



Velocity and Density Fluctuations in the Quiet Sun Corona

Michael Hahn¹ , Xiangrong Fu^{2,3}, Stefan J. Hofmeister¹ , Yifan Huang² , Alexandros Koukras¹ , and Daniel Wolf Savin¹ ¹ Columbia Astrophysics Laboratory, Columbia University, 550 West 120th Street, New York, NY 10027, USA; mhahn@astro.columbia.edu² Los Alamos National Laboratory, P.O. Box 1663, Los Alamos, NM 87545, USA³ New Mexico Consortium, 4200 W Jemez Rd #200, Los Alamos, NM 87544, USA

Received 2025 February 19; revised 2025 March 14; accepted 2025 March 14; published 2025 April 28

Abstract

We investigate the properties and relationship between Doppler velocity fluctuations and intensity fluctuations in the off-limb quiet Sun corona. These are expected to reflect the properties of Alfvénic and compressive waves, respectively. The data come from the Coronal Multichannel Polarimeter (COMP). These data were studied using spectral methods to estimate the power spectra, amplitudes, perpendicular correlation lengths, phases, trajectories, dispersion relations, and propagation speeds of both types of fluctuations. We find that most velocity fluctuations are due to Alfvénic waves but that intensity fluctuations come from a variety of sources, likely including fast and slow mode waves, as well as aperiodic variations. The relation between the velocity and intensity fluctuations differs depending on the underlying coronal structure. On short closed loops, the velocity and intensity fluctuations have similar power spectra and speeds. In contrast, on longer nearly radial trajectories, the velocity and intensity fluctuations have different power spectra, with the velocity fluctuations propagating at much faster speeds than the intensity fluctuations. Considering the temperature sensitivity of COMP, these longer structures are more likely to be closed field lines of the quiet Sun rather than cooler open field lines. That is, we find the character of the interactions of Alfvénic waves and density fluctuations depends on the length of the magnetic loop on which they are traveling.

Unified Astronomy Thesaurus concepts: Quiet sun (1322); Solar physics (1476); Solar coronal waves (1995); Quiet solar corona (1992); Atomic spectroscopy (2099); Time series analysis (1916)

1. Introduction

Wave-driven models of coronal heating consider that energy is transported to the corona by Alfvén waves, which damp in the corona and transfer that energy into particle heating. In most models, wave dissipation is mediated by a turbulent cascade of wave energy from the long length scales of Alfvén waves to smaller scales where particle interactions can heat the plasma. Turbulence arises through a nonlinear interaction between counter-propagating Alfvén waves (G. G. Howes & K. D. Nielson 2013). In closed loops, counter-propagating waves arise naturally as waves excited at the two footpoints of the loop interact near the loop top. In open structures, such as coronal holes, the origin of sunward-propagating waves is less clear. Large-scale gradients in the magnetic field and density in the corona are thought to be too weak to efficiently reflect Alfvén waves (V. Réville et al. 2018; M. Asgari-Targhi et al. 2021).

The reflection of waves may be enhanced by interaction with fluctuations in the plasma density, such as an interaction with slow mode waves (M. Asgari-Targhi et al. 2021). Alfvén waves are expected to be reflected from parallel gradients in the Alfvén speed, and density fluctuations increase these gradients (M. Heinemann & S. Olbert 1980; Z. E. Musielak et al. 1992; M. Velli 1993). Theoretical models for open magnetic field regions show that adding density fluctuations to the plasma significantly enhances wave reflection and turbulent heating compared to a smooth radial variation of the background plasma parameters (M. Asgari-Targhi et al. 2021).

Another process that may produce these sunward-propagating waves is parametric decay instability (PDI). This is a nonlinear interaction in which a forward-propagating “pump” Alfvén wave excites a forward acoustic wave and a backward secondary Alfvén wave. PDI is able to enhance turbulence and promote coronal heating in several ways. First, PDI directly generates a backward-propagating Alfvén wave, which can interact with forward-propagating Alfvén waves excited at the Sun to drive turbulence. Second, PDI produces large amplitude density fluctuations in the form of acoustic waves, which promotes Alfvén wave reflection. There has been much numerical work showing how PDI may lead to efficient coronal heating (M. Shoda & T. Yokoyama 2016; V. Réville et al. 2018; M. Shoda et al. 2018, 2019). There is also evidence suggesting that PDI occurs in the transition region (M. Hahn et al. 2022a) and indirect evidence that it has influenced the solar wind (T. E. Bowen et al. 2018; J. C. Kasper et al. 2021), although direct evidence for PDI in the solar wind is still lacking (G. P. Zank et al. 2022; L.-L. Zhao et al. 2022).

Besides these, there are many other potential interactions between Alfvén waves and density fluctuations. For example, low-frequency large amplitude waves may drive density fluctuations due to the wave magnetic pressure (J. V. Hollweg 1971a, 1971b). Additionally, Alfvén waves propagating along magnetic flux tubes with strong gradient in the density or magnetic field strength perpendicular to the mean field may undergo nonlinear phase mixing, generating fast mode waves (V. M. Nakariakov et al. 1997).

Our objective here is to identify observational relationships between velocity and density fluctuations in the quiet Sun corona. We take a global view, observing the entire off-limb low corona up to about $1.3 R_{\odot}$ during a typical time that does not appear to be affected by any remarkable events. We study



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

spectroscopic data observing an Fe XIII line, which has a formation temperature of about 1.8 MK. Given the high temperature of the line, we observe material from the quiet Sun and active regions, but not the cooler plasma associated with coronal holes and open field lines. Velocity fluctuations are observed through the Doppler shift of the emission lines, and the density fluctuations are related to the intensity of the line. Throughout, we refer to fluctuations rather than waves, as there may be other sources of variations. We expect, however, that many velocity fluctuations are due to Alfvénic waves and the density fluctuations to compressive waves. We use the term *Alfvénic* phenomenologically to refer to magnetized plasma waves propagating at about the Alfvén speed, as a more precise description of their nature, such as being kink or torsional modes, is difficult to ascertain observationally. We test the validity of our interpretations of the fluctuations as part of the analysis. Additionally, although our broad view is not limited exclusively to a particular physical process, we do test in more detail for PDI using methods developed by M. Hahn et al. (2022a), such as by looking for a frequency scaling between the power spectra of intensity and velocity fluctuations.

The rest of this paper is organized as follows. In Section 2 we give a brief overview of the instrument and describe the observations. The analysis is presented throughout Section 3. There, we aim to describe a broad range of properties of the observed density and velocity fluctuations. We begin with characteristics that can be obtained from time variation within a single pixel, then describe how we can track the wave propagation through the plane of sky, and finally derive properties of the waves that depend on both the space and time evolution of the fluctuations. In the course of the analysis, we describe relevant comparisons to other works and possible conclusions based on those measurements. We defer some interpretations that rely on multiple aspects of the analysis to the discussion in Section 4. Section 5 summarizes the main results.

2. Instrument and Observations

We use data from the Coronal Multichannel Polarimeter (COMP; S. Tomczyk et al. 2008). COMP uses a tunable filter and polarimeter to make spectropolarimetric observations at three wavelengths that span the Fe XIII 10747 Å emission line. At each of the three wavelengths, COMP produces an image of the off-limb solar corona in the four Stokes parameters; however, for the applications discussed here, we are concerned only with the spectroscopic aspects of the data. These images span a radial distance of about 1.05–1.3 $R_\odot$ with a spatial resolution of approximately 4.5 pixel^{-1} . We have used the Level-2 COMP data product, available through the Mauna Loa Solar Observatory, which provides the intensity and Doppler velocity of the line, as interpreted by assuming that the Fe XIII 10747 Å line has a Gaussian shape (see R. Sharma & R. J. Morton 2023).

Our data were obtained on 2016 July 25 from 20:25 to 21:53 UT. Figure 1 presents a context image, showing the time-averaged Fe XIII 10747 Å intensity measured during the observation. The labels in the image are intended to highlight several apparent structures. Labels 1, 3, 4, and 7 are located near short closed loops, whose loop tops are clearly visible within the COMP field of view. Labels 2 and 5 are located near much longer structures, that appear nearly radial at the low heights within the field of view. These long structures are most

likely the bases of long quiet Sun streamer loops. Although they appear “open” in this data set, the Fe XIII emission viewed by COMP is formed at high temperatures and so we do not expect to see the open magnetic fields from coronal holes. For future reference, we observationally define the short loops as those whose loop tops are below $1.3 R_\odot$ so that the entire loop is visible in the COMP data, and we define long loops as those that extend to much larger heights. Finally, on the west limb near the region labeled 6, there are several active regions (NOAA 12565, 12566, and 12567) behind the limb, resulting in complex magnetic structures in the corona. Because of the complex magnetic structure there and to improve the clarity of the presentation, we will often show figures focusing on the more quiescent East limb. A context image showing this more condensed field of view is shown in Figure 2.

During the observation, COMP obtained intensity and velocity data with a regular cadence of $\Delta t = 30 \text{ s}$. Our aim is to determine the typical relationships between velocity and density fluctuations in the corona. There was no particular scientific significance to the selected date, which was chosen at random. By coincidence, it turns out that the same data set was one of those studied by R. Sharma & R. J. Morton (2023) in their recent study of the perpendicular correlation length for Alfvénic waves.

Our data set contains a time series of two-dimensional spatial maps of the intensity and Doppler shift in the corona. The COMP Level-2 data calibration pipeline sets to zero any pixels where the intensity was below a threshold of two millionths of the sky brightness. However, these pixels are not necessarily the same in each image. We therefore masked the data set by omitting any pixel in all the maps where any of the 176 data points in the time series had an intensity that fell below that threshold. There are similar artifacts close to the occulted region, which we remove by masking out pixels below $1.08 R_\odot$. These filters remove about 15% of the initial data, or about 9000 out of 66,000 pixels.

For the analysis of fluctuations, it is necessary to have data whose mean is zero and to remove any long-term trends. This was done for the velocity fluctuations, δv , by performing a linear fit to the time series at each pixel and then subtracting that linear trend. The intensity fluctuations were treated in the same way, except that we also normalized the data by the intensity trend so that we study $\delta I/I$, which we will denote as δi throughout.

For collisionally excited lines, the intensity is proportional to the square of the electron density n_e^2 (K. J. H. Phillips et al. 2008). Consequently, density fluctuations and intensity fluctuations are related by $\delta n_e/n_e \sim \delta i/2$. We can apply this analysis to our study here because although the level population of the Fe XIII 10747 Å line in the corona is primarily set by radiative decays from higher levels, those levels are collisionally excited. Hence the Fe XIII line intensity fluctuations will exhibit the previously noted relationship to density fluctuations.

3. Analysis and Results

3.1. Amplitudes

The power spectra and amplitudes of δv and δi can be inferred from the time series of data at each good pixel in the data set. In principle, the rms amplitude of the fluctuations could be determined by taking the square root of the variance of the time-series data. However, this overestimates the




Figure 1. Context image showing the average Fe XIII intensity observed by COMP during the observations. Labels highlight several apparent coronal structures: 1, 3, 4, and 7 label regions with relatively short closed loops; 2 and 5 indicate quiet Sun regions where the field lines are long and extend well beyond the field of view; in the vicinity of the label 6, there is an active region just behind the limb that leads to complex structures in the corona. The boxes surrounding 1 and 2 indicate the regions considered for the average power spectra discussed in Section 3.2.

amplitude due to the noise contribution from nonphysical sources of variation, such as detector noise, photon counting statistics, and variations due to stray light from other locations within the image. We applied a correction to mitigate some of these factors. The same analysis is used for both δv and δi , but for brevity we discuss the analysis using δv as an example.

First, we computed the discrete Fourier power spectrum (periodogram) at each pixel for $\delta v(t)$. The resulting power spectrum $P_{\delta v}(f_j)$ describes the variance in $v(t)$ due to oscillations at frequency f_j . The total $\sum_j P_{\delta v}(f_j)$ is equal to the variance of $\delta v(t)$, so the rms amplitude of the fluctuations is the square root of the summed power spectrum. In order to reduce the statistical uncertainty in the estimated $P_{\delta v}(f_j)$, the power spectrum from a given pixel at position (x, y) is averaged with that of its neighbors. This averaging produces a more precise estimate of the power spectrum but does not remove the contributions of noise.

In order to estimate the noise level, we assume that the physical sources of variation in $\delta v(t)$ are the same for the

reference pixel and its neighbors, and that variations from noise are random and uncorrelated so that we can decompose the measurements into a physical signal, $s(t)$, plus random noise, $e(t)$, giving $\delta v(t) = s(t) + e(t)$. Then, we can average the time series to obtain $\overline{\delta v(t)}$ whose variance from physical sources remains the same but for which the contributions to the variance from random noise are reduced by a factor of $1/K$, where $K \leq 9$ represents the number of pixels, including the reference pixel and up to 8 neighboring good pixels. This relation also holds at each frequency in the power spectrum, and we can estimate the noise power spectrum as $P_e(f_j) = (P_{\delta v}(f_j) - P_{\overline{\delta v}}(f_j))/(1 - 1/K)$. The inferred noise power spectrum, $P_e(f_j)$, did not exhibit a clear frequency dependence, so to reduce statistical uncertainties, we further assume the noise level to be constant (i.e., white noise) and take the average P_e over all frequencies. Finally, we subtract this P_e from $P_{\delta v}(f_j)$, sum the power over all the frequencies, and take the square root to obtain the rms amplitude of the δv




Figure 2. Same as Figure 1, but focusing on the East limb.

fluctuations. The noise estimate reduces the average inferred amplitude of δv fluctuations by about 0.1 km s^{-1} ($\approx 15\%$) and reduces the inferred amplitudes of δi fluctuations by about 0.004 ($\approx 20\%$), compared in each case to when no correction is made.

There are clearly some limitations to this method. First, we can only remove uncorrelated random sources of noise. Any noise sources that affect neighboring pixels in a similar way cannot be removed. For example, stray light fluctuations due to variations on the solar disk are likely similar in these off-limb pixels, and their effects are not removed. Additionally, averaging the time series reduces the amplitude of any fluctuations that are out of phase with one another. Most relevant to this work, any traveling waves will have slightly different phases in neighboring pixels along the wave path and so the inferred amplitude will be slightly reduced by the averaging.

R. J. Morton et al. (2016) used a different method to estimate the noise level from the COMP data. They smoothed the data using a three-point boxcar average over time and then subtracted the smoothed fluctuations from the original time series. The standard deviation of the residuals then gives an estimate of the noise level. We found that using their method

compared to our approach produces δv amplitudes that are smaller by on average 0.07 km s^{-1} and δi amplitudes that are smaller by about 0.001 . That the amplitudes are systematically smaller may be due to the boxcar-average method implicitly assuming that all high-frequency signal is due to noise. Nevertheless, these discrepancies are relatively small, $\approx 10\%$ of the amplitudes, and so the following results are not very sensitive to the way the noise level is estimated.

Figures 3 and 4 show maps of the rms amplitudes for the δv and δi fluctuations on the east limb, respectively. The δv amplitudes are found to be in the range of about $0.5\text{--}1 \text{ km s}^{-1}$ throughout most of the field of view, which agrees with previous measurements based on the COMP data (e.g., S. Tomczyk et al. 2007; S. Tomczyk & S. W. McIntosh 2009). Many of the δv fluctuations are due to Alfvénic waves, but these measurements likely underestimate the true wave amplitudes due to the implicit line-of-sight averaging (S. W. McIntosh & B. De Pontieu 2012).

The δv amplitudes appear to be larger near the poles. One possibility is that this represents a difference in the characteristics of the waves on the long field lines bordering the coronal holes near the poles, as compared to the shorter loops found at lower latitudes. The high formation temperature for Fe XIII makes it unlikely that any of the emission observed by COMP comes from open field regions. As such, it seems more likely that the increased amplitudes at the poles is due to line-of-sight effects, such as a change in the average viewing angle of the line of sight with respect to the magnetic field. We may be observing at the poles a larger contribution from δv fluctuations in the parallel flows along the field, rather than the transverse waves that likely predominate in the lower-latitude parts of the observations. Alternatively, there may be less line-of-sight averaging near the poles because the depth of Fe XIII emitting material along the line of sight is smaller due to the presence of the cooler coronal hole material. As line-of-sight averaging is known to reduce the inferred amplitudes, having systematically less averaging at the poles would result in a systematically larger inferred amplitude.

The δi amplitudes in the low corona appear to be typically on the order of a few percent and increase systematically with height. However, these amplitudes are likely to be systematically underestimated at the largest heights, because we have not accounted for scattered light. The relative intensity fluctuation δi is the intensity fluctuation δI normalized by the time-averaged intensity at that position, but this average intensity has not been corrected for scattered light. Thus the normalization factor is overestimated and the δi amplitudes are likely underestimated. The contribution of scattered light to the total intensity is expected to grow with height, so that the underestimate in the δi amplitudes also grows with height.

The largest δi amplitudes also appear to be associated with the tops of coronal loops. This is partly due to the general trend that δi grows with height, but may also be due to turbulence. Previous observations by I. De Moortel et al. (2014) showed that the power spectrum of δv fluctuations was broader at loop tops and suggested that this was due to turbulence. Turbulence might also lead to an increase in density fluctuations. Density fluctuation amplitudes in open field regions have been observed to increase with height, and that increase was associated with the apparent damping of the Alfvénic waves (M. Hahn et al. 2018).




Figure 3. Amplitudes of δv fluctuations.

3.2. Power Spectra

In order to compare the typical power spectra for δv and δi fluctuations, we have averaged the power spectra over the boxed regions in Figure 1: region 1 with short loops and region 2 with long radial structures. To perform this averaging, we computed the power spectra at each pixel within each region. Histograms of the statistical distribution of the measured power within a given frequency bin showed that the results appeared to follow an exponential distribution. This is the expected statistical distribution for the power at a given frequency (S. Vaughan 2005). The reason is that the real and imaginary

parts of the discrete Fourier transform are expected to follow a normal distribution and so their squared sum, the power, follows a chi-squared distribution with two degrees of freedom, which is an exponential distribution. We then take the average power derived from all of those pixels within each frequency bin. Finally, we subtracted the noise power level that was inferred using the boxcar smoothing method of R. J. Morton et al. (2016) discussed previously in Section 3.1. This noise level was about $2 \times 10^{-3} \text{ km}^2 \text{ s}^{-2}$ for the δv fluctuations and 2×10^{-6} for the δi fluctuations. The exponential distribution has the property that the standard deviation is equal to the




Figure 4. Same as Figure 3, but for δi fluctuations.

mean, so the uncertainty on each power spectrum is essentially 100% (see, e.g., S. Vaughan 2005). However, the statistical uncertainty of the average power at a given frequency is the power at that frequency divided by $\sqrt{N}$, where N is the number of samples. Here $N = 3778$ samples for region 1 and $N = 5571$ samples for region 2, so the statistical uncertainties are very small. However, this does not account for pixel-to-pixel variations in the power spectrum, which is likely the dominant source of uncertainty in the final results.

Figure 5 shows the average power spectra for the δv fluctuations. The power spectrum for δv consists of a power-

law continuum with a hump near $f \approx 3.5$ mHz. The power spectra from the short loops and the long radial structures are nearly identical, both in shape and in magnitude. These results are also very similar to previous COMP measurements (e.g., S. Tomczyk & S. W. McIntosh 2009; R. J. Morton et al. 2016).

The power spectra for the δi fluctuations are shown in Figure 6. Here we find a rapid decay for low frequencies that levels off to a nearly flat white noise type spectrum above about 2 mHz. One possible reason for this difference in the overall shape of the power spectrum between δv versus δi fluctuations is that the δv fluctuations may be due largely to wavelike




Figure 5. Average power spectra for δv fluctuations in the regions labeled 1 and 2 in Figures 1 and 2, which correspond to short closed loops and long radial structures, respectively.

oscillations dominated by a few frequencies, whereas the δi fluctuations are caused by many underlying physical processes, such as waves, heating, flows, and so forth, occurring on a variety of timescales. The overall shape of the δi power spectrum is very similar between the short and long loops, except near 3 mHz. At that frequency, the short closed loops show a clear peak, whereas this peak is much less prominent in the region with the long radial structures. To verify this, we performed the same analysis using data in the vicinity of regions 4 and 5 in Figure 1, containing short versus long loops, respectively. Those regions also exhibited the relationship that the 3 mHz peak in the δi fluctuations is broader and shallower on the long streamers.

One possible interpretation of this result is that the 3 mHz power represents the perturbations at the base of the corona that excite both Alfvénic waves and density fluctuations. On the short loops, power at this frequency is injected at both ends of the loop and the influence of the boundary remains evident throughout the loop. On the longer structures, the excited Alfvénic wave power evidently does not undergo a significant change, but the density fluctuation power spectrum evolves. This may occur through damping of the initially excited acoustic modes and the excitation of new density fluctuations with different characteristics at larger heights, such as through the parametric decay of Alfvén waves or other processes.

3.3. Frequency Scaling

We next examine whether there is a frequency-scaling relationship between the δv and δi fluctuations. M. Hahn et al. (2022a) found that the power spectra of δv and δi fluctuations in the transition region were related to one another by a scaling of the frequency axis. They considered this a signature of PDI, interpreting the frequency-scaling factor to represent the ratio of the pump Alfvén wave frequency to the frequency of the PDI resonantly excited slow mode wave. Here we used the

same method to determine the scaling factor, α , by which the frequency axis of the δi power spectrum $P_{\delta i}(\alpha f)$ should be scaled to maximize the cross correlation with the δv power spectrum $P_{\delta v}(f)$. For each pixel in the data set, we have averaged the power spectra over all the neighboring good pixels, which is typically an average over $K = 9$ pixels.

Figure 7 shows the scaling factors between the δi and δv fluctuations. Here we have suppressed any pixels where the maximum correlation coefficient was below 0.7, which is about 40% of the unmasked pixels. The cutoff of 0.7 was chosen because it corresponds to a p -value of about 0.05 (i.e., the probability that the correlation arose by chance is estimated to be less than 5%). However, that estimate relies on the idealized condition that the underlying correlated variables followed a normal distribution. Here we consider it only as a way for removing distracting points from the plot of the scaling factor that most likely do not correspond to actual correlations.

We do not find a clear relationship between intensity and velocity fluctuations. We commonly find α values close to unity, especially in the coherent closed loop structure seen in the lower left quadrant. One possibility is that these oscillations represent the velocity and compressive fluctuations associated with the same wave mode. Somewhat larger values of $\alpha \approx 1.5$ –2 are seen on the long field lines at lower latitudes. This relationship might be consistent with PDI, which is discussed further in Section 4.

3.4. Phase Relationships

The phase relationship between the δv and δi fluctuations also provides clues as to the nature of these fluctuations. If the fluctuations arise from the same underlying wave mode, we expect δv and δi to have a definite phase difference. For example, a fast sausage mode observed perpendicular to the magnetic field should exhibit a phase difference of $\pi/2$ between the velocity and density oscillations




Figure 6. Same as Figure 5, but for δi fluctuations.

(M. Gruszecki et al. 2012), or a slow mode wave observed parallel to the field should exhibit a phase difference of zero (e.g., L. Ofman et al. 2000).

We have examined the phase differences between the fluctuations by computing the Fourier cross-power spectrum between δv and δi in each pixel (O. R. White & R. G. Athay 1979). The phase differences $\phi(x, y, f)$ are derived from the real and imaginary parts of the cross-power spectrum at each frequency and at each pixel. As discussed previously, the Fourier coefficients have large uncertainties, so an averaging is needed in order to assess whether the phase difference is significant. For this, we assume that the relation between the δv and δi fluctuations varies slowly in space so that the phase differences should be similar among a given pixel and its eight neighboring pixels. Then we compute the average phase $\bar{\phi}$ over the central pixel and up to eight good neighboring pixels and a quantity that measures the dispersion of these ϕ values.

The analysis of the phase angle requires special treatment because ϕ is cyclical, so one cannot simply take the standard arithmetic mean. Instead, we find the mean by averaging the unit vectors corresponding to each value of ϕ (K. V. Mardia & P. E. Jupp 2000). We compute the mean value of $\bar{C} = \langle \cos(\phi) \rangle$ and $\bar{S} = \langle \sin(\phi) \rangle$, where the angle brackets denote the arithmetic mean. Then the average value of the phase difference is $\bar{\phi} = \arctan(\bar{S}/\bar{C})$, producing angles in the range of $-\pi$ to π . The concentration of the angles is described by the mean resultant length $\bar{R} = (\bar{C}^2 + \bar{S}^2)^{1/2}$, which is a number between 0 and 1. That is, if all the angles in the set over which the average is taken are in the same direction, then $\bar{R} = 1$, and if they are all randomly distributed, $\bar{R} = 0$. It is useful to define instead the dispersion $\bar{V} = 1 - \bar{R}$, also called the circular variance, which corresponds to the more familiar interpretation that this number represents a spread of values so that a smaller number suggests a stronger relationship.

We have applied the previously noted steps to compute $\bar{\phi}(x, y, f)$ and $\bar{V}(x, y, f)$, where the average is taken over each pixel and the up to eight neighboring good pixels. Figure 8 presents a map of the dispersion, $\bar{V}$ at 3.5 mHz, corresponding to the peak in the δv and δi fluctuation power spectra. We find that there are low-dispersion structures extending out from the base of the corona that appear to largely follow the magnetic field lines. At larger heights the dispersion is large, corresponding to essentially random phases. This may be due to increasing scatter in the data from noise or to physical processes that cause the break up of coherent wave modes. Even at low heights, though, there are many locations with a high dispersion. The clear spatial structuring of the $\bar{V}$ that aligns with observed coronal structures supports that there are strong phase relationships between δv and δi fluctuations.

Figure 9 shows the absolute value of the measured phase angles $|\bar{\phi}(x, y)|$ at the same frequency of 3.5 mHz. The reason for plotting the absolute value is that $\phi = -\pi$ and $\phi = \pi$ correspond to the same phase difference. In this plot we have suppressed locations with a high dispersion by plotting only those pixels where $\bar{V} \leq 0.4$. This cutoff was chosen as it corresponds, approximately, to a p -value of 0.05 (K. V. Mardia & P. E. Jupp 2000). Although there is a great deal of scatter in the data, there appears to be a tendency for the phase values near the equator to be $\phi \sim \pi/2$, whereas at middle to high latitudes, $\phi \sim 0$ is more common. In terms of coronal structures, the equatorial region contains more long field lines, whereas the middle to high latitudes have the shorter closed loops.

3.5. Wave Tracking

For many properties of the waves, it is necessary to measure the fluctuations in both time and space along the path of the waves. Thus a first step in such an analysis is to determine the trajectories of the waves. We used a method based on




Figure 7. Frequency-scaling factor between δi and $\delta \nu$ fluctuations.

measuring the cross correlation of the fluctuations to determine the direction of wave travel at each pixel. Once the directions are known, then the wave paths in the plane of the sky can be found by stepping along the trajectories using Euler's method. In the course of this analysis, we also determine the parallel and perpendicular correlation lengths.

Our method for determining the wave direction is based on the correlation method presented by S. Tomczyk et al. (2007). Our measured power spectra (Figure 5), and previous work by others, has shown $\delta \nu$ fluctuations in the corona tend to have a peak power around 3.5 mHz, and we found δi fluctuations to

have a peak at a similar low frequency. So, we first apply a low-pass Butterworth filter to the data with a cutoff above 7 mHz to reduce the impact of high frequency noise. (Note that this filter was only applied for this wave tracking step and not in any of the following analyses of the power spectra along the wave path.) Then, for each pixel, we compute the cross correlation c_{ij} between that pixel i and all of its neighbors j within ± 12 pixels ($\pm 54''$). This reveals an oval island of high cross correlation with a long axis oriented along the direction of wave travel. The ± 12 pixel limit was selected because it was found to be large enough to contain these islands without being




Figure 8. Dispersion, $\bar{V}$, of the phase differences between the $\delta\nu$ and δi fluctuations at 3.5 mHz. A dispersion of 0 indicates that the phases of fluctuations at neighboring pixels were the same, while a dispersion of 1 indicates that the phases of neighboring pixels were all different.

so large as to be influenced by random cross correlations with distant structures. To quantify the travel direction, we can find the slope of the line that will minimize the square of the perpendicular distance to each point in the island above a threshold high correlation. Here we defined a high correlation to be a correlation coefficient >0.6 . This value was chosen by inspection, as it results in islands with a large enough number of pixels to estimate the orientation consistently while also providing a reasonable noise cutoff. The computation of the best-fit line was performed using a principal component analysis. That is, we computed the covariance matrix for the

coordinates, and then found the eigenvalues and eigenvectors of the covariance matrix. The direction parallel to the wave path is given by the eigenvector corresponding to the largest eigenvalue, and the other eigenvector gives the perpendicular direction.

This process was repeated for every pixel in the COMP data to build up a map of the wave travel directions. We can track waves through the map by selecting a starting pixel and then using Euler’s method to step forward and backward along the path to solve for the entire wave trajectory. Figure 10 shows examples of wave trajectories traced starting from initial




Figure 9. Nearest-neighbor-averaged phase differences $\bar{\phi}$ between δv and δi fluctuations at 3.5 mHz for locations with a statistically significant low dispersion.

positions located at a radius of $r = 1.1 R_{\odot}$ and varying polar angles, θ . For clarity, the plot illustrates wavetracks starting only from every fifth pixel along that arc. The resulting wave trajectories show good agreement with the apparent shapes of coronal structures seen in Figures 1 and 2, which is evidence that the method is accurately tracking the waves. One exception is at the west limb, near the active regions, where the wave trajectories take on a jagged appearance due to the complex intervening structures along the line of sight.

We will compare our results to several other publications also based on COMP, so it would be useful to discuss some of

the differences between the wave tracking method used here and those used by others. Our method is most similar to that of S. Tomczyk et al. (2007), except that S. Tomczyk et al. (2007) Fourier filtered the time series using a Gaussian filter centered at 3.5 mHz, whereas we have only used a low-pass filter. In a later paper, S. Tomczyk & S. W. McIntosh (2009) used a coherence-based method (S. W. McIntosh et al. 2008), and the same method was later used by R. Sharma & R. J. Morton (2023). The coherence between the time series is computed using the cross spectrum. Because the raw cross spectrum of two time series will always give a coherence of unity




Figure 10. Wavetracks derived from the $\delta\nu$ fluctuations. Each of these paths were traced from a starting point at $1.1 R_{\odot}$. We highlight one trajectory on the lower left in blue and use this path as an example throughout the paper in describing our methods.

(G. M. Jenkins & D. G. Watts 1968), it is necessary to average the cross spectra in some way to obtain a meaningful coherence value. For the previously noted studies, this averaging takes place at the step where the Gaussian filter is applied. The reason we have not adopted this approach is that our objective is to understand the relationship between the $\delta\nu$ and δi fluctuations, and we posit that this might involve a frequency scaling between the waves. As such, we avoided assuming that the waves have a particular frequency by not filtering the data at 3.5 mHz. A drawback is that noise may be suppressed less here than in those other studies.

3.6. Perpendicular Correlation Lengths

The perpendicular correlation length can be quantified directly from the constructed cross-correlation calculations. This length represents the extent to which the plasma is perturbed in the perpendicular direction due to the wave. It is also related to the size of the source of wave excitation, analogous to the size of the antenna in a laboratory plasma. In

the context of turbulence, it may be thought of as the largest scale eddy.

To quantify this perpendicular correlation length, we found the characteristic transverse width of the island of high correlation by fitting it with a Gaussian. The island of high correlation is peaked in the center, with the central pixel having correlation coefficient $c_{ii} = 1$, and declines toward the boundaries. This island can be modeled as a two-dimensional Gaussian distribution (e.g., R. Sharma & R. J. Morton 2023) and the width characterized by the standard deviations of the parallel and perpendicular axes, $\sigma_{\parallel}$ and $\sigma_{\perp}$. To determine the value of $\sigma_{\perp}$, we summed the c_{ij} along the parallel direction and fit the resulting distribution of c_{ij} versus perpendicular distance with a one-dimensional Gaussian. Since the integral of a two-dimensional Gaussian along one axis preserves the σ for the other axis, this summation reduces the dimensionality of the problem without affecting $\sigma_{\perp}$. R. Sharma & R. J. Morton (2023) noted that solar wind studies of perpendicular correlation lengths report the exponential $1/e$ length scales given by an autocorrelation function, and so in order to better compare to

these values, we likewise scale our Gaussian correlation lengths by a factor of $\sqrt{2}$ defining $L_{\perp} = \sqrt{2}\sigma_{\perp}$.

The perpendicular correlation lengths for δv take a broad range of values and a histogram resembles a lognormal distribution, as has been found by R. Sharma & R. J. Morton (2023). We found the median $L_{\perp} = 7.0$ Mm and the most probable value (i.e., the mode) is $L_{\perp} = 5.0$ Mm. This is somewhat smaller than found by R. Sharma & R. J. Morton (2023), who found that the most probable values were in the range of $L_{\perp} = 7.6$ – 9.3 Mm among several observations. One possible reason for the discrepancy is that R. Sharma & R. J. Morton (2023) focus specifically on fluctuations at frequencies of 3.5 mHz, whereas our data is less strongly filtered. Given the broad statistical distribution of $L_{\perp}$ values in the measurements, there is still substantial overlap in these measurements and the discrepancy is small.

The $L_{\perp}$ map derived from δv (Figure 11) shows relatively low $L_{\perp}$ areas ($L_{\perp} \lesssim 9$ Mm) that appear to follow the coronal loops in the images. There are clumps of high $L_{\perp}$, particularly at low heights and also the active regions on the west limb (not shown in the figure). One possible interpretation is that the low $L_{\perp}$ represents Alfvénic waves traveling along individual magnetic flux tubes, while the high $L_{\perp}$ represents structures undergoing large-scale oscillations so that $L_{\perp}$ is due to the entire structure being generally excited rather than to characteristics of a traveling wave mode. We found no discernable trend of $L_{\perp}$ varying with height.

With respect to intensity fluctuations (Figure 12), the range of $L_{\perp}$ derived from δi was even broader than that for δv , having a median value of $L_{\perp} = 10.8$ Mm and a mode of $L_{\perp} = 3.7$ Mm, though still comparable to the range found for the δv . The spatial distribution of $L_{\perp}$ for the δi fluctuations shows broader structures than that for δv . These may represent regions affected by large-scale intensity changes as those structures undergo heating and cooling throughout. Those processes may not lead to similar signatures in the Doppler velocities, because any associated flows are likely to be parallel to the magnetic field and so perpendicular to the line of sight at the limb, resulting in minimal Doppler shifts.

It is interesting that $L_{\perp}$ for both δv and δi fluctuations are an order of magnitude larger than the apparent widths of coronal loops. Analysis of loops in high-spatial-resolution images has shown that the typical widths of coronal loops are ≈ 0.5 – 1.5 Mm (H. Peter et al. 2013; M. J. Aschwanden & H. Peter 2017). One might have expected that phase mixing at the boundary of coronal loops would interrupt the correlations leading to wave patterns and loops having similar size distributions. One possible reason this does not occur comes from simulations of Alfvén waves on multi-stranded loops, showing that strands can influence their neighbors, which would result in $L_{\perp}$ larger than the scale of the underlying structures (M. Guo et al. 2019). Another possibility is that the images of loops may not be an accurate representation of the true coronal density structure (e.g., A. Malanushenko et al. 2022).

Similarly, we also expect to find small values of $L_{\perp}$ if the fluctuations are related to turbulence. Turbulence is thought to be important for the generation and evolution of Alfvénic waves in the solar corona. If the fluid motions that generate the Alfvénic waves are turbulent, we would expect waves to be excited with a broad range of perpendicular length scales. Moreover, if turbulence occurs in the corona, then we would

expect energy initially excited at large $L_{\perp}$ to cascade to small $L_{\perp}$. A potential physical explanation for why small values of $L_{\perp}$ are not observed is that small values of $L_{\perp}$ are efficiently dissipated.

It is possible, though, that we observe few small values of $L_{\perp}$ because the measurements have limitations from spatial resolution, noise, and line-of-sight averaging. To test this, we have studied in more detail the distribution of our $L_{\perp}$ measurements. R. Sharma & R. J. Morton (2023) commented that the distribution of $L_{\perp}$ values in the COMP data appear to approximately follow a lognormal distribution, and we find a similar distribution here. In the Appendix we investigate the statistical distribution in more detail with the aim to determine whether there may be small $L_{\perp}$ values that are not resolved. In particular, we consider the hypothesis that the actual distribution of $L_{\perp}$ in the corona follows a power law, and argue that the addition of noise and line-of-sight averaging can distort the measurements, producing a histogram that can match the observed one. The implication is that known systematic issues are capable of biasing our $L_{\perp}$ measurements to larger values. However, the arguments are speculative (which is why we relegate them to the Appendix).

Future work to understand the perpendicular correlation lengths will require higher spatial resolution measurements and a deeper understanding of systematic uncertainties in the observations. We plan to analyze measurements with higher spatial resolution instruments, such as the Daniel K. Inouye Solar Telescope (DKIST; M. P. Rast et al. 2021), to address some of these limitations.

3.7. Time-Distance Plots and k – ω Diagrams

We can measure properties of the wave as they propagate along the trajectories identified in Figure 10, allowing us to infer the wave speed and direction. To do this, we first combine the time series of fluctuations at each pixel along the wave path to construct time-distance plots $\delta i(t, s)$ and $\delta v(t, s)$, where s is the distance along the trajectory.

Next, we took the Fourier transform of the time-distance plot and constructed a wavenumber-frequency k – ω diagram. Note that it is conventional to use angular frequency ω for such a plot, even though elsewhere we usually refer to the oscillation frequency f . Figure 13 shows the k – ω diagram for δv fluctuations along the wave trace indicated in Figure 10. Here we have used the sign convention that negative ω corresponds to anti-sunward (upward) traveling waves and positive ω corresponds to sunward (downward) traveling waves.

In Figure 13 the wave power appears concentrated along two ridges. Since the phase speed of a wave is given by ω/k , these ridges have a slope that is the inverse of the wave phase speed. The dashed lines in the figure correspond to 1 Mm s^{-1} , which is a typical Alfvén speed in the corona. The dashed lines were drawn by hand to “guide the eye,” yet it does align roughly with a ridge in the power spectrum, suggesting that the δv fluctuations are due to Alfvénic waves, such as kink or torsional Alfvén modes. More precise measurements of the phase speed are presented in Section 3.9.

Figure 14 plots the k – ω diagram derived for δi fluctuations along the same trajectory. Dashed lines on the plot correspond to wave speeds of 1 and 0.2 Mm s^{-1} , representing typical Alfvén and sound speeds, respectively. Although there is some δi power along and between these lines, we do not find distinct ridges indicative of a particular wave mode. The majority of the




Figure 11. Map of $L_{\perp}$ for the δv fluctuations.

power is concentrated at low ω and low k , which likely represents long-term changes in the intensity of large-scale structures due to heating, cooling, or other evolution. That the power does not align with a clear wave mode suggests that much of the δi variation observed in the corona is not due to waves, although waves are likely also present.

3.8. Upward versus Downward Waves

We can estimate the balance of upward versus downward waves by comparing the total power P at negative versus positive frequencies, respectively, having finite wavenumber $k > 0$. We

will quantify this ratio by defining $\eta = (P_{\text{up}} - P_{\text{down}}) / (P_{\text{up}} + P_{\text{down}})$. Because noise power will affect the denominator in η , we first estimate the noise level and subtract it from each pixel. The k - ω diagram suggests that the wave power is concentrated in the ridges at low k and that the power at high k and high ω is likely due to noise. So we estimate the noise in each k and ω bin as being the average power in upper corners of the diagram containing the highest 10% of wavenumbers and angular frequencies. This value is then subtracted from each bin of the diagram before computing η . For the example shown in Figure 13, we find $\eta_{\delta v} = 0.24$, indicating predominantly upward-propagating




Figure 12. Map of $L_{\perp}$ for the δi fluctuations.

waves. We compute the imbalance of upward versus downward δi power in the same way. Similarly, for the δi fluctuations shown in Figure 14, we find $\eta_{\delta i} = 0.07$, indicating that traveling δi fluctuations tend to be moving upward. The magnitude of $\eta_{\delta i}$ is smaller than $\eta_{\delta v}$, but this is likely because a significant amount of the δi fluctuation power is not in a traveling wave.

Figure 15 shows a polar plot illustrating the wave power imbalance for all wave trajectories starting at $1.1 R_{\odot}$ like those in Figure 10. Red symbols indicate the imbalance for velocity fluctuations, $\eta_{\delta v}$, and blue symbols for intensity fluctuations, $\eta_{\delta i}$. The solid circle indicates zero, so symbols at larger radii

indicate predominantly upward-traveling waves while those at smaller radii indicate predominantly downward-traveling waves. For scale, the dotted circles are drawn at intervals of $\Delta\eta = 0.1$.

We find that $\eta_{\delta v}$ and $\eta_{\delta i}$ are usually positive. Although there appear to be exceptions at several locations, Figure 10 shows that these locations all correspond to trajectories that travel along the top of a short coronal loop. In such cases, our definition of “up” and “down” along the wave trajectory does not have a clear interpretation as being toward or away from the Sun. However, it is interesting that such loops do show a clear




Figure 13. k - ω diagram for the δv fluctuations along the trajectory shown in Figure 10. The dashed lines indicate a wave speed of 1 Mm s^{-1} , which is a typical Alfvén speed in the corona.

imbalance. One possibility is that this indicates a stronger wave driving from one footpoint of the loop compared with the other. We also find that $\eta_{\delta v}$ and $\eta_{\delta i}$ are strongly correlated with one another. Quantitatively, their correlation has a Spearman coefficient of 0.8. We find average magnitudes of $\eta_{\delta v} = 0.14$ and $\eta_{\delta i} = 0.09$, indicating waves are traveling predominantly, though not exclusively, away from the Sun. The smaller value for $\eta_{\delta i}$ is likely due to δi having relatively more power in fluctuations that are not traveling compared with δv .

3.9. Wave Speeds and Plasma Beta

We have determined the phase speeds of the waves along the inferred trajectories by separating the upward versus downward waves and using a cross-correlation method to determine the lag time between different positions along the path, following S. Tomczyk & S. W. McIntosh (2009). The separation of the upward versus downward waves is necessary because the superposition of counter-propagating waves of equal amplitude will form a standing wave, which the cross-correlation method will see as having an infinite phase speed. Even if the amplitudes are not equal, the phase speed will be biased. Thus the first step for determining the phase speeds is to construct a time-distance map that selects only one direction of waves. We do this by performing the inverse Fourier transform of the k - ω diagram using either only the negative (outward) or positive (inward) wave ω . Next we select a reference height i near the center of the wave trajectory. We then find the cross correlation of the fluctuations at each other height j with the reference as a function of lag time $c_{ij}(\tau)$. The cross correlation as a function of the lag time is fit with a parabola to find the lag time τ_{ij} that maximizes c_{ij} . Finally, we fit a linear function to the plot of lag times τ_{ij} versus distance from the reference pixel Δs_{ij} and obtain the speed as the slope of that line. Figure 16 shows an example of such a fit for upward δv fluctuations along the trajectory highlighted in Figure 10. Note that this method is

only sensitive to the components of the velocity in the plane of sky.

We applied the lag time analysis to all of the wave trajectories that pass through $1.1 R_{\odot}$. For both δi and δv fluctuations, there was no significant difference in the phase speeds in the upward versus downward direction. This makes sense, as the mean flow speeds at these low heights are expected to be small, so there is no physical reason for the measured fluctuation speeds to differ. Earlier we explored how there is more power in the upward fluctuations, and thus the upward signals are relatively less affected by noise and their speeds are typically measured with greater precision. For that reason, the rest of the discussion here will be based only on upward fluctuation speeds.

Figure 17 shows a polar plot of the measured wave speeds. The red symbols correspond to δv fluctuations and the blue symbols to δi fluctuations. The speed is proportional to the radial distance from the solid circle, which represents zero, and the dotted circles indicate intervals of 200 km s^{-1} . Because we are measuring the motions of the fluctuations as they propagate through the image, these speeds correspond only to the plane-of-sky velocity. The actual propagation speeds are necessarily greater than those we have inferred because we are ignorant of the line-of-sight component of the speed. However, as δv and δi appear to propagate in the same direction, it remains valid to compare their speeds to one another.

The relationship between the speeds of δv versus δi fluctuations appears to fall into two classes, depending on whether the fluctuations are traveling along short closed loops or long paths that do not close within the bounds of our measurements (refer to Figure 10). The short closed loops are found at middle latitudes. In these locations, the speeds of the δv and δi fluctuations are nearly the same. This suggests that both types of fluctuations in these locations reflect the same




Figure 14. k - ω diagram for the δi fluctuations along the trajectory shown in Figure 10. The dashed lines indicate wave speed of 1 and 0.2 Mm s⁻¹, corresponding to typical Alfvén and sound speeds in the corona.

underlying wave mode, such as compressive modes along these loops.

The long open trajectories, seen most clearly at moderately low latitudes on the east limb, exhibit significant differences in the wave speeds of δv versus δi fluctuations. The speeds of the δv fluctuations are greater by several hundred km s⁻¹ than those of the δi fluctuations. This suggests that the δv fluctuations here correspond to Alfvénic waves and the δi fluctuations are dominated by slow mode acoustic waves. If we adopt this interpretation, then we may estimate the plasma β (i.e., the ratio of the fluid to magnetic pressure) from the ratio of the observed sound and Alfvén speeds as $\beta \sim c_s^2/V_A^2$. The median value of β for these open trajectories is then about $\beta \approx 0.5$.

4. Discussion

The δv fluctuations observed in the corona appear to be predominantly in the form of Alfvénic waves, as has been previously recognized (S. Tomczyk et al. 2007). Fourier analysis of the time-distance maps of these waves (Figure 13) suggests that they obey a linear dispersion relation with propagation speeds of $\gtrsim 500$ km s⁻¹, which are consistent with Alfvénic waves.

The physical significance of the δi fluctuations is less clear. They likely arise due to multiple phenomena. In some regions, especially relatively short closed loops, the dominant δi fluctuations appear to have similar speeds and frequencies to the δv fluctuations. In these cases, δi may be due to density variations associated with fast compressional modes. In other regions, corresponding usually to long trajectories, we find that the δi fluctuations propagate slower than δv fluctuations and appear at different frequencies. In these cases, much of the δi fluctuation power may come from slow mode acoustic waves. However, even for these open trajectories, the k - ω diagrams for δi fluctuations do not resolve into a clear dispersion relation (e.g., Figure 14). In all regions of the corona, a significant

amount of density fluctuation power is likely due to erratic changes that are not associated with waves. Such sources may include eruptions, flows, and impulsive heating events. This ambiguity adds to the well-known problem of distinguishing slow mode waves from quasi-periodic pulsations (see D. Banerjee et al. 2021 and references therein).

In describing these characteristics, we have noted that there is an apparent distinction between the behavior of fluctuations on short loops and open trajectories. By short loops we mean, essentially, those whose maximum height is around $1.3 R_\odot$ so that the top of the loop falls within our field of view. The long trajectories are those regions where the fluctuations appear to propagate roughly radially from the base of the corona to the outer radius of our available data. It is interesting to note that these open trajectories do not correspond to open magnetic field lines. Rather, they appear to correspond to the bases of long loops whose loop tops are well beyond the field of view of the COMP data. Both of these structures are expected to have plasma properties typical of the quiet Sun, and so the apparent differences in the behavior of fluctuations along these structures is likely related to the loop length.

Much previous work has distinguished between the behavior of waves on open field lines versus closed field regions, but our observations suggest that there is an intermediate behavior on long loops. The reasons for this are not yet clear and should be investigated in future work. However, we can speculate on several explanations. One possibility is that the behavior is related to the ability to form standing modes on the loops. In this case, it may be that the short loops are capable of supporting low-order resonances, but for the longer loops the dominant wave frequencies are non-resonant. Another possibility is that the long loops are open from the waves' perspective because of dissipation at large heights. That is, on the shorter loops waves from opposite footpoints still retain significant wave power at the locations where they collide with




Figure 15. Polar plot illustrating $\eta_{\delta\nu}$ (red symbols) and $\eta_{\delta i}$ (blue symbols) for wave trajectories at varying polar angles with starting points at $1.1 R_{\odot}$. The solid circle indicates $\eta = 0$ with larger radii corresponding to $\eta > 0$ and smaller radii $\eta < 0$. Dotted circles are drawn at intervals of $\Delta\eta = 0.1$ for scale. The background image is the average Fe XIII intensity as in Figure 1.

one another. However, on the long loops, waves from opposite footpoints lose most of their energy before they encounter waves propagating up from the opposite footpoint. From the perspective of the waves, such a long loop would be similar to an open field line.

Even on the long trajectories, we found that there is a substantial fraction of Alfvénic wave energy that is traveling toward the Sun, as can be seen from the positive frequency half of the k - ω diagrams. Our estimated average $\eta_{\delta\nu} = 0.14$ corresponds to up to 40% of the wave energy traveling toward the Sun. However, this is likely an overestimate because it is sensitive to noise level in estimating the normalization for η . If it is the case that the injected wave energy from the footpoints is dissipated before it can reach the other half of the loop, then the downward wave energy must be generated within the loop as a reflection or instability of the incident upward-propagating waves. Many of the possible processes that would cause this involve interactions between the Alfvénic waves and density fluctuations, either because the density fluctuations enhance the wave reflection or due to an interaction through PDI.

One objective of this investigation was to determine the nature of the relationship between the $\delta\nu$ and δi fluctuations. We were particularly interested in seeing whether PDI is a common process in the corona, but the data here are ambiguous. This is largely because the δi fluctuations appear to be caused by several different physical processes, which impedes our ability to make a direct comparison with $\delta\nu$ fluctuations, as was done by M. Hahn et al. (2022a). Tentatively, it appears that PDI is unimportant on the shorter closed loops, whose dynamics we hypothesize to be more dominated by Alfvénic wave collisions between the primary Alfvénic waves traveling up from the footpoints. This is consistent with the relevant timescales as well. The Alfvén waves have periods of about 5 minutes, and for the measured Alfvén speed the travel time from one footpoint to the other of a short ~ 500 Mm loop is about 20 minutes, or a few wave periods. The PDI growth rate is expected to be a small fraction of the pump wave frequency (see, e.g., M. Hahn et al. 2022a), leading to growth timescales of hours. Thus Alfvén wave collisions are probably a more important effect than PDI on short loops. In contrast, there do appear to be frequency-scaling relations on the




Figure 16. Lag time vs. distance for δv fluctuations along the trajectory highlighted in Figure 10. The vertical error bars are the uncertainties in the lag time from the parabolic fit to the cross correlation $c_{ij}(\tau)$.

longer wave trajectories that are suggestive of PDI, which would very likely also apply to open field regions. Future measurements should aim to verify this by focusing on open structures.

Additionally, our analysis of the cross-power spectrum showed that there are structures where the δi and δv fluctuations at a given frequency show a definite phase difference. This suggests that they arise from the same underlying wave mode. Although there is much scatter in the data, the dichotomy between long and short loops appears also to be true in these measurements, where there is a tendency for the longer loops to exhibit phase differences of $\phi \sim \pi/2$ versus short loops having phase differences near zero.

5. Summary

We have investigated the general properties of δv and δi fluctuations in the quiet Sun corona and their interactions with one another. The δv fluctuations have a distinct dispersion relation in a k - ω diagram, indicating that the majority of those δv fluctuations are Alfvénic waves. We find typical amplitudes of 0.5 – 1 km s^{-1} growing with height, at least for the observed heights of 1.1 – $1.3 R_{\odot}$. The δv fluctuations have median perpendicular correlation lengths of $\approx 7 \text{ Mm}$. They are predominantly in the outward direction with typically $\eta_{\delta v} \approx 0.14$. These values are in agreement with previous results for δv fluctuations observed by COMP (e.g., S. Tomczyk et al. 2007; S. Tomczyk & S. W. McIntosh 2009; R. J. Morton et al. 2016; R. Sharma & R. J. Morton 2023).

The δi fluctuations are a proxy for density fluctuations, but they appear to arise from several physical processes, including possibly compressional fast waves, slow mode acoustic waves, and aperiodic variations. The δi fluctuations have typical amplitudes of a few percent, which also grow with height over the observed field of view. They have median perpendicular correlation lengths of about 11 Mm , though with a broad distribution so that this length is not significantly different from that of the δv fluctuations.

The δi fluctuation is also predominantly in the outward direction with $\eta_{\delta i} \approx 0.1$. However, $\eta_{\delta i}$ is correlated with $\eta_{\delta v}$. The apparently smaller magnitude of $\eta_{\delta i}$ compared to $\eta_{\delta v}$ is likely because much of δi fluctuation power is not in the form of propagating waves.

The relation between δv and δi is complex and differs depending on the underlying coronal structures. On short closed loops, we find that power spectra of both types of fluctuations are very similar and that they have similar propagation speeds. Examining the phase relationship at 3.5 mHz on these loops suggests that they have a phase difference near 0. These properties may indicate that both fluctuations are due to the same underlying wave mode.

On the long quiet Sun structures oriented radially within our data, the δv and δi fluctuations have different power spectra, which in some cases can be related by a frequency-scaling factor. They also have significantly different speeds, with the δv fluctuations traveling hundreds of km s^{-1} faster than the δi fluctuations. These longer structures show evidence for wave reflection or PDI, which we plan to investigate in more detail in future work. We also find evidence for another relationship on the long field lines, where we observed δv and δi fluctuations at the same frequency with consistent phase differences of roughly $\pi/2$, suggesting they are due to the same wave mode.

In order to better understand the complex variety of relationships between δv and δi fluctuations in the corona, future measurements should aim to overcome several limitations of the COMP data. First, the COMP data have rather coarse spatial resolution. A similar analysis should be performed on data with a spatial resolution sufficient to resolve the transverse structure of coronal loops. Second, the COMP data observe an Fe XIII line, which is formed at relatively high temperatures and is therefore insensitive to coronal holes. Hence we probably do not observe quiescent open magnetic field lines. Measurements of lines formed at cooler temperatures are needed to study open field structures. We next plan to




Figure 17. Plane-of-sky propagation speeds for $\delta\nu$ fluctuations (red symbols) and δi fluctuations (blue symbols). The speed is indicated by the radial distance from the solid circle, which corresponds to a speed of zero. Dotted circles are at intervals of 200 km s^{-1} .

conduct an analysis of DKIST data to address some of these limitations.

Acknowledgments

M.H. and D.W.S. acknowledge support from NSF AST grant No. AST-2005887 and NSF SHINE grant No. 2229100. X.F. acknowledges support from NASA under award No. 80NSSC23K0101, and NSF under award No. 2229101. We thank Mahboubeh Asgari-Targhi for interesting conversations about this work.

Facilities: HAO.

Software: IDL, solarsoft (S. L. Freeland & B. N. Handy 1998).

Appendix

Systematic Factors Influencing the Measurement of the Perpendicular Correlation Length

In Section 3.6 we described measurements of $L_{\perp}$ for $\delta\nu$ and δi fluctuations. Here we focus on the $\delta\nu$ measurements. To derive $L_{\perp}$, we computed the cross correlation between each

pixel and a region of neighboring pixels within a box of ± 12 pixels. We took the Gaussian width of these islands as a measure of the perpendicular correlation length. We obtained a sample of $L_{\perp}$ at $N = 45,215$ pixels with converged fits within the off-limb corona. A histogram of this set of $L_{\perp}$ measurements resembles a lognormal distribution characterized by a median value of $L_{\perp} = 7.0 \text{ Mm}$ and a most probable value of 5.0 Mm (see Figure 18). R. Sharma & R. J. Morton (2023) performed similar measurements of $L_{\perp}$ for $\delta\nu$ fluctuations, finding similar typical values for $L_{\perp}$ and remarking that their histograms also resembled a lognormal distribution.

Here our objective is to explore whether observational biases may prevent the measurement of small $L_{\perp}$ values. In particular, we identify systematic reasons that an underlying power-law distribution may be distorted by the measurements into a form that is roughly lognormal.

A.1. Analysis

Figure 18 shows a histogram of the $L_{\perp}$ measurements. In constructing the histogram, we omit values larger than 54 Mm ,




Figure 18. Histogram of the measured $L_{\perp}$ shown as a histogram in gray. The blue curve shows the best-fit lognormal distribution. The red curve illustrates the best-fit distribution obtained by convolving a power law with a Gaussian. The convolved distribution is a better match to the measurements.

because this $L_{\perp}$ corresponds to a $\sigma_{\perp}$ that exceeds the ± 12 pixel width of field of view used in the correlation analysis described in Section 3.6. We expect fits of such outlier $L_{\perp}$ values to represent poor fits or pixels where the high correlation island was not clear. We then normalize the measured $L_{\perp}$ histogram by dividing by the total number of samples ($N = 45, 215$) and the width of the binning used for the histogram, in this case $w = 0.45$ Mm. The normalization ensures that the integral of the histogram is equal to one, so that we can compare the distribution to the standard normalized form of probability distribution functions (pdf). We have verified that the bin size used in constructing the histogram does not qualitatively affect the results, and we also make use of metrics that are insensitive to the binning.

As a first step, we quantify the parameters of the lognormal distribution that best fits the observations. The pdf for the lognormal distribution is given by

$$p(x) = \frac{1}{\sigma x \sqrt{2\pi}} \exp \left[-\frac{(\ln(x) - \mu)^2}{2\sigma^2} \right]. \quad (\text{A1})$$

The parameters μ and σ represent the mean and standard deviation of $\ln(x)$, which follows a normal (Gaussian) distribution. The normal distribution occurs frequently in statistics, because it is a natural consequence of the central limit theorem, namely that the sum of independent random variables tends to follow the normal distribution (S. L. Meyer 1975). The lognormal distribution has a similar central limit theorem, except with multiplication rather than addition, because the logarithm of the product is the sum of the logarithms of each term. This situation occurs, for example, in fragmentation processes where the size of each fragment of an object is a fraction of its previous size (E. L. Crow & K. Shimizu 1988).

The blue curve in Figure 18 shows the lognormal distribution that best fits the measurements. The best-fit parameters were $\mu = 1.91 \pm 0.01$ and $\sigma = 0.439 \pm 0.008$. It would be useful to also describe the perpendicular wavenumber

$k_{\perp} \approx 2\pi/L_{\perp}$. The distribution of the inverse of a variable that is lognormally distributed is also lognormal, with the same σ , but the μ becomes negative. Accounting for the factor of 2π in the definition of $k_{\perp}$, the μ parameter for lognormal $k_{\perp}$ distribution would be $\mu = -0.072$.

A.1.1. Convolution Model

Our hypothesis is that there may be many values with small $L_{\perp}$ that we are not able to measure due to systematic factors. One distribution that is consistent with this hypothesis is a power law. The convolution of a Gaussian and a power-law distribution somewhat resembles a lognormal distribution and can arise if normally distributed random errors are added to power-law distributed data. So we next investigate whether such a convolution provides a better fit to the measured distribution. The normalized Gaussian distribution with mean zero is given by

$$G(x) = \frac{1}{\sqrt{2\pi}s} \exp \left[-\frac{x^2}{2s^2} \right], \quad (\text{A2})$$

where s represents the standard deviation of the distribution. A normalized power-law distribution can be defined described by

$$P_{\alpha}(x) = \frac{\alpha - 1}{x_0} \left(\frac{x}{x_0} \right)^{-\alpha}, \quad (\text{A3})$$

where $\alpha > 1$ is the power-law index and x_0 is a low-value cutoff, which ensures that the distribution can be normalized. The convolution of these distributions is the integral

$$P_{\alpha} \star G = \int_{x_0}^{\infty} P_{\alpha}(x') G(x - x') dx'. \quad (\text{A4})$$

The integration was performed numerically. We will refer to the resulting distribution as the Convolved distribution henceforth. We note that K. Parodi & T. Bortfeld (2006) have obtained closed form solutions in terms of special functions for a different formulation of the power-law distribution that is valid for power-law indices between 0 and 1.

We performed a least-squares fit to the histogram of the $L_{\perp}$ measurements to find the parameters of the best-fit convolved distribution parameters. The resulting fit is illustrated by the red curve in Figure 18. The best-fit parameters were $\alpha = 2.45 \pm 0.03$, $x_0 = 4.22 \pm 0.02$ Mm, and $s = 0.97 \pm 0.03$ Mm.

The convolved distribution is a better fit to the data than the lognormal distribution. This is clear from Figure 18, as well as more quantitative metrics. The least-squares fit minimizes the sum of the squared differences between the observed and model distribution. This value is smaller for the convolved distribution (1.3×10^{-3}) than the lognormal distribution (6.8×10^{-3}).

We have also used an alternative metric to compare the fits, the Akaike information criterion (AIC), which confirms that the convolved distribution is the better fit. The AIC is computed by first computing the likelihood function $\mathcal{L}$, given by

$$\mathcal{L} = \sum_i^N \ln f(x_i), \quad (\text{A5})$$

where f is the pdf representing a given model and x_i is the value of the i th measured data point, of which there are N in total. The AIC is defined by

$$\text{AIC} = 2k - 2 \ln \mathcal{L}, \quad (\text{A6})$$

where k is the number of parameters in the model: 2 (μ and σ) for the lognormal distribution and 3 (α , x_0 , and s) for the convolved distribution. A lower value of the AIC indicates a superior fit. In this case, the AIC values are comparable but favor the Convolved distribution with $\text{AIC}=2.5 \times 10^5$ for the Convolved distribution versus 2.6×10^5 for the lognormal distribution. The AIC has two benefits: First, it uses the data directly and is thereby independent of the binning used in constructing the histogram. Second, it accounts for the different number of parameters in the models. That is, one might expect that a three-parameter model will give a better fit than a two-parameter model, simply because there is more freedom to adjust the model. The AIC test accounts for this advantage.

As a third metric, we have also computed the Kolmogorov–Smirnov (K-S) test statistic for both models. The K-S test is designed to test whether two sets of data come from the same pdf. It uses the cumulative distribution function, $\text{cdf}(x)$, which is the integral of the pdf up to x . The K-S statistic is the absolute value of the maximum difference between the cdf of the two distributions being tested. Once again, the K-S test shows that the convolved distribution ($\text{KS} = 0.051$) is a better match to the data than the lognormal distribution ($\text{KS} = 0.097$).

A.1.2. A Simple Aggregation Model

An alternative model that provides a reasonable match to the observed distribution is that the actual distribution of $L_{\perp}$ follows a power-law distribution, but the measured distribution is distorted by sampling and averaging over that distribution. We will refer to this as the aggregation model. For this model, suppose that the data come from a power-law distribution described by Equation (A3). The measured value, though, represents an average of n samples drawn from the power-law distribution with n assumed to follow a Poisson distribution. Physically this model could represent the averaging of multiple structures along the line of sight in the measurements of the optically thin corona.

In order to determine whether this model is consistent with the data and the best-fit parameters, we simulated the process described previously. That is, we selected random numbers from a power-law distribution with parameters α and x_0 and averaged over n of those numbers, where n was a random number drawn from a Poisson distribution with a mean value of n_0 . This process was repeated 1000 times and then the distribution of the resulting averages was examined. In order to identify the best-fit parameters (α , x_0 , and n_0), we used a differential evolution algorithm to minimize the KS statistic comparing the model distribution to the measurements of $L_{\perp}$. Minimizing the KS metric has several advantages over performing a least-squares fit. One is that the KS metric is independent of the binning for the histogram. More importantly, the KS metric is robust to statistical variations in the model distribution. A least-squares fit requires many more samples to converge, because the samples are generated stochastically and so χ^2 varies with each iteration, even if the fit parameters remain the same. Because of these random variations, the fit parameters vary slightly with each run, so we estimate the uncertainties on the fit parameters by taking the average and standard deviation of the best-fit values over 20 runs.

Figure 19 compares the aggregation model to the measurements for the averaged best-fit parameters. These parameters were $\alpha = 2.02 \pm 0.09$, $x_0 = 1.2 \pm 0.4$ Mm, and $n = 1.0 \pm 0.3 \times 10^2$. From the figure it is clear that the aggregate model provides a good fit for the data, and in fact the average aggregate model $\text{KS} = 0.046$ was comparable to that of the convolved distribution. It is not possible to compare the other metrics discussed previously, because we have not derived an explicit expression for the pdf corresponding to the aggregation model.

A.2. Discussion

The provided models demonstrate that even if the actual distribution of $L_{\perp}$ were statistically distributed as a power law, we could obtain a distribution that approximates a lognormal distribution due to systematic biases in the measurement. One way this can happen is if the measurement involves the addition of normally distributed random noise, which will result in a convolution of the power law with a Gaussian. A second way is if the measurement involves an implicit aggregation or averaging of the power-law distributed values.

Our measurement process does involve steps that are consistent with these factors. Considering the convolution model, the analysis of the correlation lengths involves computing a cross correlation of time-series data, which has photon counting noise as well as noise from the detector. There are also uncertainties involved in the step where the cross correlation is fit with a Gaussian to infer $\sigma_{\perp}$. From the perspective of the aggregation model, the observed Doppler velocity fluctuations are measured using optically thin emission lines, and those Doppler velocities are implicitly averaged along the line of sight. This averaging is known to have a significant effect on the amplitudes of the Doppler fluctuations (S. W. McIntosh & B. De Pontieu 2012). Other observational factors may also contribute to the observed distribution. For example, the instrument resolution may limit the ability to measure below a certain length scale. It is suggestive that the convolved fit identifies a cutoff value of $x_0 = 4.2$ Mm, which is very close to the effective pixel scale 4.5 Mm. There may also




Figure 19. Histogram of the measured $L_{\perp}$ shown as a histogram in gray. The black curve shows the best-fit aggregation model, which is obtained by averaging data values drawn from a power-law distribution.

be physical limits to the distribution, for example if Alfvénic waves with small $L_{\perp}$ are preferentially dissipated.

The properties of the distribution of $L_{\perp}$ have physical implications. A power-law distribution of $L_{\perp}$ suggests that most of the Alfvénic wave energy is concentrated at small $L_{\perp}$ and that large $L_{\perp}$ values occur sporadically. In contrast, a lognormal distribution suggests that there is a preferred scale for the $L_{\perp}$. For example, there is a peak in the solar p-mode spectrum at spherical harmonics with angular wavenumbers of $l \sim 300$, which corresponds to lengths of ~ 10 Mm (e.g., I. G. Hernández et al. 1998). This would be consistent with the median measured values of $L_{\perp}$ found here.

In solar physics, the question of distinguishing a power law from a lognormal distribution also occurs in the study of the distribution of flare energies. Many studies of flare energies, or proxies such as the X-ray flux, have inferred that these energies follow a power-law distribution. If $\alpha > 2$ for these energies, then there are sufficiently numerous small events to heat the corona, which is the basis for the nanoflare theory of coronal heating. This power-law form of the nanoflare distribution has been invoked in support of models of nanoflares based on self-organized criticality (P. Charbonneau et al. 2001) or turbulence (e.g., G. Boffetta et al. 1999). However, it has also been argued that the flare distribution is not a power law. C. Verbeecq et al. (2019) showed that some proxies, such as solar flare EUV intensities, are better fit by a lognormal distribution, although other proxies, such as the X-ray flux, do seem consistent with a power law.

Some of the models that we have discussed in terms of the $L_{\perp}$ measurements may have relevance to the problem of understanding the distribution of flares. E. D’Huys et al. (2016) argued that some of the power-law fits obtained were not performed in a rigorous way. One factor is that many of these fits are performed only for data that exceed a certain threshold. E. Buchlin et al. (2005) supported imposing such a threshold in the analysis definition as being consistent with the power-law distribution, because line-of-sight averaging in observational studies causes small events to be not observed. This is

essentially the same issue that we have addressed in the aggregation model without imposing an ad hoc threshold.

Another example in solar physics of the connection between lognormal and power-law distributions concerns the intensity distribution of the quiet Sun and its relationship to nanoflare coronal heating. A. Pauluhn et al. (2000) studied the intensity distribution of the quiet Sun chromosphere, transition region, and corona. They found that the distribution of these intensities was well-described by a lognormal distribution. However, nanoflare heating models that assume a power-law energy distribution are able to reproduce these lognormal intensity distributions (A. Pauluhn & S. K. Solanki 2007; M. Hahn et al. 2022b). The processes that allow the nanoflare power-law distribution to generate lognormally distributed intensities involve steps that have a similar effect as our convolution and aggregation models.

A.3. Summary

The statistical distribution of measurements of the perpendicular correlation length of Doppler velocity fluctuations in the corona resembles a lognormal distribution. However, the observations can be better fit by models that consider the underlying distribution to be a power law. These models can represent the addition of noise and the implicit averaging in the observation, along with other systematic factors, such as the influence of spatial resolution.

These interpretations are tentative. There is insufficient information to determine objectively the true distribution of the data. These $L_{\perp}$ measurements come from different locations throughout the corona. If the characteristics of Alfvénic wave excitation are not uniform throughout the corona, structural differences in the corona may cause real variations $L_{\perp}$, and possibly this accounts for the observed distribution. Alternatively, the peak in the $L_{\perp}$ distribution may represent a preferred length scale for injecting wave energy into the corona, such as a peak in the p-mode spectrum. Further insight into the nature of the $L_{\perp}$ measurements will be gained by making measurements with different instruments having different limitations.

Our conclusion is thus limited to demonstrating the possibility that existing measurements have not resolved the smallest $L_{\perp}$ values. This possibility has important implications, because it implies that there may be a large population of Alfvénic waves in the corona with small perpendicular wavelengths (high $k_{\perp}$). Such waves could arise from the turbulent excitation of the Alfvénic waves or Alfvénic turbulence in situ driving waves to high $k_{\perp}$. Alternatively, the Alfvénic waves might be excited by nanoflares occurring on small, unresolved spatial scales. However, our inferred power-law parameters should not be directly interpreted to represent the power spectrum of turbulence as a function of $k_{\perp}$; the assumed $L_{\perp}$ power law slopes in the opposite direction of that predicted for turbulence, which predicts greater power at larger values of $L_{\perp}$. Accurate measurements of the perpendicular correlation lengths or wavenumbers in the corona will require higher-resolution measurements and a more-comprehensive modeling of systematic factors.

ORCID iDs

Michael Hahn  <https://orcid.org/0000-0001-7748-4179>

Stefan J. Hofmeister  <https://orcid.org/0000-0001-7662-1960>

Yifan Huang  <https://orcid.org/0000-0002-0422-2925>

Alexandros Koukras  <https://orcid.org/0000-0002-6232-5527>

Daniel Wolf Savin  <https://orcid.org/0000-0002-1111-6610>

References

- Aschwanden, M. J., & Peter, H. 2017, *ApJ*, **840**, 4
- Asgari-Targhi, M., Asgari-Targhi, A., Hahn, M., & Savin, D. W. 2021, *ApJ*, **911**, 63
- Banerjee, D., Krishna Prasad, S., Pant, V., et al. 2021, *SSRv*, **217**, 76
- Boffetta, G., Carbone, V., Giuliani, P., Veltri, P., & Vulpiani, A. 1999, *PhRvL*, **83**, 4662
- Bowen, T. E., Badman, S., Hellinger, P., & Bale, S. D. 2018, *ApJL*, **854**, L33
- Buchlin, E., Galtier, S., & Velli, M. 2005, *A&A*, **436**, 355
- Charbonneau, P., McIntosh, S. W., Liu, H.-L., & Bogdan, T. J. 2001, *SoPh*, **203**, 321
- Crow, E. L., & Shimizu, K. 1988, *Lognormal Distributions: Theory and Applications* (New York: Marcel Dekker)
- De Moortel, I., McIntosh, S. W., Threlfall, J., Bethge, C., & Liu, J. 2014, *ApJL*, **782**, 34
- D’Huys, E., Berghmans, D., Seaton, D. B., & Poedts, S. 2016, *SoPh*, **291**, 1561
- Freeland, S. L., & Handy, B. N. 1998, *SoPh*, **182**, 497
- Gruszecki, M., Nakariakov, V. M., & Van Doorselaere, T. 2012, *A&A*, **543**, 12
- Guo, M., Van Doorselaere, T., Karamelas, K., & Li, B. 2019, *ApJ*, **883**, 20
- Hahn, M., D’Huys, E., & Savin, D. W. 2018, *ApJ*, **860**, 34
- Hahn, M., Fu, X., & Savin, D. W. 2022a, *ApJ*, **933**, 52
- Hahn, M., Ho, B., & Savin, D. W. 2022b, *ApJ*, **936**, 113
- Heinemann, M., & Olbert, S. 1980, *JGR*, **85**, 1311
- Hernández, I. G., Patrón, J., Chou, D.-Y., & TON Team. 1998, *ApJ*, **501**, 408
- Hollweg, J. V. 1971a, *PhRvL*, **27**, 1349
- Hollweg, J. V. 1971b, *JGR*, **76**, 5155
- Howes, G. G., & Nielson, K. D. 2013, *PhPl*, **20**, 072302
- Jenkins, G. M., & Watts, D. G. 1968, *Spectral Analysis and its Applications* (San Francisco, CA: Holden-Day)
- Kasper, J. C., Klein, K. G., Lichko, E., et al. 2021, *PhRvL*, **127**, 255101
- Malanushenko, A., Cheung, M. C. M., DeForest, C. E., Klimchuk, J. A., & Rempel, M. 2022, *ApJ*, **927**, 1
- Mardia, K. V., & Jupp, P. E. 2000, *Directional Statistics* (New York: Wiley)
- McIntosh, S. W., & De Pontieu, B. 2012, *ApJ*, **761**, 138
- McIntosh, S. W., De Pontieu, B., & Tomczyk, S. 2008, *SoPh*, **252**, 321
- Meyer, S. L. 1975, *Data Analysis for Scientists and Engineers* (New York: Wiley)
- Morton, R. J., Tomczyk, S., & Pinto, R. F. 2016, *ApJ*, **828**, 89
- Musielak, Z. E., Fontenla, J. M., & Moore, R. L. 1992, *PhFIB*, **4**, 13
- Nakariakov, V. M., Roberts, B., & Murawski, K. 1997, *SoPh*, **175**, 93
- Ofman, L., Nakariakov, V. M., & Sehgal, N. 2000, *ApJ*, **533**, 1071
- Parodi, K., & Bortfeld, T. 2006, *PMB*, **51**, 1991
- Pauluhn, A., & Solanki, S. K. 2007, *A&A*, **462**, 311
- Pauluhn, A., Solanki, S. K., Rüedi, I., Landi, E., & Schühle, U. 2000, *A&A*, **362**, 737
- Peter, H., Bingert, S., Klimchuk, J. A., et al. 2013, *A&A*, **556**, 104
- Phillips, K. J. H., Feldman, U., & Landi, E. 2008, *Ultraviolet and X-Ray Spectroscopy of the Solar Atmosphere* (Cambridge: Cambridge Univ. Press)
- Rast, M. P., Bello González, N., Bellot Rubio, L., et al. 2021, *SoPh*, **296**, 70
- Réville, V., Tenerani, A., & Velli, M. 2018, *ApJ*, **866**, 38
- Sharma, R., & Morton, R. J. 2023, *NatAs*, **7**, 1301
- Shoda, M., Suzuki, T. K., Asgari-Targhi, M., & Yokoyama, T. 2019, *ApJ*, **880**, L2
- Shoda, M., & Yokoyama, T. 2016, *ApJ*, **820**, 123
- Shoda, M., Yokoyama, T., & Suzuki, T. K. 2018, *ApJ*, **860**, 17
- Tomczyk, S., & McIntosh, S. W. 2009, *ApJ*, **697**, 1384
- Tomczyk, S., McIntosh, S. W., Keil, S. L., et al. 2007, *Sci*, **317**, 1192
- Tomczyk, S., Card, G. L., Darnell, T., et al. 2008, *SoPh*, **247**, 411
- Vaughan, S. 2005, *A&A*, **431**, 391
- Velli, M. 1993, *A&A*, **270**, 304
- Verbeeck, C., Kraaikamp, E., Ryan, D. F., & Podladchikova, O. 2019, *ApJ*, **884**, 50
- White, O. R., & Athay, R. G. 1979, *ApJS*, **39**, 317
- Zank, G. P., Zhao, L.-L., Adhikari, L., et al. 2022, *ApJ*, **926**, L16
- Zhao, L.-L., Zank, G. P., Telloni, D., et al. 2022, *ApJ*, **928**, L15