitudes

Figure 5.2a shows a small region of the inversion solution at a depth of 0.25 Mm, and compares our helioseismically determined flow field to a map of the temporal average of the absolute, line-of-sight, magnetic field strength as measured by HMI over the same region. Magnetic field elements seen in photospheric observations are advected by the horizontal flows inside the Sun and tend to gather in the boundaries between neighboring convection cells. Averaged over the course of a day, the magnetic field map traces out the locations of supergranules, which persist for a comparable amount of time. Since the magnetic field map is an independent measure of the size and locations of supergranules, it serves as a useful comparison to our horizontal flow field. Figure 5.2a shows excellent agreement between the two, with collections of magnetic field sitting in the regions of converging flow. Figure 5.2b shows a larger portion of the inversion result at a depth of 10 Mm.

In Figure 5.3, we present the velocity spectrum of the inversion results at four depths (0.25 Mm, 5 Mm, 15 Mm, 30 Mm) as a function of spherical harmonic degree, and show the spectrum of the averaging kernels at each of these depths. Since the inversion flow field at a given depth is a horizontal convolution of the true flow field in the Sun with the averaging kernel at that depth (Equation 4.1), the spectrum of the averaging kernel provides the linear sensitivity of our procedure to solar flows at each harmonic degree. By construction, the averaging kernels integrate to unity, so the spectral sensitivity peaks at $\ell = 0$ at a value of 1. The sensitivity remains constant up to some value of ℓ , then drops steadily thereafter. The location in spectral space at which this transition occurs is determined by the width of the averaging kernels (smaller width, larger ℓ). We can therefore interpret the velocity spectrum at each depth as being reliable where the averaging kernel spectrum is flat, and artificially diminished for higher ℓ .

At a depth of 0.25 Mm, the velocity spectrum in Figure 5.3 shows a prominent bump between $\ell = 50$ and $\ell = 200$, consistent with the prominent supergranulation seen in Figure 5.2a. The location and amplitude of this feature is also consistent with measurements using other methods (e.g., Hathaway et al. 2000; Roudier et al. 2012). The averaging kernel spectrum at this depth confirms that we are able to resolve these scales, and that the steady drop in power above $\ell = 200$ is due to limitations in resolution as opposed to an intrinsic




Figure 5.2: (a) Vector field for a subsection of the horizontal flow map at a depth of 0.25 Mm. This map is extracted from the center of analysis period A2 (see Table 2.1) and is overlaid on a map of magnetic field strength. Darker colors indicate stronger magnetic field, and the color table saturates at 25 Gauss. The vectors have been subsampled to a resolution of 0.5° (6 Mm) to roughly match the width of the averaging kernel for this depth. The magnetic field map shown is an average of the magnetic field strength (absolute value) as measured by HMI over the same time period that is averaged over in the velocity fields. We see a strong correspondence between the horizontal flow found through our analysis and the advection of magnetic field at the photosphere. (b) Vector field for a larger subsection of the inversion solution taken at a depth of 10 Mm and subsampled to a resolution of 2° (24 Mm) to match the averaging kernel at this depth. The gray box in panel (b) indicates the position and spatial extent of the map in panel (a).

drop in velocity at smaller scales. Deeper in the Sun, we lose the ability to sample supergranular scales as the horizontal resolution degrades. At a depth of 30 Mm, we are capable of resolving only the lowest harmonic degrees ($\ell < 40$).

Figure 5.4 shows the root-mean-square (rms) horizontal velocity plotted as a function of averaging-kernel depth along with both the error on an individual solution point (shading) and the propagated error on the rms (error bars). The rms velocity peaks at 427 m s^{-1} near the surface and diminishes rapidly with depth. Within the first 10 Mm below the photosphere, the velocity drops to around 200 m s^{-1} and continues a steady but slow decline until a depth of 30 Mm. The small perturbations that appear within these deeper layers are within the error bounds and are a consequence of the specific realization of convective flows that we sample.




Figure 5.3: Spectrum of inversion results at a few depths as a function of harmonic degree. These results are from the central portion of analysis period A2 (see Table 2.1). Solid lines are the azimuthally integrated spectra of the horizontal velocity field, and color differentiates the central depth of the averaging kernel. Vertical error bars on the black and red curves indicate the 95% confidence interval, and those for the blue and green curves show similar trends. The dashed lines are the upper envelope of the averaging kernels for each depth and demonstrate the relative sensitivity of our inversion result to the true flows in the Sun. The sensitivity for every averaging kernel is unity at $\ell = 0$ by definition, but the dashed lines have been normalized to arbitrary values for visual clarity.

5.4 Rossby Number

Next, we consider the degree of rotational influence on the observed flow fields at each depth independently. We estimate this effect through the use of the Rossby number, which describes the balance of advection forces to Coriolis forces. The typical definition of the Rossby number $U/\Omega L$ requires one to select a single velocity scale U and length scale L representative of the flow field. Due to the nature of the present observations, selecting a single length scale is problematic.

At any given depth, the inversion result has a critical spatial scale. Scales larger than this critical value are well-resolved. Small scales beyond this critical value are artificially reduced in amplitude due to smoothing. The scale at which this transition occurs depends on depth, and the power spectrum of the averaging kernel for a given depth (Figure 5.3) provides a simple way of determining the critical scale. This resolution effect introduces a characteristic length scale into our results that has little to do with the true flows inside the Sun. Instead, this scale is largely determined by the error statistics of the frequency-shift




Figure 5.4: Root-mean-square horizontal velocity as a function of depth. The shaded region indicates the uncertainty on a single solution point while the vertical error bars indicate the 3-sigma values for the propagated uncertainty on the rms value. These results are from the central portion of analysis period A2 (see Table 2.1).

measurements and how we tune the inversion to prevent noise amplification. This creates a problem when using the standard definition of the Rossby number $U/\Omega L$, where it is common to pick a single length scale that is relevant to the flow field.

To avoid the issue of picking a single length scale (or a single velocity scale), we define a Rossby number that better represents the balance of forces in Equation 5.1 and is an integral measure over all resolved spatial scales:

$$\text{Ro}^2(z) \equiv \frac{1}{A} \int \int k \, dk \, d\theta \frac{|\hat{w}_x(\mathbf{k}, z)|^2 + |\hat{w}_y(\mathbf{k}, z)|^2}{\Omega^2} k^2. \quad (5.2)$$

Here, $\hat{w}_x(\mathbf{k}, z)$ and $\hat{w}_y(\mathbf{k}, z)$ are the two-dimensional Fourier transforms of the east–west and north–south components (respectively) of the inversion solution at each depth z . The constant A represents the area of the Sun over which the Rossby number is computed and the area element in Fourier space that is integrated over is written in polar coordinates $k, dk, d\theta$. This definition of the Rossby number scales with velocity and length scale in the same way as the original definition in Equation 5.1. We have defined the square of the Rossby number so that it can be viewed as a spectral density:

$$\text{Ro}^2(z) = \int dk \, \mathcal{R}^2(k, z) \quad (5.3)$$

$$\mathcal{R}^2(k, z) \equiv \frac{1}{A} \int d\theta \frac{|\hat{w}_x(\mathbf{k}, z)|^2 + |\hat{w}_y(\mathbf{k}, z)|^2}{\Omega^2} k^3. \quad (5.4)$$

This provides a simple way of understanding what spatial scales at each depth contribute to the Rossby number, which is computed as an integral of the Rossby integrand $\mathcal{R}$ over a range of spatial scales.

Figure 5.5 shows the Rossby integrand $\mathcal{R}(k, z)$ as a function of harmonic degree $\ell = kR_\odot$ for a few

depths. These results were computed by independently calculating the Rossby integrand using the full spatial extent of all of the eleven available analysis periods, then averaging the result over all periods. Near the surface, the spectrum peaks near $\ell = 100$ just as in Figure 5.3, indicating that the scales associated with supergranulation offer a significant contribution to the integrated Rossby number. Deeper into the Sun, we see that the dominant scales that contribute to the Rossby number tend to be lower ℓ than this (~ 30). The ringing seen in these curves is due to the compact nature of the inversion sensitivity kernels (see Chapter 4). It is more apparent here than in the previous spectra (Figure 5.3) due to the larger analysis region used and therefore the higher wavenumber resolution.

A useful feature of expressing the Rossby number as an integral over harmonic degree is that it becomes directly comparable to the spectrum of the inversion averaging kernel. The horizontal spectrum of the inversion averaging kernel for a given depth tells us which spatial scales are well-resolved and which suffer an artificial degradation of power. The averaging kernel spectra originally shown in Figure 5.3 are reproduced in Figure 5.5 to once again demonstrate which spatial scales are well-resolved in the inversion solution. Here, the full spectrum is shown as opposed to the upper envelope in Figure 5.3. Since the Rossby integrand depends linearly on the flow velocity, the averaging kernel spectrum reveals the degree to which the integrand is artificially diminished due to resolution effects. The total Rossby number that we infer primarily contains contributions from scales larger than these cut-offs at each depth (smaller values of ℓ).

The Rossby number Ro is computed as an integral of the Rossby integrand $\mathcal{R}$ over a chosen range of spatial scales. Figure 5.6 shows the Rossby number as defined in Equation 5.4 as a function of depth for two different ranges of spatial scales. The black curve represents an integral over all available spatial scales, ignoring the effects of the variable effective resolution with depth. This curve shows a value of around 3 at the surface, indicating flow structures that are not strongly influenced by the rotation of the Sun. Directly below the photosphere, the Rossby number integrated over all scales drops monotonically until reaching a value of 0.1 by a depth of 10 Mm. Below this depth, the Rossby number stays relatively constant until the bottom of our analysis range at 30 Mm. Since spatial resolution plays an important role in both the inversion procedure and in computing the Rossby number, Figure 5.6 also includes the Rossby number as a function of depth using only length scales $\ell < 20$, which are resolved uniformly at all depths (blue curve). This provides




Figure 5.5: Rossby integrand (Equation 5.4, solid lines) as a function of harmonic degree for various depths. The dashed lines indicate the power spectrum of the inversion averaging kernel. These are equal to unity at $\ell = 0$ by construction, but have been shifted vertically for clarity. In Figure 5.3, the upper envelope of the averaging kernel spectra were presented. Here, I show the full structure of the averaging kernel spectra, which exhibit strong ringing due to the compact nature of the inversion sensitivity kernels. The Rossby integrand shows similar oscillations in power for the same reason. These results are computed using all available analysis periods (see Table 2.1).

an estimate of the effective Rossby number relevant to large-scale convective motions. Not only does this demonstrate that the largest scales of motion are rotationally influenced at all depths, but that it is the small-scale motions ($\ell > 20$) near the surface that cause the overall Rossby number to be greater than one.

5.5 Discussion

5.5.1 Comparison of Convective Amplitudes

The results presented here show the spectrum of convective amplitudes throughout the NSSL along with an estimate of the degree of rotational constraint on these flows. A primary finding is that the speed of solar convective flows exceeds 120 m s^{-1} throughout the NSSL. This finding is in stark contradiction with the previous helioseismic study of HDS12, which presented anomalously slow velocities. Figure 5.7 compares




Figure 5.6: Rossby number as a function of depth below the photosphere obtained by integrating over all observed spatial scales (black curve) and over only the largest scales ($\ell < 20$, blue curve). At each depth, the value of the Rossby number squared has been computed independently for each of the eleven analysis periods and then averaged together. The plotted Rossby number is the square root of this average. The error bars represent the standard deviation of the Rossby number over these independent analysis periods.

the velocity spectra from Figure 5.3 for low harmonic degree (red) with the time-distance helioseismic result of HDS12 (orange). The velocity spectrum for the numerical simulation of global convection used for comparison in HDS12 is shown in green (Miesch et al. 2008). Further, the spectrum of motions from a more recent numerical simulation of the global convection zone (Featherstone and Miesch 2015) is indicated (purple). This particular simulation evinces convective motion capable of sustaining a solar-like differential rotation, possessing a pole-to-equator contrast of $\Delta\Omega/\Omega \approx 15\%$. This model was computed by solving the anelastic equations in a rotating spherical shell using the numerical algorithms described in Clune et al. (1999) and a prescription for boundary conditions and radiative heating as described in Featherstone and Miesch (2015). The simulation domain spans from the base of the convection zone at $0.72R_\odot$ to a height of $0.97R_\odot$ with a resolution of $128 \times 384 \times 768$ ($n_r \times n_\theta \times n_\phi$). The two helioseismic spectra and the newer simulation spectrum are taken at a depth of approximately 30 Mm ($0.96 R_\odot$), and the spectrum from Miesch et al. (2008) is taken at a depth of 14 Mm ($0.98 R_\odot$).

Note, the time-distance study of HDS12 does not claim to have directly detected the convective flow signal. Instead, the indicated spectrum is an upper limit that depends on a specific noise model. The drastic




Figure 5.7: Comparison of horizontal velocity spectra as a function of harmonic degree. The red curve indicates the inverted flow field (same as Figure 5.3), and error bars indicate the 95% confidence interval at each value of harmonic degree. These results are from the central portion of analysis period A2 (see Table 2.1). The orange curve is the upper limit on convective amplitudes as appears in HDS12. The purple curve is from the numerical hydrodynamic simulation in Featherstone and Miesch (2015) and is described in Section 3.4. These three spectra are taken at a depth of approximately 30 Mm ($0.96 R_{\odot}$). The green curve is from the numerical simulation in Miesch et al. (2008) at a depth of 14 Mm ($0.98 R_{\odot}$).

difference between the two helioseismically determined spectra is particularly interesting, since each set of results use full-disk Dopplergrams from the same instrument. Research comparing results from time-distance analysis to those from ring-diagram analysis have generally shown good agreement (Kosovichev et al. 2011; Hindman et al. 2003; 2004), and differences seen between the two techniques are usually far less substantial than what is shown in Figure 5.7.

One of the key steps to the analysis in HDS12 is the attempted removal of uncorrelated noise in the variance of the velocity measurements. This is accomplished based on the assumption that the relevant flows produce a constant signal as the time duration of the analysis increases, while simultaneously the level of uncorrelated noise decreases. If we were to assume that the discrepancy between the red and orange curves in Figure 5.7 is simply due to the presence of noise in our data (which we have made no attempt to remove), the signal-to-noise ratio for the results presented would be at most 0.01. However, by propagating the uncertainty measured directly from each power spectrum through the inversion procedure, we find our final signal-to-noise ratio to be around 2 (Figure 5.4). Perhaps, a more likely option is that the flow structures seen in this study evolve with time in such a way that they are removed during the noise-subtraction procedure of HDS12.

5.5.2 Rotational Influence on Convective Motions

Just as global helioseismic measurements of the subsurface differential rotation have guided numerical models of solar convection, our measurements of the convective amplitude provide useful observational constraints in the near-surface shear layer. The convective amplitude in the deepest layers that we sample (30 Mm) are particularly instructive, as these results are beginning to sample the deep flow structures responsible for the Sun’s differential rotation and global meridional circulation. Such organized, large-scale motions are the result of Reynolds stresses induced by Coriolis deflection of the convective motions. The level of rotational influence felt by the convection, characterized by the Rossby number Ro , is thus a crucial ingredient in simulations of the solar convection zone. Values of Ro greater than unity are characteristic of convection that only weakly senses the rotation, whereas values less than unity indicate rotationally constrained convection.

From the measurements presented here of the horizontal flow velocities below the surface of the Sun, we see a transition from rotationally unconstrained flows at the surface to flows that feel the influence of rotation deeper down. It is interesting that this transition not only occurs within the NSSL, but actually exhibits most of the change in the upper 10 Mm. Our measurements indicate that the depth at which advective forces roughly match Coriolis forces is at around 5 Mm.

The fact that the Rossby number at the surface is greater than unity is not surprising. Existing measurements of near-surface flows have also revealed velocities and length scales consistent with a Rossby number greater than one (300 m s^{-1} , 35 Mm; Hathaway et al. 2002). Further, the characteristic lifetimes of surface convective features (10 minutes for granules, 30 hours for supergranules) are so much shorter than the solar rotation period of 27 days that we wouldn't expect rotation to play a significant role in their evolution.

Global convection models tuned to simulate solar conditions often demonstrate rotationally influenced convective structures throughout the convection zone (e.g., Miesch et al. 2008). Paired with observations of near-surface convection, we would expect to find a transition of the Rossby number from greater than one to less than one somewhere in the convection zone. The measurements presented here suggest that this transition occurs within the NSSL, which has long been a region of interest due to its implications with angular momentum balance. This discovery of a transition from a Rossby number greater than one near the surface to less than one by the bottom of the NSSL may prove useful in understanding the origin and dynamics of the layer.

The Rossby number is a simple metric for determining the degree of rotational influence in a complex fluid system, but in a sense is a prediction of this influence and not a direct observation. A rotationally-constrained convective flow (like the deeper layers of the NSSL and the bulk of the convection zone) should exhibit correlations between the flow components due to deflections mediated by the Coriolis force. Directly measuring the Reynolds stresses for the full three-dimensional flow field in the NSSL would allow a more direct confirmation of the rotational influence on near-surface convection. The number of independent convective realizations necessary to make such a measurement is much greater than the number presented in this study, and will be the focus of future efforts.

The analysis presented here provides not only a robust measure of the horizontal convective velocities

in the solar near-surface shear layer, but also a proof-of-concept for further work in high-resolution ring-diagram helioseismology. While the computational burden of producing highly-overlapped tiles is significant, the ability to recover fine structures without sacrificing depth is of great importance. With larger tiles, it may be possible to extend this type of analysis deeper into the Sun to sample the largest scales of convection with adequate resolution.

Chapter 6

Imaging Supergranular Flows in the Near-Surface Shear Layer

6.1 Introduction

A primary goal of this thesis is to present new methods for ring-diagram helioseismology that enable a new class of observations of the solar interior. These new observations are higher resolution, lower noise, and most importantly, easier to interpret than the observations generated with previous methods. The first scientific results from these observations are presented in the previous chapter (Chapter 5). There, I provide a set of horizontal flow fields in the context of convective amplitudes and global solar dynamics. These results span the upper 30 Mm of the solar convection zone.

The use of velocity spectra and bulk characterizations like the Rossby number strip the observed flow field of phase information. That kind of analysis is easy to compare to both observations and simulations, as it does not depend on the particular convective realization being measured. However, if these flows are truly representative of the convective motions inside the Sun, the phase information carries the proof. A focus of this chapter is to analyze the observed flow fields in a way that emphasizes the convective nature of the observed subsurface flows.

6.1.1 Supergranulation

Supergranulation is the term assigned to a consistent pattern of horizontal flow seen in the photosphere. In Hart (1956), fluctuations in the photospheric velocity field were found to have a characteristic horizontal scale of around 26 Mm and with the publication of the first Doppler images of the Sun (Leighton et al. 1962) supergranulation was found to cover the entire solar disk. These velocity fluctuations form a cell-like pattern




Figure 6.1: Supergranulation as seen in direct Doppler imaging. The 32° by 32° region has been tracked over 25.6 hours of HMI Dopplergram images and then averaged in time at each projected location. See Chapter 2 for a description of the tracking process. The region shown is centered 45° from the center of the disk, causing the apparent alignment of supergranular structures. The line-of-sight Dopplergrams measure a mixture of horizontal and vertical velocity here, but supergranules have a dominant signal in the horizontal component. Red indicates a velocity towards the observer and blue indicates away. The color table saturates at $\pm 300 \text{ m s}^{-1}$.

that repeats across the solar disk, with broad areas of horizontal divergence separated by narrow lanes of convergence. This pattern is indicative of a characteristic scale of convection, and has been treated as such in most studies. Much is still unknown about supergranules six decades after their discovery, such as what sets their size, what their three-dimensional structure looks like, and how they interact with other scales of turbulent convection.

The unanswered questions regarding supergranules have been tackled by observational studies, analytic theory, and numerical simulations. The three key methods for observing supergranules are direct Doppler, photospheric tracking, and local helioseismology. Direct Doppler observations involve using the full-disk line-of-sight Dopplergrams to directly measure the photospheric flows associated with supergranules (Figure 6.1). Measurements from this technique are limited to the height of formation of the Dopplergram image, which is typically within a few hundred kilometers of the photosphere. Tracking (either with distinct features or with local correlations) involves measuring the advection of features (often granules or magnetic elements) seen in photospheric images. While advection seen in the photosphere may be influenced by subsurface flows, feature tracking does not allow one to separate out the influence of different depths within the Sun. Local helioseismology (the focus of this thesis) is the only method that permits measurements of horizontal flows at isolated depths below the photosphere.

Surface observations have provided robust measures of the horizontal scale, temporal lifetime, and velocity amplitude of supergranules. While the list of studies measuring any of these three quantities is long,

I will provide a brief summary of these numbers just to have the relevant scales in mind. A more exhaustive list can be found in Rieutord and Rincon (2015).

The horizontal size of a supergranule can range between 20 Mm and 60 Mm, with a typical size around 35 Mm (Hathaway et al. 2000; 2002; De Rosa and Toomre 2004; Del Moro et al. 2004; Hirzberger et al. 2008; Rieutord et al. 2008). The photospheric pattern of supergranulation is constantly evolving, with a coherence timescale of around 1.7 days (Hirzberger et al. 2008). As shown in De Rosa and Toomre (2004), some supergranules are recognizable even after 4 to 5 days. Under the simplifying assumption that these convective structures maintain their individual identity between being created and destroyed at discrete times, this coherence time can be treated as a supergranule lifetime. Recent measurements of supergranular velocities show typical horizontal velocities of $\sim 300 \text{ m s}^{-1}$ (Hathaway et al. 2002) and vertical velocities of $\sim 30 \text{ m s}^{-1}$ (Hathaway et al. 2002; Rieutord et al. 2010).

While these observations have provided useful measurements of the photospheric structure of supergranules, the subsurface structure remains less certain. Numerous studies have reported on the vertical extent of supergranulation using helioseismology to measure subsurface flows. These have reported depths between 1 Mm and 15 Mm, with most of them agreeing on a depth of around 7 Mm (Duvall 1998; Zhao and Kosovichev 2003; Woodard 2007; Sekii et al. 2007; Duvall et al. 2014). In these studies, the photospheric characteristics of identifiable supergranules are compared to subsurface flows to determine how far in depth the supergranules remain correlated. While some studies take the depth at which they cease to detect a supergranule-like pattern as the end-point of supergranules (Woodard 2007; Sekii et al. 2007), other studies detect a reversal of the supergranular flow that is interpreted as the return flow of a cell-like convective structure (Duvall 1998; Zhao and Kosovichev 2003).

The analysis of a possible return flow at the bottom of supergranulation implies a three-dimensional structure that contains a well-defined bottom surface. In a Rayleigh-Bénard convection, the vertical extent of convection cells is set by the presence of physical boundaries that enforce zero vertical velocity above and below. Applying this rationale to supergranulation is difficult, as there is no clear reason behind what sets their depth. Simon and Leighton (1964) suggested that He^{++} recombination could act as a driver of supergranular scales at depth, but this theory has not found support in numerical simulations. Granular-

scale simulations such as those in Stein and Nordlund (1989; 1998) have revealed an alternate picture of how the surface patterns of convection may connect to deeper motions. The intense density stratification and rapid cooling near the photosphere sets up a large entropy contrast between warm upflows and cool downflows. The downflow network seen at the surface is the source of narrow, filamentary downflow plumes that penetrate the background of broad, warm upflows. These simulations tend to show downflow plumes that merge together as they travel to form larger horizontal scales. In this picture of solar convection, we might expect that supergranules do not possess coherent return flows at a constant depth, but instead initiate a characteristic horizontal scale of downflows that morph as they permeate the convection zone.

There is an additional complication to this interpretation when comparing the lifetime of supergranulation to an estimate of the convective turnover time. One might expect that a coherent cell-like structure can only exist if the turnover time is comparable to or shorter than the lifetime of the cell. Using a 30 m s^{-1} vertical velocity and a vertical round-trip distance of 14 Mm, the estimated turnover time is around 5.5 days. This is considerably longer than the observed lifetime of 1.7 days. This mismatch of timescales is problematic. How supergranular cells can maintain a coherent return flow with a surface pattern that evolves much faster than the vertical equilibration time is a mystery.

On the theoretical side, it has been difficult to find a mechanism that can create such a distinct scale of convection amidst the high degree of turbulence expected in the solar interior. Analytic models of convection—such as Mixing Length Theory (Böhm-Vitense 1958)—predict scales of convective motion proportional to the local scale height. Since the thermodynamic properties of the solar interior vary smoothly with depth, it is difficult to create a single prominent scale of convection that stands out against the background turbulence. The highly nonlinear nature of solar magneto-convection makes numerical simulation difficult. The advancement of high-performance supercomputers has led to increasingly sophisticated numerical models (e.g., Stein et al. 2009; Ustyugov 2010; Lord et al. 2014), but a simple and consistent explanation of supergranulation remains elusive.

With measurements of subsurface flows, we may attempt to further probe the character of supergranules in this context. The technique of high-resolution helioseismology employed in this thesis is capable of resolving supergranular flows near the surface (see Chapter 5) and can produce flow fields at isolated depths

from the photosphere to a depth of 30 Mm. Applying our helioseismic procedure to successive analysis periods allows us to consider both the spatial and temporal evolution of these structures. In this chapter, I present results that link the photospheric pattern of supergranulation to convective motions throughout the upper 30 Mm and suggest a model of surface-driven convection.

6.1.2 Near-Surface Shear Layer

The upper boundary layer of the solar convection zone is a region of strong rotational shear known as the near-surface shear layer (NSSL). The region spans the upper $0.05 R_{\odot}$ (35 Mm) and exhibits a 3% decrease in the global rotation rate compared to the interior. While global helioseismology has provided precise measurements of the rotation rate in this region (Howe et al. 2000; Thompson et al. 2003), the physical origin of the layer is unknown. It has been proposed that the distinction between this layer and the bulk of the convection zone is that convective flows in the upper 35 Mm are fast enough to traverse the layer without considerable deflection by Coriolis forces (Gilman 1977; Hathaway 1982; Aurnou et al. 2007), unlike flows that exist deeper in the Sun. While this theory may play a part in understanding the physical origin of the NSSL, it currently lacks observational support.

The data set presented here spans the upper 30 Mm of the convection zone, allowing us to consider the structure of convection through the majority of the NSSL. In the previous chapter, I characterized the degree of rotational influence on convective motions through the Rossby number. Using only the bulk horizontal convective velocities and characteristic length scales, I determined that the NSSL is a transition region for the Rossby number. This suggests that the interplay between convection and rotation may be an important ingredient in maintaining the NSSL. However, using a Rossby number computed independently at each depth ignores the possibility of rotational influence on larger-scale convective structures that span multiple depths. The use of multiple consecutive analysis periods permits us to consider the evolution timescale of convective flows compared to the rotation rate. This allows a more direct measure of the degree of rotational influence on convection.

The vertical flows themselves likely play a key role in the influence of rotation on convective flows and the transport of angular momentum in the NSSL. The spatial scales at which these flows operate, along

with their radial speed, can help constrain global dynamics models and lead to a greater understanding of the NSSL. Traditionally, local helioseismology is insensitive to the vertical component of subsurface flows. While there are modifications that can be made to isolate the vertical component from observations, these do not produce results with the same fidelity as horizontal flows. Alternatively, the anelastic approximation can be applied to the observed horizontal flows to create an inferred vertical velocity field, but the validity of these flows relies heavily on the quality of the observations. In this chapter, I will present inferred vertical flows derived from the observations made possible with this new implementation of ring-diagram analysis.

In Section 6.2, I demonstrate how some aspects of the measured flow field match our assumptions about magnetic fields and bulk rotation, and thus lend credence to the observations. In Section 6.3, I show how the pattern of convection at any depth within the observed domain can be traced backward in time to the surface pattern of supergranulation. This is a new discovery made possible with the methods presented in this thesis. In Section 6.4, I will discuss the implications of this finding for both theories of solar convection and other helioseismic observations.

6.2 Matching to Independent Observations

Figure 6.2 shows vector maps of the horizontal velocity recovered from the inversions. The colors underlying the vector field indicate the value of the horizontal divergence. The vectors have been subsampled to the resolution of the averaging kernel at each depth, so each vector is nearly independent of its neighbors. These independent points reveal maps at all depths that contain spatially coherent structures. The vector field near the surface seems to be comprised of cell-like structures with sizes comparable to those of supergranules, lending some authority to the pattern. However, the flows recovered deeper into the Sun cannot be independently confirmed as easily.

The lack of independent observations for deeper flows (5 – 30 Mm depth) is due to two issues. The first is that non-helioseismic observations of the Sun are relegated to the photosphere and above. While some well-known surface phenomena are thought to be rooted in the deep solar interior (i.e., sunspots), there are currently no observations that extend our knowledge of these features any deeper than the photosphere. The second issue is that the methods used in this thesis are in practice able to measure flows with much




Figure 6.2: Horizontal flow field and horizontal divergence maps for a few depths of analysis period A4 (see Table 2.1). The vector field at each depth has been sub-sampled to the horizontal resolution of the averaging kernels. For the divergence map, red indicates positive divergence, blue is negative, and white is zero. The color map has been scaled individually at each depth.

higher horizontal resolution than most other helioseismic procedures. While other helioseismic methods are able to measure flows throughout the entire convection zone, few have the horizontal resolution to see the scales relevant for convection. The observations of deeper flows presented here are currently lacking any possibility for independent confirmation; therefore, the best chance for verifying deeper flows is to look for self-consistency between the flows at every depth.

Proper interpretation of these observations and comparison to other observations depends on understanding the resolution with which they are made. Since I have used a fully three-dimensional inversion procedure to produce the flow fields, both the horizontal and vertical resolutions are of interest. As seen previously, the resolution of the flows at any given depth are determined by the size of the inversion averaging kernel (see Appendix B for a broad selection of averaging kernels). Figure 6.3a shows the horizontal FWHM of the inversion averaging kernels as a function of target depth. The horizontal resolution begins at around 1° (10 Mm) at the surface and degrades steadily with depth. By a depth of 30 Mm, the horizontal resolution has dropped to around 5.5° (70 Mm). Figure 6.3b demonstrates that the vertical resolution (as measured with the vertical FWHM) degrades nearly linearly with depth. At any given depth, the averaging kernel has considerable sensitivity to $\pm 1/3$ of the central depth.




Figure 6.3: (a) Horizontal FWHM of the inversion averaging kernels as a function of depth. (b) Vertical FWHM of the same averaging kernels as a function of depth.

The analysis periods used in the analysis of this chapter are listed in Table 2.1. A key attribute of these regions is that they each span a 90° by 90° square centered on disk center. The Sun's differential rotation will cause convective patterns to rotate at different rates depending on latitude. As opposed to

striving to map convective flows across the entire solar disk, limiting the latitudinal span of the regions to $\pm 45^\circ$ reduces the influence of differential rotation and allows for easier analysis across multiple analysis periods. Further, each analysis period is taken from a time of relatively low magnetic activity. This is done both to minimize any systematic effects in the frequency measurements having to do with magnetic field (Hindman et al. 2000; Howe et al. 2005; Featherstone 2011) and to avoid the effects that strong localized magnetic fields may have on the flows themselves (Hindman et al. 2009; Featherstone 2011).

6.2.1 Direct Doppler Imaging

The flows observed near the photosphere can be compared to the line-of-sight, Doppler-velocity field measured by HMI. Figure 6.4 shows the photospheric line-of-sight velocity using (a) direct Doppler imaging and (b) ring-diagram analysis for the same region of the solar disk. The Doppler image has been integrated in time over the same analysis period used for the helioseismic flow map and convolved with the inversion averaging kernel to approximate the lower horizontal resolution. The helioseismic vector field has been projected onto the line-of-sight vector to mimic the signal seen in the Dopplergrams. The two maps show very similar flows structures despite the vast differences in observational methods. Direct Doppler imaging of solar flows is limited in depth and projected velocity component, but is an extremely robust method. Some of the subtle differences between the two flow maps could be due to the unmeasured vertical flows, which are roughly an order of magnitude weaker than the horizontal flows (Hathaway et al. 2002). In all, the strong similarity between the observed flows using the two methods allows us to put faith in the inversion results for near-surface flows.

6.2.2 Advection of Magnetic Field

One of the more obvious non-helioseismic measures of near-surface convection is the horizontal advection of magnetic field. Line-of-sight magnetograms provide a measure of the magnetic field strength at each point on the surface of the Sun. Just like with Dopplergrams, HMI captures full-disk magnetograms with a cadence of 45 seconds. In these, small magnetic elements are observed to be advected by horizontal flows; particularly those of supergranules (Parker 1963; Galloway et al. 1977; Galloway and Weiss 1981). Tracking




Figure 6.4: (a) Line-of-sight photospheric Doppler velocity as seen with the HMI instrument averaged over the same time as analysis period A2 (Table 2.1). The result has been blurred to approximate the surface resolution of the helioseismic observations presented in this chapter. The dominant velocity signal is from the horizontal flows of supergranulation. Supergranules have weak vertical flows, so the line-of-sight velocity diminishes near disk center. Due to a B-angle of around 5° during this analysis period, disk center falls to the north side of the equator. (b) Line-of-sight projection of photopheric horizontal flows obtained using high-resolution ring-diagram analysis. The two maps show supergranular flows and are strikingly similar (correlation coefficient of 0.73) to each other despite the vast difference in analysis procedure.

individual magnetic elements can be done to infer a horizontal velocity, as in Simon et al. (1988); Wang et al. (1995); Jafarzadeh et al. (2014). Due to the cell-like pattern of supergranulation, these magnetic elements are swept away from cell centers and converge into downflow lanes. The interstices of these lanes often show the highest concentration of advected magnetic elements. Thus, the average magnetic field strength as a function of disk position provides a rough measure of where the horizontal flow field has negative divergence. Comparing average magnetic field strength to horizontal flow fields has been done successfully for near-surface convection using helioseismology (e.g., Langfellner et al. 2015a) and feature tracking (Orozco Suárez et al. 2012).

Figure 6.5a shows the photospheric horizontal divergence map beneath a contour map of the average magnetic field strength over the same area. The average magnetic field at each location has been computed by averaging the absolute magnitude of the line-of-sight magnetic field observed with HMI. The average is compiled over the same 25.6 hours of observations over which the helioseismic flows were collected, and the region has been tracked in longitude as a function of time in the same manner as well (see Chapter

2 for a discussion on tracking). Generally, the regions of strong magnetic field tend to lie in regions of negative divergence (colored blue). This is not an exact correspondence, as the continual production of surface magnetic field (generated at the surface or advected from below) and the formation and death of supergranules precludes a steady-state pattern. Figure 6.5b shows the scatter plot density of the average magnetic field strength versus the horizontal divergence for each area element shown in Figure 6.5a. While the presence of low-level magnetic fields is nearly independent of the horizontal divergence, regions of strong magnetic field tend to coincide with regions of negative divergence. The rough correlation between magnetic field strength and horizontal divergence verifies the assertion that the observed flow fields (near the surface) are representative of convective flows.

6.2.3 Rotation Rate

The rotation rate of the solar interior has been measured with high precision using global helioseismology. The rotation rate within the NSSL varies smoothly with latitude and depth, but only varies by about 3% of the global rotation rate. We therefore expect that convective patterns near the bottom of the NSSL should rotate by nearly the same rate by which the surface patterns rotate. Simulations have shown the possibility for deep convective patterns that “swim” prograde (Miesch et al. 2000; 2008), but this is a small effect when compared to the overall rotation rate.

The separate analysis periods are spaced in time such that a feature rotating at exactly the Carrington rate (424.3 nHz) will over the period move 15° in longitude relative to disk center. This can be confirmed by correlating divergence maps measured one day apart, which are observed over the same region relative to disk center. The divergence map is used because it combines both flow components into a single scalar value that is easy to manipulate. Figure 6.6a shows the correlation coefficient as a function of relative longitudinal shift for divergence maps separated by one analysis period. The correlation is computed for maps created at the same depth, and curves are shown for a few different depths. In general, we compute the correlation coefficient for two signals as such:

$$r_{xy} = \frac{1}{N} \sum_i^N \frac{(x_i - \bar{x})(y_i - \bar{y})}{\sigma_x \sigma_y}, \quad (6.1)$$

where $\bar{x}$ and $\bar{y}$ are the means of the sets x_i and y_i respectively, σ_x and σ_y are the rms values of the signals,




Figure 6.5: (a) Contours of average line-of-sight magnetic field strength overlaid on a map of helioseismically determined horizontal divergence at a depth of 1 Mm. The contour lines are set at 10, 20, and 30 Gauss. On the divergence map, red indicates positive horizontal divergence and blue indicated negative. (b) Scatter plot density of average line-of-sight magnetic field strength versus horizontal divergence shown in panel (a). The value of the horizontal divergence has been normalized by the rms value over the region considered. The horizontal advection of magnetic fields causes larger amounts of field strength to accumulate in areas of negative divergence.

and the sum is taken over all N points that the two signals have in common.




Figure 6.6: (a) Correlation coefficient between divergence maps recorded on consecutive analysis periods as a function of longitudinal shift. A shift of 0° corresponds to correlating maps in disk-center coordinates, where the disk centers of each map are aligned. Since the separation of consecutive analysis periods is exactly 27.2753 hours, a longitudinal shift of 15° corresponds to the Carrington rotation rate. The colors indicate the depth of the divergence maps used. (b) Mean and standard deviation of correlation peak amplitude (vertical error bars) and peak longitudinal shift (horizontal error bars). This has been computed by splitting the domain into latitudinal bands and computing the correlation coefficient as a function of longitudinal shift for each independently. The peaks are located for each band at each depth, and the distributions of peak locations are represented by the vertical and horizontal error bars.

The curves in Figure 6.6a showing the correlation at different depths all show a peak in correlation somewhere near 15° . To understand the significance of the deviations from exactly 15° (as in the curve for 30 Mm), the correlations versus longitudinal shift have been computed over restricted latitudinal bands spanning the analysis region. Each band extends 10° in latitude and is used to create its own correlation plot as in Figure 6.6a. This creates a distribution of peak locations (longitudinal shift) and peak correlation values at each depth. Figure 6.6a is the correlation averaged across all bands, while Figure 6.6b shows the mean and standard deviations of the peak locations and correlation values from each of the different bands computed independently. This provides an estimate of the significance of the deviations away from exactly 15° shifts, as each latitudinal band is an independent realization. The prograde correlation at a depth of 1 Mm is the only correlation with significance, as even the peak correlation for a depth of 30 Mm is consistent with a shift of 15° . The value of the longitudinal shift that causes maximum correlation (0.4) corresponds

to a steady advection 50 m s^{-1} faster than the Carrington rotation rate. This is consistent with Duvall (1980), which presents a supergranule rotation rate 3% faster than the global rotation rate.

The center-to-limb systematic detailed in Appendix A (which manifests as an additive bias to the frequency-shift measurements) affects both horizontal velocity components and shows up as a large-scale divergence pattern. This pattern, if not removed, would create a positive correlation between the various divergence maps for a longitudinal shift of 0° . The act of removing the 11-day mean frequency shifts from each mode at each disk location (see Chapter 5) reduces this effect, but introduces a slight negative correlation at a shift of 0° . This is due to the fact that each measurement has a slight negative contribution from the measurements made at the same disk location in each other analysis period. With more averaging in the systematic removal, this slight negative correlation will decrease in amplitude.

The clarity of the correlation peaks in Figure 6.6 allows us to move forward with the analysis of these flows. The fact that flows observed at a depth of 30 Mm are seen to rotate is in and of itself an important verification. The flow maps fit the assumption that convective patterns rotate with the rest of the Sun at a rate close to the Carrington rate. With more data, it may be possible to further constrain the rate of pattern rotation as a function of depth. Since the peak at each depth for this data set is consistent with the Carrington rate to within the horizontal resolution, further correlation analysis in this chapter will assume that value for all depths in order to align different analysis periods.

6.3 Propagation of Convective Patterns

Thus far, the observed flow field near the surface has been matched against an independent measure of convective flows (Figures 6.4 and 6.5), and the deeper flow fields have been matched against themselves to confirm the assumption that the patterns persist over time and rotate with the Sun (Figure 6.6). As mentioned before, there are currently no independent measures of deep convection (that retain phase information) with which to compare, so we must look for self-consistency within the observed domain to fully justify these new observations. The fact that we are looking at a convecting system implies that there is a causal relationship between different depths. Convection dictates that some plasma must travel radially inward or outward through the Sun, transferring entropy through different depths. The vertical flows asso-

ciated with this transport imprint themselves in the horizontal divergence pattern at a given depth as the surrounding medium is either pushed aside or entrained along with them.

Near the surface, we can use the observed horizontal divergence to estimate the magnitude of the vertical velocities that carry this information. Assuming the velocities remain much slower than the speed of sound, we can use the anelastic approximation to relate the two:

$$\frac{\partial}{\partial z}(\rho v_z) = -\rho \nabla_h \cdot \mathbf{v}_h. \quad (6.2)$$

Here, the subscript h indicates a vector or operator containing only the horizontal components. The use of this approximation assumes we have captured the relevant horizontal scales that set up the dominant vertical velocities. The majority of the divergence signal near the surface is due to supergranulation, which is the predominant convective pattern observed here. The profile for ρ as a function of depth is taken from model S (Christensen-Dalsgaard et al. 1996). Using an upper boundary condition of $v_z = 0$, we can integrate Equation 6.2 through the extent of the NSSL to get an estimate of the vertical velocity field

Through this, we find that the vertical velocities in the upper few megameters have a roughly Gaussian distribution with an rms of 30 m s^{-1} (Figure 6.8). This is consistent with estimates of the photospheric vertical velocity using direct Doppler measurements (Hathaway et al. 2002). If we assume that the vertical flows responsible for convective transport of energy have this characteristic speed, we find that they can only propagate 3 Mm vertically over the duration of the analysis period. Under this assumption, we would not expect a single analysis period to show coherent structures throughout the entire 30 Mm depth domain. Instead, we should find convective patterns that have limited vertical coherence within a single period with a coherence length set by the vertical velocities. Since we have found that the pattern of horizontal divergence at each depth individually persist across analysis periods, we should also expect that the vertically-propagating signals should imprint themselves across different days as they travel through the NSSL.

While the patterns seen in the vertical velocity field are enticing, the inferred vertical velocity is not the optimal choice when looking for correlations between different depths. Through the anelastic approximation, the vertical velocity field at any given depth contains contributions from each shallower depth. While the physical assumption behind the use of the anelastic approximation is justified, it does not account for




Figure 6.7: (a,b,c) Cuts of the inferred vertical velocity field at the three depths 1 Mm, 15 Mm, and 30 Mm. Each cut has been normalized by the rms vertical velocity at each respective depth. These rms values are 16 m s^{-1} (a), 6 m s^{-1} (b), and 5 m s^{-1} (c). (d) Cut in the longitude-depth plane of the vertical velocity. Just as in (a,b,c), the velocities at each depth have been normalized by the rms vertical velocity at that depth. The dashed lines in (a,b,c) indicate the latitude (0°) of the cut in (d).

uncertainty in the measurements. Any features in the observed flow field that are not strictly due to true subsurface flows (such as random errors) will imprint themselves over many depths. As such, I will use the horizontal divergence when correlating measurements made at different depths. Not only does this allow for easier comparison to the previous results looking at correlations as a function of longitudinal shift, but the divergence maps made at different depths are as independent as their averaging kernels allow. The inversion has been tuned to create isolated measures of the horizontal velocity at each depth, so this is the best choice given the data at hand.

To look for coherent patterns introduced by vertically-propagating flows, I have computed the corre-




Figure 6.8: Distribution of vertical velocities inferred through the anelastic approximation for a few depths. Generally, the distributions do not show a significant bias towards positive or negative, and are roughly Gaussian in shape.

lation of the horizontal divergence as a function of both depth and time. The divergence map at a selected target depth is correlated with the divergence maps within the same analysis period (with no longitudinal shift) and with the divergence maps of other distinct analysis periods (with a longitudinal shift of 15° per day of time lag). The computed correlation is then averaged over each possible realization pair. The top row of Figure 6.9 shows this correlation for four selected target depths. The maximum time lag shown is determined by the longitudinal span of the regions (58° post-inversion) and the longitudinal shift per day (15°). Analysis periods separated by more than three days have no common longitudes over which to compute the correlation, as this corresponds to a 60° relative shift between analysis regions.

The plots in Figure 6.9 provide an abundance of information about the horizontal divergence patterns, so it is useful to first pick out simple features. For a target depth of 1 Mm (top left panel), we can consider the value of the correlation near the surface simply as a function of time lag (Figure 6.10a). Knowing that the surface divergence pattern is dominated by supergranules, this cut provides the time over which supergranules are coherent at the surface. Figure 6.9 shows that this correlation lasts around 2 days, which is consistent with previous studies of supergranule lifetimes (Hirzberger et al. 2008; De Rosa and Toomre 2004). Outside of this 2-day coherence, we see a divergence pattern that is slightly anti-correlated. This is consistent with the idea that the surface pattern of supergranulation evolves through new downflow lanes




Figure 6.9: (a) Correlation coefficient between divergence maps as a function of time lag and depth. For each target depth in the four panels, this shows the correlation coefficient when the target depth is correlated with each other depth over a span of time lags. The correlation shows a downward propagation of the divergence pattern around each target depth. Since the maximum time lag is limited to 3 days, the target depth is shifted deeper in order to capture the full correlation result. A longitudinal shift of 15° per day separation has been included at all depths. (b) Similar to (a), but using a structure correlation as defined in Equation 6.3. This shows a slightly clearer correlation band. The crosses indicate the mean depth of positive correlation (≥ 0.05) at each time lag value. No cross is drawn if there is no positive correlation at any depth for a given time lag.

being established that cleave existing upflow basins (De Rosa and Toomre 2004). This process tends to replace positive divergence with negative divergence, thus creating an anti-correlation at larger time lags.

Taking the value of the correlation from a target depth of 1 Mm as a function of depth with no time lag, we see coherence from the surface down to a depth of around 7 Mm (Figure 6.10b). In this figure, we can see how the vertical extent of the averaging kernels (shown as horizontal error bars) influences the




Figure 6.10: One-dimensional cuts of the 1 Mm target depth correlation in Figure 6.9a. (a) Divergence correlation for a depth of 1 Mm as a function of time separation. The pattern of supergranulation observed at the surface has a lifetime of around 2 days (b) Divergence correlation for a target depth of 1 Mm at zero time lag as a function of depth. The horizontal error bars indicate the FWHM of the averaging kernels at the depths shown. The supergranular signal persists to a depth of around 7 Mm in this cut.

result. Measurements made in different analysis periods are completely independent from one another, but measurements made at different depths within the same period may not be fully independent. The inversion averaging kernels can be used to determine which depths are correlated with each other by looking at the vertical extent at each depth (Figure 6.3b). The vertical extents of the averaging kernels at these depths are small (about half the central depth); thus the divergence maps at 1 Mm and 7 Mm are largely independent observations. This measure of the direct vertical extent of the supergranulation signal is consistent with those of Duvall (1998); Duvall et al. (2014), which also look at the pattern within a single analysis period.

The benefit of looking at the divergence correlation as a function of both time and space simultaneously is that one can look for patterns that propagate. For a target depth of 1 Mm we find that for positive time lags the correlation is maximized at a depth greater than for zero time lag. After a time lag of two analysis periods, the divergence pattern from a depth of 1 Mm correlates best at a depth of 8 Mm, and, in fact, shows a negative correlation closer to the surface. This indicates a downward propagation of the divergence pattern. When considering only a single analysis period for the correlations (zero time lag), we find that the patterns below 7 Mm do not appear to be coherent with the surface pattern of supergranulation. However,

when we allow for a finite propagation time, we find that these deeper patterns are in fact coherent to some degree with the surface supergranulation pattern from the target analysis period. It is interesting to note that no upward-propagating signal is apparent in this set of measurements.

Due to the limitations of the analysis domain, the divergence pattern from any target depth can only be correlated to ± 3 analysis periods. With this limitation, we are unable to trace the correlation of a single target depth out until the positive correlation band disappears. Thus we cannot determine the exact vertical extent of the correlation induced by photospheric motions. However, we can continue tracking the downward-propagating signal by resetting the target depth to deeper in the NSSL. In this way, we consider the convective structures that correlate weakly with the original pattern at 1 Mm and how they correlate with structures deeper down. The right three panels of the top row of Figure 6.9 show the same measure of divergence correlation for three deeper target depths. In each, we find the same trend of downward-propagating signals.

Each target depth reveals a range of other depths that are positively correlated within the 6-day span of time lags. Across the four chosen target depths, we can trace a seamless connection between the patterns seen at the photosphere and the patterns seen at the bottom of the NSSL through the downward-propagating correlation band. While we cannot claim that the divergence signal at the photosphere has a direct correlation with the pattern at 30 Mm, we can infer a directed causal relationship between the two. The flows at every depth are influenced by flows above them and influence the flows below them.

6.3.1 Structure Similarity

Tracing the correlation of the divergence through time and depth is informative, but it begs the question of what part of the flow field is actually carrying the correlation to deeper depths as time goes forward. We find that while divergence amplitudes may change from one day to the next, the sizes and locations of divergence features tend to remain constant. Thus, we use a modified measure of correlation that emphasizes the spatial patterns seen in the data, as opposed to the direct amplitudes at each point.

The correlation coefficient from Equation 6.1 describes the degree of linear dependence between the two divergence patterns being considered. This quantity includes all points in the inversion solution regardless of their amplitude. At each depth in the inversion solution, the estimated uncertainty has a value that is

generally half the rms value of the flow field velocity. This implies that features in the divergence pattern that have an amplitude less than half the rms amplitude might not be believable. To make the correlation measure between divergence maps more robust to uncertainty, I define the modified correlation coefficient:

$$r'_{xy} = \frac{1}{N} \sum_i^N x'_i y'_i, \quad (6.3)$$

where the modified divergence signal x' is defined as

$$x'_i = \begin{cases} 1 & \text{for } x_i - \bar{x} \geq \sigma_x \\ -1 & \text{for } x_i - \bar{x} \leq -\sigma_x \\ 0 & \text{otherwise.} \end{cases} \quad (6.4)$$

As this quantity acts like a standard correlation but describes the similarity between structures, I call it the structure similarity coefficient.




Figure 6.11: (a) Horizontal divergence at a depth of 1 Mm from an inversion solution. Blue indicates positive divergence, red is negative. (b) Horizontal divergence from the same depth and same region of the Sun observed one analysis period later than (a). Positive and negative thresholds have been applied to each map at a magnitude equal to the rms of the original maps. (c) Multiplication of (a) and (b) showing correlations (green), anticorrelations (orange), and no correlation (white). The strongest divergence features tend to show either a positive correlation or no correlation day-to-day, and rarely show an anti-correlation.

Figure 6.11 shows an example of the structure similarity for two near-surface divergence maps. Panels (a) and (b) show the structure maps, where a threshold has been applied to each divergence field based on the standard deviation. Panel (c) shows the multiplication of these two co-aligned maps. The obvious preference for positive structure similarity (green) over negative similarity (orange) indicates that these two

divergence maps contain many similar features. While it is no surprise that the divergence pattern at 1 Mm correlates well day-to-day, Figure 6.12 shows a similar degree of structure correlation for inversion results separated by two days and 6 Mm in depth.

Returning to Figure 6.9, the bottom row shows the structure similarity coefficient mapped in the same way as the standard correlation in the upper row. These panels show much the same patterns as the upper row, but the value of the modified correlation is higher, leading to a cleaner measurement. This is to be expected, as the structure similarity coefficient is more robust to random errors. The data once again shows a clear downward propagation of patterns from the surface through the NSSL.




Figure 6.12: (a) Horizontal divergence at a depth of 7 Mm from the inversion solution of analysis period A1. Blue indicates positive divergence, red is negative. (b) Horizontal divergence at a depth of 13 Mm from analysis period A3 (two days later than the map shown in (a)). The longitudinal span of (a) and (b) have been picked to show the region of the Sun that is common between the two analysis periods. Just as in Figure 6.11, positive and negative thresholds have been applied to each map at a magnitude equal to the rms of the original maps. (c) Multiplication of (a) and (b) showing correlations, anticorrelations, and no correlation with the same color coding as in Figure 6.11. The strongest divergence features tend to show either a positive correlation or no correlation day-to-day, and rarely show an anti-correlation.

6.3.2 Inferred Vertical Velocities

Since the correlation band in Figure 6.9 clearly indicates a downward propagation of information, we now wish to measure the speed of this propagation. Allowing the speed to change with depth, we define it as $v_c(r)$ (the correlation velocity). The crosses in the bottom row of Figure 6.9 are a measure of the central depth of correlation for a given time lag. The depth of each cross is computed by taking the correlation-

weighted mean depth at each time lag, ignoring correlation values less than 0.05. This provides a reasonable estimate of the depth at which the correlation is maximized.

For each target depth, we fit a line with an unknown slope (change in depth per time lag) and offset to the crosses. To each cross, we assign an uncertainty equal to the vertical extent of the averaging kernel at that depth. This allows us to estimate both the correlation velocity v_c and its uncertainty. Repeating this analysis for each target depth between the photosphere and 30 Mm, we can measure the correlation velocity as a function of depth.

Figure 6.13 shows the measured correlation velocity (black line) as well as a few guides for the vertical velocity inferred using the anelastic approximation. Since the anelastic approximation provides a distribution of vertical velocities instead of a single characteristic velocity, the colored lines in Figure 6.13a are the bounds of the distribution that contain 50%, 85%, and 99% of the velocities. As we are interested in downward-propagating information, the negative bounds are plotted.




Figure 6.13: Vertical velocities as determined by the anelastic approximation (colored lines) and the divergence correlation in Figure 6.9 (black line). The anelastic approximation provides a distribution of velocities that are both positive and negative (see Figure 6.8), whereas the inferred velocity from correlation is negative. Thus, the colored lines indicate the negative bounds of the anelastic velocity distribution that contain 50%, 85%, and 99% of the distribution. The correlation velocity is measured by fitting a straight line through the crosses in Figure 6.9b at each target depth. The error bars on the crosses are assumed to be the vertical FWHM of the averaging kernel at each depth, and the error bars in this figure are a consequence of those.

It is important to note that these two methods of estimating the vertical velocities are nearly indepen-

dent. The magnitude of the vertical velocities inferred from the anelastic approximation are entirely set by the horizontal convective speeds observed within a single analysis period. On the other hand, the correlation velocity is set by the similarity in convective structure between completely independent analysis periods. An arbitrary scaling of the horizontal velocities results in a linear response for the anelastic velocities, but no change in the correlation (since the divergence is normalized before correlation).

Both methods of estimating the vertical velocities show an increase in speed over the first few megameters beneath the photosphere. After reaching a peak downflow speed of 35 m s^{-1} at a depth of 3 or 4 Mm, the downflow speed decreases steadily as a function of depth. By a depth of 15 Mm, the downflow speed has decelerated to around 10 m s^{-1} and stays roughly constant until the bottom of the observational domain.

6.4 Discussion

6.4.1 Self-Consistent Flow Fields

Relatively little is known observationally about the convective dynamics of the solar interior compared to what can be seen in the photosphere. Comparing new observations to well-established observations is difficult, and we are often left comparing to only the most basic traits of the convection zone, alongside global measures of subsurface flows like the internal rotation rate. Despite the lack of independent observations of convection, these more global traits can still provide consistency checks. Matching the observed photospheric flow pattern to the direct Doppler measurements as in Figure 6.4 and the advection of magnetic fields as in Figure 6.3 provide sanity checks for the flows observed at the surface. The line-of-sight magnetic field is a robust measure of photospheric advection and allows us to put faith in the helioseismically determined flow pattern at the top of the observational domain.

Below the photosphere, there are no other methods that can provide an independent measure of convective flows. Thus, we must rely on self-consistency within the observed data set. Each analysis region is centered on disk center, and we expect flow structures to rotate around the Sun by $\sim 15^\circ$ per analysis period. Thus, assuming the observed flow patterns retain some degree of coherence between subsequent

analysis periods, we expect to see the same convective pattern shifted by around 15° between two analysis periods. As seen in Figure 6.6, the dominant pattern observed in consecutive analysis periods is not only one that shifts in longitude at a rate consistent with the global rotation rate, but has a lifetime long enough to be observed coherently across these periods. Global helioseismic measurements of the bulk rotation rate suggest that the structures at the bottom of the NSSL may rotate around 3% faster than the surface rate. This corresponds to around 0.5° of longitudinal shift relative to the surface rotation rate, which is below the precision of the current measurements.

One further self-consistency check could be the use of estimated uncertainties on the inverted flow field. The non-linear ring-fitting method that measures frequency shifts produces estimates of the uncertainties in each measurement. These uncertainties are propagated through the linear inversion and apply to each location in the flow field. Formally, the signal-to-noise in each pixel of the inverted flow fields is around two (at all depths, see Chapter 4). While this allows us to generally believe any of the stronger features seen in maps such as those in Figure 6.2, it requires faith in the original uncertainty estimates. As mentioned in Chapter 3, the uncertainty estimates from the fitting code are likely influenced by the choice of fitting model. This reduces the confidence that the final propagated uncertainties on the inversion solution are truly representative of the random error on a given velocity estimate. Thus, I have decided to forego the use of formal uncertainty estimates in favor of independent observations and internal consistency checks.

6.4.2 Supergranular Flows and Surface-Driven Convection

Through correlating the divergence patterns observed at different depths in the NSSL in time, I have discovered a preferred direction of influence for convection in the upper 30 Mm of the solar interior. The convective motions observed at any given depth feel the influence of convective motions from above and exert their influence on the convective motions below. While the limitations of the analysis procedure prevent me from exploring how deep into the Sun this chain of influence extends, I find that it begins near the photosphere where supergranules dominate the convective pattern. Thus, the convective dynamics of the NSSL are intimately linked to the structure and evolution of supergranules.

Despite the wealth of information about supergranular flows seen in the photosphere, the subsurface

structure of supergranulation has remained a mystery for decades. Existing helioseismic studies of the subsurface structure have largely been concerned with the vertical extent of supergranulation (Duvall 1998; Zhao and Kosovichev 2003; Rieutord et al. 2010; Duvall and Hanasoge 2013), but have not included the possibility of the temporal evolution of convective flows in the analysis. The results presented here reveal the importance of the temporal evolution of supergranular flows as they permeate the interior. The present discovery that the influence of supergranular flows can be traced throughout the NSSL relies on the repeated analysis of convective motions through time and reveals that supergranulation plays a critical role in the convective dynamics of the NSSL.

The pattern of supergranulation observed near the photosphere can be linked to the patterns seen deeper in the NSSL, allowing for a finite propagation time between the depths. One conclusion that can be made about supergranulation from Figure 6.9 is that there is no identifiable bottom of supergranules or return flow that is observed within the upper 30 Mm. Instead, we find a convective pattern that is generated at the photosphere with no observed influence from other structures and propagates downward with time. Due to the limited range of time lags available in the correlation maps, we cannot determine whether the surface pattern itself has a direct influence at a depth of 30 Mm. The divergence pattern at any given depth is strongly influenced by the signal propagating from above, but this signal likely changes slowly as it propagates. Thus the patterns found deep in the NSSL are causally dependent on the pattern of supergranulation at the surface even if they do not look similar or are significantly correlated.

An upward propagating supergranule return flow would appear in Figure 6.9 with the opposite slope of the correlation branch seen. While upflows must exist alongside the observed downflows, their absence from the correlation maps implies that they do not have significant influence on the convective patterns observed. The results presented here suggest that the upflows and downflows play different roles in the overall structure of near-surface convection. The lack of an upward propagating signal indicates that the downflows are the important driver of convective patterns in the upper 30 Mm. Supergranules form at the surface with no apparent influences and produce strong downflows that impact the layers below. This supports the theory of surface-driven convection, where the structure of solar convection is dominated by narrow, fast downflows that remain coherent over many scale heights (neglecting the possible influence of rotation), while the broad,

slow upflows merely fill the space between.

This picture of near-surface convection explains why investigations into the depth of supergranules have returned such varied results. Studies that have claimed to see the bottom of supergranulation based on the disappearance of patterns that correlates with the surface (Woodard 2007; Sekii et al. 2007) do not consider the temporal evolution of these flows. As seen in Figure 6.9, a single snapshot of convection will reveal a pattern of supergranulation that only persists a few megameters below the surface. This is due to the finite propagation velocity of the supergranular downflows, which take time to reach the deeper layers. Studying the subsurface properties of supergranules requires measurements of not only the three-dimensional structure of flows within the Sun, but also the temporal evolution.

The cell-like nature of supergranulation is only apparent at the photosphere, where the upflows turn over against a boundary and push against each other horizontally. The three-dimensional structure is likely one with tendrils of cool plasma forming a meandering and morphing network of downflows that permeate the NSSL. The extended reach of convective plumes also indicates a convective system that has more nonlocal influence than an ergodic turbulent convection model. What sets the horizontal scale and lifetime of supergranules is still unknown, but the process that creates them must operate in the upper few megameters of the Sun.

6.4.3 Rotational Influence on Convection in the NSSL

In the previous chapter, I ascertained the degree of rotational constraint on the convective flows observed within the NSSL. This estimate was made through the Rossby number, which uses the speed and spatial scales of the observed velocity field to estimate the influence of the Coriolis force on the flow structures. With this, I found that the flows near the photosphere are unlikely to be influenced significantly by rotation, while deeper flows past 10 Mm in depth probably are. While this result is intriguing, it relies on a simplified view of rotational influence. The use of the Rossby number provides a rough guideline for determining the degree to which rotation impacts convective motions, but ignores the role of large-scale, coherent structures. As seen in the results presented here, the downflow plumes generated near the photosphere create long-lived and large-scale nonlocal correlations in the NSSL that are not well-described by a single velocity or length




Figure 6.14: Propagation time from the photosphere to any depth in the NSSL using the vertical velocities in Figure 6.13. Assuming a fluid element starts at a depth of 0 Mm and travels downward into the Sun at a rate specified by the correlation velocity, this shows how long it takes to get to any given depth. The dashed line indicates the global rotation period of the Sun. The time taken for a parcel of fluid to traverse the depth of the NSSL is nearly equal to the time it takes for the Sun to effect one full rotation. This suggests a connection between the depth of the NSSL and the balance between convective motions and global rotation.

scale. The influence of rotation on these downflow plumes must be considered separately from the bulk motions previously studied.

The results presented in this chapter show evidence for downflows that penetrate the vertical extent of the NSSL. These extended structures take a few days to traverse the NSSL, and their typical vertical velocity has been measured as a function of depth (Figure 6.13). This measurement enables one to estimate the time it takes for a downflow plume to travel from the photosphere to a depth of 30 Mm. Comparing this time to the global rotation period of the Sun, we can get an idea of whether these downflows feel the influence of rotation. A coherent downflow plume will be deflected due to the Coriolis force over a timescale comparable to the rotation rate. Thus, the lifetime of a downflow plume provides an estimate of how strongly rotation may influence the trajectory of the plume. The distance that the downflow travels over this timespan indicates the depth at which rotation may play an important role in convective dynamics.

Figure 6.14 compares the average global solar rotation period of 27.28 days to the time it takes for the influence of a downflow plume initiated at the photosphere to reach any given depth. We find that these downflow plumes are able to reach a depth of 30 Mm in just a little longer than a solar rotation period. Within the upper ten or so megameters of the convection zone, we may say that the downflows do not feel a significant influence from the Coriolis force, as they are able to traverse this distance in a time far shorter than the rotation period. However, by the time they have reached the bottom of the NSSL, they have persisted long enough to feel this effect and become deflected. The intersection of these two timescales at a

depth comparable to the depth of the NSSL suggests that there is a fundamental difference in the degree of rotational influence for flows near the photosphere and flows crossing the bottom of the NSSL.

This comparison of timescales lends credence to theories of the origin of the NSSL that propose the degree of rotational influence on convective flows distinguishes the layer from the bulk of the convection zone (see Section 6.1.2). The downflow plumes generated near the photosphere penetrate a large fraction of the NSSL in just a few days, much too fast to feel any significant deflection. By the time a downflow plume has reached a depth of 30 Mm, the Coriolis force has had time to deflect the plume and introduce Reynolds stresses that play into the global balance of angular momentum in the NSSL. The results presented here are not sufficient to explain the full dynamics of the NSSL, but they do suggest a distinct difference in the degree of rotational influence on convective flows within this layer and those below it.

The results presented here reveal the dynamics of supergranular flows as they permeate the NSSL. The downflows associated with this characteristic scale of convection traverse the layer unimpeded and create a linkage between photospheric convection and deeper motions. This strong observed association between supergranules and downflow structures impacts our understanding of the physical origin of both supergranules and the rotational shear of the NSSL.

Chapter 7

Conclusions and Future Directions

7.1 Thesis Conclusions

7.1.1 Improvements to Ring-Diagram Helioseismology

The scientific discoveries presented in this thesis were made possible with the advent of a new implementation of high-resolution ring-diagram helioseismology. The key aspects of this new procedure can be classified into three main components. First, using a grid of large analysis tiles with exceedingly small displacements (thanks to the suggestion of Douglas Gough) permits the analysis of deep flows without the loss of horizontal resolution. Second, the use of the Multi-Ridge Fitting (MRF) procedure to measure frequency shifts of wave modes significantly increases the reliability and consistency of the subsurface flow measurements. Third, the use of an efficient three-dimensional inversion procedure (MORDI) that can take full advantage of the previous two improvements unlocks the full potential of high-resolution ring-diagram analysis.

The novel methods of ring-diagram helioseismology that I have presented in this thesis enable a new class of observations of subsurface flows with measurable improvements over previous results. The use of a high-efficiency combined tracking, projecting, and ring-fitting code (ATLAS) has resulted in a considerable increase in the number of independent, large tiles available for analysis. This increase in computational efficiency has been used to increase the horizontal grid spacing of frequency-shift measurements on the solar surface by a factor of 30 in both the east–west and north–south directions. The use of my MRF procedure has led to a significant increase in the accuracy and reliability of the frequency shifts measured from each tile.

While the previous ring-fitting method suffered from a high failure rate and systematic bias of the frequency shift measurements that varied over the disk, this new method has increased the fitting success rate over most of the measured wave modes to over 95% and decreased the systematic bias of frequency shift measurements to below 5%. This increase in reliability across the mode set has allowed consistent measurements of subsurface flows to twice the depth of the previous method. The three-dimensional inversion procedure takes full advantage of these improvements to produce high-resolution maps of the subsurface flows throughout the NSSL. The regularization scheme implemented in MORDI allows finer control of the averaging kernels than previously possible, resulting in inversion averaging kernels that are much more compact and inversion results that are easier to interpret.

The tools developed to achieve these improvements are publicly available with accompanying documentation. Each of the key tools used in this thesis are listed in Appendix C. Accompanying each entry is a brief description of how the code works as well as a URL where the code and associated documentation can be found.

7.1.2 Properties of Near-Surface Solar Convection

Using this improved, high-resolution implementation of ring-diagram helioseismology, I have measured the characteristics of convective flows throughout the upper 30 Mm of the solar interior. These observations permit a wide range of analyses from bulk motion properties to the detailed analysis of transient convective patterns. The speed of these convective flows are found to be significantly faster than previous estimates from local helioseismic observations that placed an upper limit of 1 m s^{-1} on the convective motions at a depth of 30 Mm. Here, I have presented convective velocities that exceed 100 m s^{-1} throughout the upper 30 Mm of the solar interior and are in general agreement to the flow speeds generated in numerical simulations of global solar convection. These measurements of fast convective flows in the NSSL along with their observed spatial scales provide constraints for models of solar convection and large-scale dynamics.

To confirm that the observed flow fields are representative of the true subsurface flow field, I have compared the measured flows to independent observations, such as direct Doppler imaging and the advection of surface magnetic field. These comparisons are enabled by the high resolution of the observational proce-

ture, as both of the independent measures predominantly reveal small-scale convective motions like those of supergranules. I have confirmed that the flow patterns found deeper in the Sun (which lack independent observations) rotate roughly at the solar rotation rate, further confirming the validity of the observations.

I have analyzed the nature of supergranular flows seen near the photosphere of the observation set and found them to be a driving influence for convective patterns throughout the entire NSSL. While the spatial extent and lifetime of supergranulation observed at the photosphere is consistent with previous studies, the depth profile is found to be significantly more complicated than previously assumed. Instead of a simple cell-like profile in depth, I have found that supergranules exhibit characteristics of surface-driven convection, where downflow plumes generated at the surface and with the particular spatial scale associated with supergranulation penetrate the NSSL and impart their pattern on each depth. These downflow plumes are found to initiate near the photosphere and continually evolve as they propagate downward with time. By correlating the observed divergence patterns across time lags of up to three days, I have found that the downflow plumes travel 10 to 30 m s⁻¹ through the NSSL. The upflows associated with supergranulation are not found to contribute significantly to the observed propagation of convective patterns. This picture of supergranulation as a surface-driven phenomena along with the properties of the associated flow structures deeper down will enable further understanding as to the origin of supergranulation and its impact on large-scale dynamics.

7.1.3 Influence of Rotation on Convection

Through two different methods, I have inferred the degree of rotational influence on convective motions in the NSSL. The first method relies on estimates of the Rossby number, which describes the balance between advective and Coriolis forces on convection. Using the observed flow field at each depth, I have determined that the NSSL is a transition region for the influence of rotation; flows near the surface are fast and small and therefore are not influenced by rotation, while the deeper, slow flows at the bottom of the NSSL likely are. The use of a spectral Rossby number has allowed us to consider the contributions of different spatial scales to the total Rossby number independently at any given depth as well as account for the finite and variable horizontal resolution of the measurement procedure.

The second method of inferring the influence of rotation is through the detection of long-lived downflow plumes that traverse the upper 30 Mm of the solar interior. Through detailed analysis of the propagation of convective patterns in time and space, I have measured the vertical propagation velocity and resulting travel time of convective patterns across the NSSL. I have found that the time it takes for surface-driven convective downflow plumes to reach the bottom of the NSSL is comparable to the rotation rate of the Sun. This measured travel-time reveals that these newly-observed downflow plumes live long enough to feel the influence of rotation and therefore may experience Coriolis deflection around a depth of 30 Mm. Both the bulk convective speeds observed throughout the NSSL as well as the propagation of downflow plumes suggest that the NSSL is a region of transition for the influence of rotation on convection. These results will play an important role in our understanding of the physical origin and dynamics of the NSSL.

7.2 Future Directions in High-Resolution Ring-Diagram Analysis

7.2.1 Spatial Extent of Analysis Regions

The pattern propagation shown in Figure 6.11 connects the divergence pattern seen at the photosphere to the patterns seen at 30 Mm. It would be very interesting to see how much farther into the convection zone this pattern can be traced. If the idea that some of the the structure of solar convection is determined by compact downflow plumes, with enough data it may be possible to trace these patterns deeper into the convection zone to see how they interact with larger global convection scales. The choice of using 16° tiles for the analysis procedure limits the depth to which we can observe flows to 30 Mm. Larger tile sizes provide frequency-shift measurements for lower wavenumber modes which penetrate deeper into the Sun. Repeating the entire analysis procedure for a tile size twice as big will likely increase the depth limit by a factor of two. This comes at a significant computational cost, as many of the analysis steps scale in time as the linear size of the tiles squared. Thus a factor of two in depth will likely require four times as much computing time to achieve.

Increasing the horizontal size of the analysis regions not only provides a larger swath of convective motions to analyze, but also allows tracking individual convective features for longer as they rotate across

the disk. The longitudinal extent of the analysis regions can be extended fairly easily, as the fitting procedure presented in Chapter 3 produces reliable frequency-shift measurements out to $\sim 70^\circ$ from disk center. This limit is largely due to foreshortening of the Dopplergram data as one approaches the solar limb. While the fitting procedure does this in both the longitudinal and latitudinal directions, increasing the latitudinal extent of the analysis regions has an additional complication. The latitudinal extent of the analysis regions ($\pm 45^\circ$) has been chosen to limit the effects of sphericity on the inversion. The three-dimensional inversion procedure assumes translational invariance of both the sensitivity kernels and the estimated uncertainties for frequency-shift measurements of each wave mode. Implicit in this assumption is the fact that adjacent measurements must have the same covariance regardless of their position on the disk. Since longitude lines converge towards the poles, measurements at adjacent longitudes have increased covariance away from the equator due to the higher overlap of their respective tiles. Accounting for this effect would break the assumptions that allow the inversion to scale to such large and high-resolution data sets. Thus, it is most advantageous to maintain the current latitudinal extent of the analysis regions and instead focus on increasing the longitudinal extent.

7.2.2 Horizontal Resolution

Currently, the inversion solution has just enough horizontal resolution to see the effects of supergranulation near the photosphere. Deeper into the NSSL, the horizontal resolution degrades significantly, limiting the range of scales available for analysis. Increasing the horizontal resolution of the results may reveal more about supergranulation and the NSSL, but the road to a higher-resolution inversion solution is less clear. Currently, the horizontal resolution of the solution is limited by the magnitude of the estimated uncertainties. The inversion is tuned to achieve an optimal trade-off between resolution and uncertainty (see Chapter 4 for details). As such, a decrease in the frequency-shift uncertainties used in the inversion could be traded for an increase in horizontal resolution at every depth. A decrease in the uncertainty of frequency-shift measurements for a constant tile size would have to result from improvements to the ring-fitting method, and is likely a difficult task. An easier solution is to use larger tile sizes, which generally provide lower-uncertainty measurements for the same wave modes measured in smaller tiles (see Chapter 3, in particular Figure 3.5). While counterintuitive, using larger ring-diagram tiles will likely allow higher-resolution observations of subsurface

flows.

While the current limit of horizontal resolution is set by the estimated uncertainties on the frequency-shift measurements, it is useful to consider other limits that may become important in the future. The first potential limit is the ability to decrease the marginal shift between adjacent ring-diagram tiles. Currently, this is set to 0.25° (3 Mm), which is around six times the pixel resolution of HMI Dopplergrams near disk center (~ 500 km). This resolution decreases towards the limb of the disk, so a full factor of six in the horizontal resolution is likely not achievable. Allowing for a significant increase in the computational cost of producing such a tile set, one could take advantage of the full resolution of the HMI Dopplergrams and enable higher resolution results.

This new high-resolution implementation of ring-diagram helioseismology has shown to be capable of directly imaging convective flows in the solar interior with unprecedented results. The resolution, depth extent, and temporal range of the data set presented here have been chosen as a trade-off between scientific usefulness and computational cost. I believe that the results presented in this chapter justify an increase in these critical quantities in order to improve upon and extend these scientific discoveries. While this thesis presents an unprecedented look at near-surface convection, it can be viewed as a collection of preliminary results in using observational means to unravel some of the mysteries of solar dynamics. The path forward will involve utilizing this new machinery and applying it to the vast number of unanswered questions in solar dynamics.

Bibliography

- E. R. Anderson, T. L. Duvall, Jr., and S. M. Jefferies. Modeling of solar oscillation power spectra. *Astrophys. J.*, 364:699–705, December 1990. doi: 10.1086/169452.
- H. M. Antia and S. Basu. Temporal Variations of the Rotation Rate in the Solar Interior. *Astrophys. J.*, 541:442–448, September 2000. doi: 10.1086/309421.
- K. Augustson, M. Rast, R. Trampedach, and J. Toomre. Modeling the Near-Surface Shear Layer: Diffusion Schemes Studied With CSS. *Journal of Physics Conference Series*, 271(1):012070, January 2011. doi: 10.1088/1742-6596/271/1/012070.
- J. Aurnou, M. Heimpel, and J. Wicht. The effects of vigorous mixing in a convective model of zonal flow on the ice giants. *Icarus*, 190:110–126, September 2007. doi: 10.1016/j.icarus.2007.02.024.
- C. S. Baldner and J. Schou. Effects of Asymmetric Flows in Solar Convection on Oscillation Modes. *Astrophys. J. Lett.*, 760:L1, November 2012. doi: 10.1088/2041-8205/760/1/L1.
- S. Basu and H. M. Antia. Solar Cycle Variations of Large-Scale Flows in the sun. *Solar Phys.*, 192:469–480, March 2000. doi: 10.1023/A:1005256817534.
- S. Basu, H. M. Antia, and S. C. Tripathy. Ring Diagram Analysis of Velocity Fields within the Solar Convection Zone. In S. Korzenik, editor, *Structure and Dynamics of the Interior of the Sun and Sun-like Stars*, volume 418 of *ESA Special Publication*, page 705, 1998.
- S. Basu, H. M. Antia, and S. C. Tripathy. Ring Diagram Analysis of Near-Surface Flows in the Sun. *Astrophys. J.*, 512:458–470, February 1999. doi: 10.1086/306765.
- J. G. Beck, L. Gizon, and T. L. Duvall, Jr. A New Component of Solar Dynamics: North-South Diverging Flows Migrating toward the Equator with an 11 Year Period. *Astrophys. J. Lett.*, 575:L47–L50, August 2002. doi: 10.1086/342636.
- A. C. Birch, A. G. Kosovichev, and T. L. Duvall, Jr. Sensitivity of Acoustic Wave Travel Times to Sound-Speed Perturbations in the Solar Interior. *Astrophys. J.*, 608:580–600, June 2004. doi: 10.1086/386361.
- A. C. Birch, L. Gizon, B. W. Hindman, and D. A. Haber. The linear sensitivity of helioseismic ring diagrams to local flows. *Astrophys. J.*, 662:730–737, June 2007. doi: 10.1086/513683.
- R. S. Bogart, C. Baldner, S. Basu, D. A. Haber, and M. C. Rabello-Soares. HMI ring diagram analysis I. The processing pipeline. *J. Phys. Conf. Ser.*, 271(1):012008, January 2011a. doi: 10.1088/1742-6596/271/1/012008.
- R. S. Bogart, C. Baldner, S. Basu, D. A. Haber, and M. C. Rabello-Soares. HMI ring diagram analysis II. Data products. *J Phys. Conf. Ser.*, 271(1):012009, January 2011b. doi: 10.1088/1742-6596/271/1/012009.
- E. Böhm-Vitense. Über die Wasserstoffkonvektionszone in Sternen verschiedener Effektivtemperaturen und Leuchtkräfte. Mit 5 Textabbildungen. *Zeitschrift fuer Astrophysik*, 46:108, 1958.

- D. C. Braun. Scattering of p-Modes by Sunspots. I. Observations. *Astrophys. J.*, 451:859, October 1995. doi: 10.1086/176272.
- D. C. Braun and C. Lindsey. Seismic Imaging of the Far Hemisphere of the Sun. *Astrophys. J. Lett.*, 560: L189–L192, October 2001. doi: 10.1086/324323.
- D. C. Braun, T. L. Duvall, Jr., and B. J. Labonte. Acoustic absorption by sunspots. *Astrophys. J. Lett.*, 319:L27–L31, August 1987. doi: 10.1086/184949.
- D. C. Braun, T. L. Duvall, Jr., B. J. Labonte, S. M. Jefferies, J. W. Harvey, and M. A. Pomerantz. Scattering of p-modes by a sunspot. *Astrophys. J. Lett.*, 391:L113–L116, June 1992. doi: 10.1086/186410.
- D. C. Braun, A. C. Birch, and C. Lindsey. Local Helioseismology of Near-Surface Flows. In D. Danesy, editor, *SOHO 14 Helio- and Asteroseismology: Towards a Golden Future*, volume 559 of *ESA Special Publication*, page 337, October 2004.
- T. M. Brown, J. Christensen-Dalsgaard, W. A. Dziembowski, P. Goode, D. O. Gough, and C. A. Morrow. Inferring the sun’s internal angular velocity from observed p-mode frequency splittings. *Astrophys. J.*, 343:526–546, August 1989. doi: 10.1086/167727.
- J. Christensen-Dalsgaard. Helioseismology. *Reviews of Modern Physics*, 74:1073–1129, November 2002. doi: 10.1103/RevModPhys.74.1073.
- J. Christensen-Dalsgaard, J. Schou, and M. J. Thompson. A comparison of methods for inverting helioseismic data. *Mon. Not. Roy. Astron. Soc.*, 242:353–369, February 1990.
- J. Christensen-Dalsgaard, W. Dappen, S. V. Ajukov, E. R. Anderson, H. M. Antia, S. Basu, and et al. The current state of solar modeling. *Science*, 272:1286–1292, May 1996. doi: 10.1126/science.272.5266.1286.
- T. C. Clune, J. R. Elliot, M. S. Miesch, J. Toomre, and G. A. Glatzmaier. Computational aspects of a code to study rotating turbulent convection in spherical shells. *Parallel Computing*, 25:361–380, April 1999. doi: 10.1016/S0167-8191(99)00009-5.
- S. Couvidat, S. P. Rajaguru, R. Wachter, K. Sankarasubramanian, J. Schou, and P. H. Scherrer. Line-of-Sight Observables Algorithms for the Helioseismic and Magnetic Imager (HMI) Instrument Tested with Interferometric Bidimensional Spectrometer (IBIS) Observations. *Solar Phys.*, 278:217–240, May 2012. doi: 10.1007/s11207-011-9927-y.
- A. D. Crouch, P. Charbonneau, and K. Thibault. Supergranulation as an Emergent Length Scale. *Astrophys. J.*, 662:715–729, June 2007. doi: 10.1086/515564.
- M. L. De Rosa and J. Toomre. Evolution of Solar Supergranulation. *Astrophys. J.*, 616:1242–1260, December 2004. doi: 10.1086/424920.
- K. DeGrave and J. Jackiewicz. Helioseismic Investigation of Modeled and Observed Supergranule Structure. *Solar Phys.*, 290:1547–1568, June 2015. doi: 10.1007/s11207-015-0693-0.
- D. Del Moro, F. Berrilli, T. L. Duvall, Jr., and A. G. Kosovichev. Dynamics and Structure of Supergranulation. *Solar Phys.*, 221:23–32, May 2004. doi: 10.1023/B:SOLA.0000033363.15641.8f.
- F.-L. Deubner and D. Gough. Helioseismology: Oscillations as a Diagnostic of the Solar Interior. *ARA&A*, 22:593–619, 1984. doi: 10.1146/annurev.aa.22.090184.003113.
- M. Dikpati and P. Charbonneau. A Babcock-Leighton Flux Transport Dynamo with Solar-like Differential Rotation. *Astrophys. J.*, 518:508–520, June 1999. doi: 10.1086/307269.
- T. L. Duvall and S. M. Hanasoge. Subsurface Supergranular Vertical Flows as Measured Using Large Distance Separations in Time-Distance Helioseismology. *Solar Phys.*, 287:71–83, October 2013. doi: 10.1007/s11207-012-0010-0.

- T. L. Duvall, S. M. Hanasoge, and S. Chakraborty. Additional Evidence Supporting a Model of Shallow, High-Speed Supergranulation. *Solar Phys.*, 289:3421–3433, September 2014. doi: 10.1007/s11207-014-0537-3.
- T. L. Duvall, Jr. The equatorial rotation rate of the supergranulation cells. *Solar Phys.*, 66:213–221, June 1980. doi: 10.1007/BF00150578.
- T. L. Duvall, Jr. Recent Results and Theoretical Advances in Local Helioseismology. In S. Korzennik, editor, *Structure and Dynamics of the Interior of the Sun and Sun-like Stars*, volume 418 of *ESA Special Publication*, page 581, 1998.
- T. L. Duvall, Jr. and A. C. Birch. The Vertical Component of the Supergranular Motion. *Astrophys. J. Lett.*, 725:L47, December 2010. doi: 10.1088/2041-8205/725/1/L47.
- T. L. Duvall, Jr. and S. M. Hanasoge. Measuring Meridional Circulation in the Sun. In M. Dikpati, T. Arentoft, I. González Hernández, C. Lindsey, and F. Hill, editors, *Solar-Stellar Dynamos as Revealed by Helio- and Asteroseismology: GONG 2008/SOHO 21*, volume 416 of *Astronomical Society of the Pacific Conference Series*, page 103, December 2009.
- T. L. Duvall, Jr., S. M. Jefferies, J. W. Harvey, and M. A. Pomerantz. Time-distance helioseismology. *Nature*, 362:430–432, April 1993. doi: 10.1038/362430a0.
- T. L. Duvall, Jr., A. G. Kosovichev, P. H. Scherrer, R. S. Bogart, R. I. Bush, C. de Forest, J. T. Hoeksema, J. Schou, J. L. R. Saba, T. D. Tarbell, A. M. Title, C. J. Wolfson, and P. N. Milford. Time-Distance Helioseismology with the MDI Instrument: Initial Results. *Solar Phys.*, 170:63–73, 1997. doi: 10.1023/A:1004907220393.
- N. A. Featherstone. *Exploring Convection and Dynamos in the Cores and Envelopes of Stars*. PhD thesis, University of Colorado at Boulder, 2011.
- N. A. Featherstone and M. S. Miesch. Meridional Circulation in Solar and Stellar Convection Zones. *Astrophys. J.*, 804:67, May 2015. doi: 10.1088/0004-637X/804/1/67.
- B. Fleck, S. Couvidat, and T. Straus. On the Formation Height of the SDO/HMI Fe 6173 Å Doppler Signal. *Solar Phys.*, 271:27–40, July 2011. doi: 10.1007/s11207-011-9783-9.
- D. J. Galloway and N. O. Weiss. Convection and magnetic fields in stars. *Astrophys. J.*, 243:945–953, February 1981. doi: 10.1086/158659.
- D. J. Galloway, M. R. E. Proctor, and N. O. Weiss. Formation of intense magnetic fields near the surface of the sun. *Nature*, 266:686–689, April 1977. doi: 10.1038/266686a0.
- P. M. Giles, T. L. Duvall, P. H. Scherrer, and R. S. Bogart. A subsurface flow of material from the Sun’s equator to its poles. *Nature*, 390:52–54, November 1997. doi: 10.1038/36294.
- P. M. Giles, T. L. Duvall, Jr., and P. H. Scherrer. Time-Distance Measurements of Subsurface Rotation and Meridional Flow. In S. Korzennik, editor, *Structure and Dynamics of the Interior of the Sun and Sun-like Stars*, volume 418 of *ESA Special Publication*, page 775, 1998.
- P. A. Gilman. Nonlinear Dynamics of Boussinesq Convection in a Deep Rotating Spherical Shell. I. *Geophysical and Astrophysical Fluid Dynamics*, 8:93–135, 1977. doi: 10.1080/03091927708240373.
- L. Gizon and A. C. Birch. Local helioseismology, 2005. URL <http://www.livingreviews.org/lrsp-2005-5>. [Online Article]: cited 13-August-2015.
- L. Gizon, T. L. Duvall, Jr., and R. M. Larsen. Probing Surface Flows and Magnetic Activity with Time-Distance Helioseismology. In P. Brekke, B. Fleck, and J. B. Gurman, editors, *Recent Insights into the Physics of the Sun and Heliosphere: Highlights from SOHO and Other Space Missions*, volume 203 of *IAU Symposium*, page 189, 2001.

- P. Goldreich and D. A. Keeley. Solar seismology. II - The stochastic excitation of the solar p-modes by turbulent convection. *Astrophys. J.*, 212:243–251, February 1977. doi: 10.1086/155043.
- I. González Hernández, J. Patrón, R. S. Bogart, and T. SOI Ring Diagram Team. Meridional Flows from Ring Diagram Analysis. *Astrophys. J. Lett.*, 510:L153–L156, January 1999. doi: 10.1086/311811.
- I. González Hernández, J. Patrón, T. Roca Cortés, R. S. Bogart, F. Hill, and E. J. Rhodes, Jr. A Synoptic View of the Subphotospheric Horizontal Velocity Flows in the Sun. *Astrophys. J.*, 535:454–463, May 2000. doi: 10.1086/308812.
- I. Gonzalez-Hernandez, C. Lindsey, and F. Hill. Far-side Helioseismic Holography: Calibrating The Signature Of Active Regions. In *American Astronomical Society Meeting Abstracts #210*, volume 39 of *Bulletin of the American Astronomical Society*, page 126, May 2007.
- I. González Hernández, R. Howe, R. Komm, and F. Hill. Meridional circulation during the extended solar minimum: Another component of the torsional oscillation? *Astrophys. J. Lett.*, 713:L16–L20, April 2010. doi: 10.1088/2041-8205/713/1/L16.
- P. R. Goode, D. Gough, and A. G. Kosovichev. Localized excitation of solar oscillations. *Astrophys. J.*, 387:707–711, March 1992. doi: 10.1086/171118.
- D. Gough and J. Toomre. Seismic observations of the solar interior. *ARA&A*, 29:627–685, 1991. doi: 10.1146/annurev.aa.29.090191.003211.
- B. Greer, B. Hindman, and J. Toomre. Center-to-Limb Velocity Systematic in Ring-Diagram Analysis. In K. Jain, S. C. Tripathy, F. Hill, J. W. Leibacher, and A. A. Pevtsov, editors, *Fifty Years of Seismology of the Sun and Stars*, volume 478 of *Astronomical Society of the Pacific Conference Series*, page 199, December 2013.
- B. J. Greer, B. W. Hindman, and J. Toomre. Multi-Ridge Fitting for Ring-Diagram Helioseismology. *Solar Phys.*, 289:2823–2843, August 2014. doi: 10.1007/s11207-014-0514-x.
- B. J. Greer, B. W. Hindman, N. A. Featherstone, and J. Toomre. Helioseismic Imaging of Fast Convective Flows throughout the Near-surface Shear Layer. *Astrophys. J. Lett.*, 803:L17, April 2015. doi: 10.1088/2041-8205/803/2/L17.
- D. A. Haber. Subphotospheric Structure and Dynamics Deduced from Ring Analysis. In D. Danesy, editor, *SOHO 14 Helio- and Asteroseismology: Towards a Golden Future*, volume 559 of *ESA Special Publication*, pages 676–685, October 2004.
- D. A. Haber, B. W. Hindman, J. Toomre, R. S. Bogart, M. J. Thompson, and F. Hill. Solar shear flows deduced from helioseismic dense-pack samplings of ring diagrams. *Solar Phys.*, 192:335–350, March 2000. doi: 10.1023/A:1005235610132.
- D. A. Haber, B. W. Hindman, J. Toomre, R. S. Bogart, and F. Hill. Daily variations of large-scale subsurface flows and global synoptic flow maps from dense-pack ring-diagram analyses. In A. Wilson and P. L. Pallé, editors, *SOHO 10/GONG 2000 Workshop: Helio- and Asteroseismology at the Dawn of the Millennium*, volume 464 of *ESA Special Publication*, pages 209–212, January 2001.
- D. A. Haber, B. W. Hindman, J. Toomre, R. S. Bogart, R. M. Larsen, and F. Hill. Evolving submerged meridional circulation cells within the upper convection zone revealed by ring-diagram analysis. *Astrophys. J.*, 570:855–864, May 2002. doi: 10.1086/339631.
- S. M. Hanasoge, T. L. Duvall, and K. R. Sreenivasan. Anomalously weak solar convection. *Proceedings of the National Academy of Science*, 109:11928–11932, July 2012. doi: 10.1073/pnas.1206570109.
- A. B. Hart. Motions in the Sun at the photospheric level. VI. Large-scale motions in the equatorial region. *Mon. Not. Roy. Astron. Soc.*, 116:38, 1956.

- J. Harvey. High-resolution helioseismology. In E. Rolfe and B. Battrock, editors, Future missions in solar, heliospheric & space plasma physics, volume 235 of ESA Special Publication, pages 199–208, June 1985.
- D. H. Hathaway. Nonlinear simulations of solar rotation effects in supergranules. *Solar Phys.*, 77:341–356, April 1982. doi: 10.1007/BF00156116.
- D. H. Hathaway. Doppler Measurements of the Sun’s Meridional Flow. *Astrophys. J.*, 460:1027, April 1996. doi: 10.1086/177029.
- D. H. Hathaway. The Sun’s Shallow Meridional Circulation. ArXiv e-prints, March 2011.
- D. H. Hathaway. The solar cycle, 2010. URL <http://www.livingreviews.org/lrsp-2010-1>. [Online Article]: cited 13-August-2015.
- D. H. Hathaway, J. G. Beck, R. S. Bogart, K. T. Bachmann, G. Khatri, J. M. Petitto, S. Han, and J. Raymond. The Photospheric Convection Spectrum. *Solar Phys.*, 193:299–312, April 2000. doi: 10.1023/A:1005200809766.
- D. H. Hathaway, J. G. Beck, S. Han, and J. Raymond. Radial Flows in Supergranules. *Solar Phys.*, 205: 25–38, January 2002. doi: 10.1023/A:1013881213279.
- F. Hill. Rings and trumpets - Three-dimensional power spectra of solar oscillations. *Astrophys. J.*, 333: 996–1013, October 1988. doi: 10.1086/166807.
- B. Hindman, D. Haber, J. Toomre, and R. Bogart. Local Fractional Frequency Shifts Used as Tracers of Magnetic Activity. *Solar Phys.*, 192:363–372, March 2000. doi: 10.1023/A:1005283302728.
- B. Hindman, L. Gizon, D. Haber, T. Duval, Jr., and J. Toomre. Comparison of near-surface flows assessed by ring-diagram and f-mode time-distance analyses. In H. Sawaya-Lacoste, editor, GONG+ 2002. Local and Global Helioseismology: the Present and Future, volume 517 of ESA Special Publication, pages 299–302, February 2003.
- B. W. Hindman, L. Gizon, T. L. Duvall, Jr., D. A. Haber, and J. Toomre. Comparison of Solar Subsurface Flows Assessed by Ring and Time-Distance Analyses. *Astrophys. J.*, 613:1253–1262, October 2004. doi: 10.1086/423263.
- B. W. Hindman, D. Gough, M. J. Thompson, and J. Toomre. Helioseismic ring analyses of artificial data computed for two-dimensional shearing flows. *Astrophys. J.*, 621:512–523, March 2005. doi: 10.1086/427400.
- B. W. Hindman, D. A. Haber, and J. Toomre. Helioseismically Determined Near-Surface Flows Underlying a Quiescent Filament. *Astrophys. J.*, 653:725–732, December 2006. doi: 10.1086/508603.
- B. W. Hindman, D. A. Haber, and J. Toomre. Subsurface Circulations within Active Regions. *Astrophys. J.*, 698:1749–1760, June 2009. doi: 10.1088/0004-637X/698/2/1749.
- J. Hirzberger, L. Gizon, S. K. Solanki, and T. L. Duvall. Structure and Evolution of Supergranulation from Local Helioseismology. *Solar Phys.*, 251:417–437, September 2008. doi: 10.1007/s11207-008-9206-8.
- H. Hotta, M. Rempel, and T. Yokoyama. High-resolution calculation of the solar global convection with the reduced speed of sound technique: II. Near surface shear layer with the rotation. ArXiv e-prints, October 2014.
- R. Howard and B. J. Labonte. The sun is observed to be a torsional oscillator with a period of 11 years. *Astrophys. J. Lett.*, 239:L33–L36, July 1980. doi: 10.1086/183286.
- R. Howe. Solar interior rotation and its variation, 2009. URL <http://www.livingreviews.org/lrsp-2009-1>. [Online Article]: cited 13-August-2015.

- R. Howe, J. Christensen-Dalsgaard, F. Hill, R. W. Komm, R. M. Larsen, J. Schou, M. J. Thompson, and J. Toomre. Dynamic Variations at the Base of the Solar Convection Zone. *Science*, 287:2456–2460, March 2000. doi: 10.1126/science.287.5462.2456.
- R. Howe, R. W. Komm, I. Gonzalez Hernandez, F. Hill, D. A. Haber, and B. W. Hindman. Local Helioseismic Mode Frequency Shifts With Magnetic Activity, From GONG and MDI. *AGU Spring Meeting Abstracts*, page B6, May 2005.
- R. Howe, R. Komm, F. Hill, R. Ulrich, D. A. Haber, B. W. Hindman, J. Schou, and M. J. Thompson. Large-Scale Zonal Flows Near the Solar Surface. *Solar Phys.*, 235:1–15, May 2006. doi: 10.1007/s11207-006-0117-2.
- J. Jackiewicz, A. C. Birch, L. Gizon, S. M. Hanasoge, T. Hohage, J.-B. Ruffio, and M. Švanda. Multichannel Three-Dimensional SOLA Inversion for Local Helioseismology. *Solar Phys.*, 276:19–33, February 2012. doi: 10.1007/s11207-011-9873-8.
- S. Jafarzadeh, R. H. Cameron, S. K. Solanki, A. Pietarila, A. Feller, A. Lagg, and A. Gandorfer. Migration of Ca II H bright points in the internetwork. *Astron. Astrophys.*, 563:A101, March 2014. doi: 10.1051/0004-6361/201323011.
- S. Kholikov, A. Serebryanskiy, and J. Jackiewicz. Meridional Flow in the Solar Convection Zone. I. Measurements from GONG Data. *Astrophys. J.*, 784:145, April 2014. doi: 10.1088/0004-637X/784/2/145.
- R. Komm, I. González Hernández, F. Hill, R. Bogart, M. C. Rabello-Soares, and D. Haber. Subsurface meridional flow from HMI using the ring-diagram pipeline. *Solar Phys.*, 287:85–106, October 2013. doi: 10.1007/s11207-012-0073-y.
- A. G. Kosovichev and T. L. Duvall, Jr. Acoustic tomography of solar convective flows and structures. In F. P. Pijpers, J. Christensen-Dalsgaard, and C. S. Rosenthal, editors, *SCORe’96 : Solar Convection and Oscillations and their Relationship*, volume 225 of *Astrophysics and Space Science Library*, pages 241–260, December 1997.
- A. G. Kosovichev, S. Basu, R. Bogart, T. L. Duvall, Jr., I. Gonzalez-Hernandez, D. Haber, T. Hartlep, R. Howe, R. Komm, S. Kholikov, K. V. Parchevsky, S. Tripathy, and J. Zhao. Local helioseismology of sunspot regions: Comparison of ring-diagram and time-distance results. *Journal of Physics Conference Series*, 271(1):012005, January 2011. doi: 10.1088/1742-6596/271/1/012005.
- J. Langfellner, L. Gizon, and A. C. Birch. Anisotropy of the solar network magnetic field around the average supergranule. *Astron. Astrophys.*, 579:L7, July 2015a. doi: 10.1051/0004-6361/201526422.
- J. Langfellner, L. Gizon, and A. C. Birch. Spatially resolved vertical vorticity in solar supergranulation using helioseismology and local correlation tracking. *Astron. Astrophys.*, 581:A67, September 2015b. doi: 10.1051/0004-6361/201526024.
- R. B. Leighton, R. W. Noyes, and G. W. Simon. Velocity Fields in the Solar Atmosphere. I. Preliminary Report. *Astrophys. J.*, 135:474, March 1962. doi: 10.1086/147285.
- K. G. Libbrecht, B. D. Popp, J. M. Kaufman, and M. J. Penn. The excitation and damping of solar oscillations. *Nature*, 323:235–238, September 1986. doi: 10.1038/323235a0.
- C. Lindsey and D. C. Braun. Helioseismic Holography. *Astrophys. J.*, 485:895–903, August 1997.
- C. Lindsey and D. C. Braun. Seismic Images of the Far Side of the Sun. *Science*, 287:1799–1801, March 2000. doi: 10.1126/science.287.5459.1799.
- J. W. Lord, R. H. Cameron, M. P. Rast, M. Rempel, and T. Roudier. The Role of Subsurface Flows in Solar Surface Convection: Modeling the Spectrum of Supergranular and Larger Scale Flows. *Astrophys. J.*, 793:24, September 2014. doi: 10.1088/0004-637X/793/1/24.

- M. S. Miesch. Large-scale dynamics of the convection zone and tachocline, 2005. URL <http://www.livingreviews.org/lrsp-2005-1>. [Online Article]: cited 13-August-2015.
- M. S. Miesch, J. R. Elliott, J. Toomre, T. L. Chune, G. A. Glatzmaier, and P. A. Gilman. Three-dimensional Spherical Simulations of Solar Convection. I. Differential Rotation and Pattern Evolution Achieved with Laminar and Turbulent States. *Astrophys. J.*, 532:593–615, March 2000. doi: 10.1086/308555.
- M. S. Miesch, A. S. Brun, M. L. De Rosa, and J. Toomre. Structure and Evolution of Giant Cells in Global Models of Solar Convection. *Astrophys. J.*, 673:557–575, January 2008. doi: 10.1086/523838.
- J. O. Murphy. The Effect of a Magnetic Field on the Onset of Thermal Convection when Constant Flux Boundary Conditions Apply. *Proceedings of the Astronomical Society of Australia*, 3:164, September 1977.
- D. Orozco Suárez, Y. Katsukawa, and L. R. Bellot Rubio. The Connection between Internetwork Magnetic Elements and Supergranular Flows. *Astrophys. J. Lett.*, 758:L38, October 2012. doi: 10.1088/2041-8205/758/2/L38.
- E. N. Parker. Kinematical Hydromagnetic Theory and its Application to the Low Solar Photosphere. *Astrophys. J.*, 138:552, August 1963. doi: 10.1086/147663.
- S. P. Rajaguru and H. M. Antia. Meridional circulation in the solar convection zone: time-distance helioseismic inferences from four years of HMI/SDO observations. *ArXiv e-prints*, October 2015.
- M. P. Rast. The Scales of Granulation, Mesogranulation, and Supergranulation. *Astrophys. J.*, 597:1200–1210, November 2003. doi: 10.1086/381221.
- M. Rempel, M. Schüssler, and M. Knölker. Radiative Magnetohydrodynamic Simulation of Sunspot Structure. *Astrophys. J.*, 691:640–649, January 2009. doi: 10.1088/0004-637X/691/1/640.
- M. Rieutord and F. Rincon. The suns supergranulation, 2015. URL <http://www.livingreviews.org/lrsp-2010-2>. [Online Article]: cited 13-August-2015.
- M. Rieutord and J.-P. Zahn. Turbulent plumes in stellar convective envelopes. *Astron. Astrophys.*, 296:127, April 1995.
- M. Rieutord, N. Meunier, T. Roudier, S. Rondi, F. Beigbeder, and L. Parès. Solar supergranulation revealed by granule tracking. *Astron. Astrophys.*, 479:L17–L20, February 2008. doi: 10.1051/0004-6361/20079077.
- M. Rieutord, T. Roudier, F. Rincon, J.-M. Malherbe, N. Meunier, T. Berger, and Z. Frank. On the power spectrum of solar surface flows. *Astron. Astrophys.*, 512:A4, March 2010. doi: 10.1051/0004-6361/200913303.
- L. Rightmire-Upton, D. H. Hathaway, and K. Kosak. Measurements of the sun’s high-latitude meridional circulation. *Astrophys. J. Lett.*, 761:L14, December 2012. doi: 10.1088/2041-8205/761/1/L14.
- F. Rincon and M. Rieutord. Stability of a compressible fluid layer in a magnetic field: a simple model for supergranulation. In F. Combes, D. Barret, T. Contini, and L. Pagani, editors, *SF2A-2003: Semaine de l’Astrophysique Française*, page 103, 2003.
- T. Roudier, M. Rieutord, J. M. Malherbe, N. Renon, T. Berger, Z. Frank, V. Prat, L. Gizon, and M. Švanda. Quasi full-disk maps of solar horizontal velocities using SDO/HMI data. *Astron. Astrophys.*, 540:A88, April 2012. doi: 10.1051/0004-6361/201118678.
- S. Routh, D. A. Haber, B. W. Hindman, R. S. Bogart, and J. Toomre. The influence of tracking rate on helioseismic flow inferences. *J. Phys. Conf. Ser.*, 271(1):012014, January 2011. doi: 10.1088/1742-6596/271/1/012014.

- J. Schou, J. Christensen-Dalsgaard, and M. J. Thompson. On comparing helioseismic two-dimensional inversion methods. *Astrophys. J.*, 433:389–416, September 1994. doi: 10.1086/174653.
- J. Schou, H. M. Antia, S. Basu, R. S. Bogart, R. I. Bush, S. M. Chitre, J. Christensen-Dalsgaard, M. P. Di Mauro, W. A. Dziembowski, A. Eff-Darwich, D. O. Gough, D. A. Haber, J. T. Hoeksema, R. Howe, S. G. Korzennik, A. G. Kosovichev, R. M. Larsen, F. P. Pijpers, P. H. Scherrer, T. Sekii, T. D. Tarbell, A. M. Title, M. J. Thompson, and J. Toomre. Helioseismic Studies of Differential Rotation in the Solar Envelope by the Solar Oscillations Investigation Using the Michelson Doppler Imager. *Astrophys. J.*, 505:390–417, September 1998. doi: 10.1086/306146.
- M. Schwarzschild. On the scale of photospheric convection in red giants and supergiants. *Astrophys. J.*, 195:137–144, January 1975. doi: 10.1086/153313.
- T. Sekii, A. G. Kosovichev, J. Zhao, S. Tsuneta, H. Shibahashi, T. E. Berger, K. Ichimoto, Y. Katsukawa, B. Lites, S. Nagata, T. Shimizu, R. A. Shine, Y. Suematsu, T. D. Tarbell, and A. M. Title. Initial Helioseismic Observations by Hinode/SOT. *Pub. Astron. Soc. Japan*, 59:637, November 2007. doi: 10.1093/pasj/59.sp3.S637.
- G. W. Simon and R. B. Leighton. Velocity Fields in the Solar Atmosphere. III. Large-Scale Motions, the Chromospheric Network, and Magnetic Fields. *Astrophys. J.*, 140:1120, October 1964. doi: 10.1086/148010.
- G. W. Simon, A. M. Title, K. P. Topka, T. D. Tarbell, R. A. Shine, S. H. Ferguson, H. Zirin, and SOUP Team. On the relation between photospheric flow fields and the magnetic field distribution on the solar surface. *Astrophys. J.*, 327:964–967, April 1988. doi: 10.1086/166253.
- H. B. Snodgrass. Separation of large-scale photospheric Doppler patterns. *Solar Phys.*, 94:13–31, August 1984. doi: 10.1007/BF00154804.
- R. F. Stein and A. Nordlund. Topology of convection beneath the solar surface. *Astrophys. J. Lett.*, 342:L95–L98, July 1989. doi: 10.1086/185493.
- R. F. Stein and Å. Nordlund. Simulations of Solar Granulation. I. General Properties. *Astrophys. J.*, 499:914–933, May 1998.
- R. F. Stein, D. Georgobiani, W. Schafenberger, Å. Nordlund, and D. Benson. Supergranulation Scale Convection Simulations. In E. Stempels, editor, *15th Cambridge Workshop on Cool Stars, Stellar Systems, and the Sun*, volume 1094 of *American Institute of Physics Conference Series*, pages 764–767, February 2009. doi: 10.1063/1.3099227.
- M. J. Thompson, J. Toomre, E. R. Anderson, H. M. Antia, G. Berthomieu, D. Burtonclay, S. M. Chitre, J. Christensen-Dalsgaard, T. Corbard, M. De Rosa, C. R. Genovese, D. O. Gough, D. A. Haber, J. W. Harvey, F. Hill, R. Howe, S. G. Korzennik, A. G. Kosovichev, J. W. Leibacher, F. P. Pijpers, J. Provost, E. J. Rhodes, Jr., J. Schou, T. Sekii, P. B. Stark, and P. R. Wilson. Differential Rotation and Dynamics of the Solar Interior. *Science*, 272:1300–1305, May 1996. doi: 10.1126/science.272.5266.1300.
- M. J. Thompson, J. Christensen-Dalsgaard, M. S. Miesch, and J. Toomre. The Internal Rotation of the Sun. *ARA&A*, 41:599–643, 2003. doi: 10.1146/annurev.astro.41.011802.094848.
- J. Toomre, J. Christensen-Dalsgaard, R. Howe, R. M. Larsen, J. Schou, and M. J. Thompson. Time Variability of Rotation in Solar Convection Zone From soi-mdi. *Solar Phys.*, 192:437–448, March 2000. doi: 10.1023/A:1005223308315.
- R. Trampedach and R. F. Stein. The Mass Mixing Length in Convective Stellar Envelopes. *Astrophys. J.*, 731:78, April 2011. doi: 10.1088/0004-637X/731/2/78.
- S. D. Ustyugov. Realistic magnetohydrodynamical simulation of solar local supergranulation. *Physica Scripta Volume T*, 142(1):014031, December 2010. doi: 10.1088/0031-8949/2010/T142/014031.

- M. Švanda. Inversions for Average Supergranular Flows Using Finite-frequency Kernels. *Astrophys. J. Lett.*, 759:L29, November 2012. doi: 10.1088/2041-8205/759/2/L29.
- M. Švanda. Issues with time-distance inversions for supergranular flows. *Astron. Astrophys.*, 575:A122, March 2015. doi: 10.1051/0004-6361/201425203.
- A. Vögler, S. Shelyag, M. Schüssler, F. Cattaneo, T. Emonet, and T. Linde. Simulations of magneto-convection in the solar photosphere. Equations, methods, and results of the MURaM code. *Astron. Astrophys.*, 429:335–351, January 2005. doi: 10.1051/0004-6361:20041507.
- Y. Wang, R. W. Noyes, T. D. Tarbell, and A. M. Title. Vorticity and Divergence in the Solar Photosphere. *Astrophys. J.*, 447:419, July 1995. doi: 10.1086/175886.
- M. A. Weber, Y. Fan, and M. S. Miesch. The Rise of Active Region Flux Tubes in the Turbulent Solar Convective Envelope. *Astrophys. J.*, 741:11, November 2011. doi: 10.1088/0004-637X/741/1/11.
- M. F. Woodard. Probing Supergranular Flow in the Solar Interior. *Astrophys. J.*, 668:1189–1195, October 2007. doi: 10.1086/521391.
- M. F. Woodard and K. G. Libbrecht. Observations of time variation in the sun’s rotation. *Science*, 260: 1778–1781, June 1993. doi: 10.1126/science.260.5115.1778.
- J. Zhao and A. G. Kosovichev. On the inference of supergranular flows by time-distance helioseismology. In H. Sawaya-Lacoste, editor, *GONG+ 2002. Local and Global Helioseismology: the Present and Future*, volume 517 of *ESA Special Publication*, pages 417–420, February 2003.
- J. Zhao and A. G. Kosovichev. Torsional Oscillation, Meridional Flows, and Vorticity Inferred in the Upper Convection Zone of the Sun by Time-Distance Helioseismology. *Astrophys. J.*, 603:776–784, March 2004. doi: 10.1086/381489.
- J. Zhao, K. Nagashima, R. S. Bogart, A. G. Kosovichev, and T. L. Duvall, Jr. Systematic center-to-limb variation in measured helioseismic travel times and its effect on inferences of solar interior meridional flows. *Astrophys. J. Lett.*, 749:L5, April 2012. doi: 10.1088/2041-8205/749/1/L5.
- J. Zhao, R. S. Bogart, A. G. Kosovichev, T. L. Duvall, Jr., and T. Hartlep. Detection of Equatorward Meridional Flow and Evidence of Double-cell Meridional Circulation inside the Sun. *Astrophys. J. Lett.*, 774:L29, September 2013. doi: 10.1088/2041-8205/774/2/L29.

Appendix A

Center-to-Limb Velocity Systematic

This appendix is based on work previously published in Greer et al. (2013) and is largely a restatement of that paper with some modifications to the language to make it more consistent with the rest of this thesis. This work is referenced in Chapter 5 when considering the processing of frequency shift measurements prior to inversion. While the process detailed here was designed with the measurement of global-scale flows in mind, it also is useful when considering convective flows over a broad range of scales.

We use HMI Ring-Diagram Pipeline data to measure a center-to-limb systematic effect seen in frequency shift measurements. To separate this signal from persistent flow patterns of physical origin, we perform a least-squares fit to the data for each wave mode. We fit to a model that includes both the radially symmetric systematic as well as global zonal and meridional flow components. The resulting measurements of the systematic error reveal a smooth trend as both a function of mode frequency and phase speed. The systematic error at 45° from disk center ranges from 20 m s^{-1} radially inward to 50 m s^{-1} radially outward. The implications for determining global scale meridional flows is discussed.

A.1 Introduction

Since the introduction of local helioseismology, the character and strength of the meridional flow has been of keen interest. Measurements of the circulation strength have both constrained global-scale numerical models as well as informed flux transport dynamo models. Both ring-diagram and time-distance helioseismology have been able to reliably provide these measurements for the upper few percent in radius of the solar interior.

Comparison between time-distance helioseismology and ring-diagram analysis has shown good agreement, but only after removing the large-scale systematics through subtraction of time-averaged velocities (Hindman et al. 2004). While the residuals of the two procedures are highly correlated with each other, it is unclear whether the observed center-to-limb velocity systematic is identical between the two methods. Recent time-distance work has made progress in accounting for this effect by using east-west travel-time differences to correct north-south measurements of deep meridional flows (Zhao et al. 2012). This correction is made assuming the effect is a position-dependent phase shift in the helioseismic Dopplergram measurements, which translates to a frequency shift and therefore subsurface velocity in ring-diagram analysis. There has been comparatively little progress in ring-diagram analysis in accounting for this effect.

In this appendix, we expand on these measurements of the center-to-limb frequency shift systematic in ring-diagram analysis and provide a method of separating the systematic signal from various real global-scale flows. Analysis of the anomalous frequency shifts introduced by this systematic error shows an inconvenient trend with phase speed that can easily be confused with subsurface flows.

A.2 Temporally-Averaged Frequency Shift Data

The data for this analysis is from the standard HMI Ring-Diagram Pipeline products (Bogart et al. 2011a;b). We use the ring-diagram frequency shift measurements from two tile sizes (15° , 30°) arranged in mosaics covering a large fraction of the solar disk. Temporal averages at each disk position for each mode are performed over three Carrington rotations for 15° tiles and over nine rotations for 30° tiles.

The resulting frequency shift maps over the solar disk reveal a problematic center-to-limb variation that must arise from an unidentified systematic error. The radial symmetry of the systematic error is more apparent if the longitudinal average of the zonal flow is first removed (Figure A.1). Modes of low frequency tend to have an erroneous outflow from disk center (Figure A.1a) while high frequency modes have an inflow (Figure A.1b).

The measured frequency shifts can be broken into four distinct components to aid in separating the systematic error from known real flows in the Sun. The first component is due to the random convective flow, which we assume is uncorrelated at a given disk position as the Sun rotates by. We treat this component




Figure A.1: (a) Residual f -mode frequency shift measurements at $\nu \approx 3.2$ mHz after subtraction of differential rotation. (b) Similar residual for p_4 frequency shift measurements at $\nu \approx 4.8$ mHz.

as random noise, and assume that in a long time average it disappears. The second component is due to the large-scale zonal flow that includes both the signature of differential rotation as well as a monotonically increasing prograde flow in depth corresponding to the near-surface shear layer. The third component is due to the large-scale meridional flow that is expected to be primarily poleward in each hemisphere. The final component is the systematic radial flow pattern that persists in a long time average along with the zonal and meridional flows (Zhao et al. 2012).

In order to isolate and measure this anomalous signal, we must first make some assumptions about it. The first is that this effect is linear and of small enough amplitude that we can treat the systematic frequency shift as an additive effect. The second assumption is that the added anomalous frequency shift is directed radially inward or outward from disk center, and that the amplitude is purely a function of distance from disk center. Figure A.1 demonstrates a primarily radial frequency shift systematic, but also shows a slight radial asymmetry for the higher frequency mode. For the purposes of this paper we utilize a simple model that ignores any radial asymmetry.

A.3 Fitting Procedure

A time average of measured frequency shifts will not only demonstrate the radial systematic, but also the longitudinally invariant zonal and meridional flows of physical origin. It is for this reason that we

cannot simply subtract an average frequency shift map from our measurements, as this would remove any persistent large-scale flows of interest. One possibility for isolating the radial systematic is to consider east-west asymmetry in the zonal flow as in Zhao et al. (2012). A long time average of the subsurface zonal flows should not show any significant variation with longitude, so any large-scale pattern along this direction can be attributed to a systematic error. Since we assume the effect to be a function of distance from disk center, this method can only be used near the equator. This method provides a simple measure of the center-to-limb systematic error, but discards useful data from most of the solar disk. Instead, we use our assumption that it is a radially symmetric effect to perform a whole-disk least squares fit to both the zonal and meridional frequency shift components simultaneously.

In order to successfully isolate the radial systematic from the large-scale flows in each frequency shift component, we fit polynomials to the mean zonal and meridional components at the same time that we are fitting the radial component. The complete function used is then a sum of these three components:

$$\mathbf{u}^{(k,n)}(\phi, \lambda) = u_r^{(k,n)}(r)\hat{\mathbf{r}} + u_m^{(k,n)}(\phi)\hat{\boldsymbol{\phi}} + u_z^{(k,n)}(\phi)\hat{\boldsymbol{\lambda}}, \quad (\text{A.1})$$

$$u_m^{(k,n)}(\phi) = a_0^{(k,n)} + \sum_{j=1}^4 a_j^{(k,n)} \sin(j\phi), \quad (\text{A.2})$$

$$u_z^{(k,n)}(\phi) = \sum_{j=0}^2 b_j^{(k,n)} \phi^{2j}, \quad (\text{A.3})$$

$$u_r^{(k,n)}(r) = \sum_{j=1}^4 c_j^{(k,n)} r^j. \quad (\text{A.4})$$

Here, λ is longitude, ϕ is latitude, r is the great-circle distance from disk center (not from 0° latitude due to the nonzero solar B-angle), k is the mode wavenumber, n is the mode order, and u_r , u_m , and u_z are the radial, meridional, and zonal frequency shift components, respectively. We anticipate that both the meridional and zonal flows change as a function of depth in the Sun, so the fit is performed independently on every mode (k, n) since each mode samples different depths. We restrict the mode set to frequencies below the acoustic cutoff frequency (5 mHz), since only these modes enter into helioseismic inversions. The




Figure A.2: Examples of successful fits for various modes, wavenumbers are in units of Mm^{-1} . Solid lines are obtained from 30° tiles, dashed are from 15° .

mode set is different for each tile size, and the systematic error might be sensitive to the size of the analysis domain. Therefore, we perform the fit independently for both sizes. The final result is a characterization of the radial systematic that is comprised of four polynomial coefficients $c_j^{(k,n)}$ for each mode and for each tile size.

A.4 Results & Analysis

Fits for four example modes of both tile sizes are shown in Figure A.2. The radial profiles demonstrate a smooth transition between two typical types of radial systematic seen. For low frequencies, it is a positive (radially outward) frequency shift that peaks in amplitude around 40° from disk center, while at higher frequencies it is a negative (radially inward) frequency shift that increases in amplitude with increasing distance from disk center. The meridional profiles for these modes show poleward flows of around 15 m s^{-1} that are roughly constant out to $\pm 60^\circ$. Since the HMI Ring-Diagram Pipeline data is tracked at the Carrington rate which ignores surface differential rotation, the zonal flow profiles have picked up the $\sim 180 \text{ m s}^{-1}$ shear between disk center and $\pm 60^\circ$ latitude. The higher radial order modes sample deeper into the Sun, causing the slight increase in zonal frequency shift with increasing order due to the near surface shear layer. Finally, we see that these frequency shift profiles are consistent between the two tile sizes considered here.

To look for trends in the amplitude of the measured systematic, we now switch to measurements made using 30° tiles. These larger analysis regions provide us with better measurements of the systematic, with both smaller errors and a significantly expanded mode set. The value of the radial profile for each mode at a distance from disk center of 45° is shown in Figure A.3.

The systematic error is positive (radially outward) for frequencies below 4 mHz and begins to trend negative (radially inward) as the frequency increases above this (Figure A.3a). The magnitude follows a smooth curve for each radial order, and the zero-crossing happens at slightly different frequencies for each. As a function of phase speed (Figure A.3b), the different radial orders collapse onto a common curve for low frequencies, creating a nearly linear increase in magnitude. At higher frequencies, each mode ridge diverges from the common curve.

A.5 Discussion & Conclusions

The two distinct regimes of the radial systematic (high and low frequency) present an interesting challenge for interpreting helioseismic results. Of particular interest is the determination of the meridional flows throughout the bulk of the convection zone. The high-frequency modes introduce a radially inward frequency shift that is strongest near the limb. This has the possibility for creating the appearance of high-latitude counter cells of equatorward flow. Low-frequency modes instead have a radially outward systematic that increases steadily as a function of phase speed (often used as a proxy for depth). This can significantly alter measurements of the magnitude and direction of the meridional flows even at mid-latitudes. Helioseismic measurements estimate the meridional flow to be around $15\text{--}20 \text{ m s}^{-1}$ poleward for most depths and latitudes (Giles et al. 1997; Haber et al. 2002; Zhao and Kosovichev 2004). With an unaccounted-for systematic frequency shift of around $5\text{--}10 \text{ m s}^{-1}$ poleward in the upper 20 Mm, these measurements are likely to be overestimates of the true meridional speed. Below 50 Mm, the systematic may be the dominant signal in any determination of meridional flows.

It is more difficult to say what pattern the systematic will imprint in an inversion. While using phase speed is not a precise method of determining the depth of a measurement, Figure A.3b still demonstrates that there will be contradictory frequency shifts introduced at some depths. Velocity measurements made




Figure A.3: (a) Amplitude of the radial systematic frequency shift using 30° tiles at $r = 45^\circ$ from disk center as a function of cyclic frequency. (b) Same as above, but as a function of horizontal phase speed (ω/k). Colors denote different radial orders.

near 3 mHz typically have smaller errors than those made at higher frequencies, so a standard inversion technique will put more weight in the low-frequency trend.




The radial dependence of this effect lends credence to the idea that it is rooted in a phase shift of the observed waves as a function of height in the atmosphere. As the height of observation changes from disk center to the limb, the phase shift of a wave mode as a function of height becomes a phase shift as a function of disk position. A first order linear approximation of the phase shift across an analysis tile will cause an additive frequency shift that is identical in form to any existing subsurface flow. While this basic mechanism seems likely, it is still uncertain what the origin of the phase shift is. Baldner and Schou (2012) suggest asymmetric flows as a possible origin.




We can clearly see the radial systematic error present in data, and we are able to reliably measure the amplitude of it for each wave mode. Considering the magnitudes presented, it is crucial to consider this effect when measuring subsurface flows, particularly global meridional flows. While it would be beneficial to have a theoretical understanding of the cause, we feel it is sufficient to use the radial profile obtained from our fitting procedure as a correction to ring-diagram frequency shift measurements.




Appendix B



Averaging Kernels for Three-Dimensional OLA Inversion

The following figures show the averaging kernels for the full three-dimensional ring-diagram inversion detailed in Chapter 4. Each figure demonstrates a few qualities of a distinct averaging kernel, and the format of each figure is the same. Each averaging kernel was created with a different target depth, indicated by the arrow in the left-most panel (a). This panel shows the horizontally-integrated sensitivity of the averaging kernel, normalized by the peak value. The panel to the right of this (b) shows the averaging kernel integrated in latitude and shown as a function of depth and longitude. Red indicates positive sensitivity, blue indicates negative. Panel (c) shows a longitudinal cut of the averaging kernel at the depth of peak sensitivity in panel (a). Panel (d) is a cut at the same depth shown as a function of both longitude and latitude. The text in each figure indicates the target depth of the averaging kernel, the FWHM in depth of the peak in panel (a), and the FWHM in longitude of the peak in panel (c).









Appendix C

Code Listing

CODE NAME: FRACK

LANGUAGE: FORTRAN, 1300 lines

URL: <http://lcd.colorado.edu/bgreer/codes/FRACK>

PURPOSE: This code performs the Dopplergram projecting and tracking necessary to create ring-diagram tiles (see Chapter 2). The code reads in full-disk Dopplergram images and outputs a set of three-dimensional data cubes corresponding to each tile. The size of the tiles and locations on the disk for each tile are provided to the code through command-line options. The code relies on distributed-memory (MPI) parallelism in order to track multiple tiles from the same set of Dopplergrams simultaneously.

CODE NAME: PSPEC

LANGUAGE: FORTRAN, 350 lines

URL: <http://lcd.colorado.edu/bgreer/codes/PSPEC>

PURPOSE: This code takes a single tracked ring-diagram tile and produces a power spectrum from the three-dimensional Fourier transform. This code is based closely on a similar code found in the HMI Ring-Diagram Pipeline (Bogart et al. 2011a). The input tile is multiplied by a pre-determined apodization function, Fourier transformed, and then the power is computed. The power spectrum is then interpolated on to polar coordinates (k, θ, ω) and sub-sampled in θ . The code relies on shared-memory (OpenMP) parallelism to speed up various steps of the process and has been found to scale efficiently up to around 10 processors.

CODE NAME: Multi-Ridge Fitting Code (MRF)

LANGUAGE: C, 3500 lines

URL: <http://lcd.colorado.edu/bgreer/codes/MRF>

PURPOSE: This code performs the non-linear ring-fitting algorithm described in Chapter 3 on power spectra produced by the PSPEC code. There is no parallelism built into this code, as the typical usage is to run many instances of the code in parallel on different power spectra.

CODE NAME: ATLAS

LANGUAGE: FORTRAN / C, 4000 lines

URL: <http://lcd.colorado.edu/bgreer/codes/ATLAS>

PURPOSE: This code performs all of the tracking, projecting, power-spectrum-making, and ring-fitting needed to turn a set of Dopplergram images into a set of frequency shift measurements. The details of these steps are covered between Chapters 2 and 3, and this code is essentially a combination of the previous three codes (FRACK, PSPEC, MRF). The code reads in a list of longitude/latitude coordinates corresponding to the desired tile centers and a set of full-disk Dopplergram images and outputs frequency shift measurements from each wave mode of each tile. The code relies on both distributed-memory (MPI) and shared-memory (OpenMP) parallelism to scale up to around 1000 processes. Due to the immense volume of data produced by the tracking and projecting steps, the intermediate data products (tiles, power spectra) are never written out.

CODE NAME: Massively-Overlapped Ring-Diagram Inversion (MORDI)

LANGUAGE: C++, 2000 lines

URL: <http://lcd.colorado.edu/bgreer/codes/MORDI>

PURPOSE: This code performs three-dimensional ring-diagram inversions, as explained in Chapter 4. The code reads in frequency shift measurements and their associated sensitivity kernels and outputs two-dimensional slices of the subsurface flow field at a constant depth and (optionally) the associated averaging kernels. It relies on both distributed-memory (MPI) and shared-memory (OpenMP) parallelism to scale efficiently up to a few thousand processors, but can also run reasonably well on small machines (1-4 cpus). The actions of the code are modified by command-line parameters, which enable a significant amount of flexibility when setting up an inversion.