Precision and systematic errors in global helioseismology mode fitting and inversions: Leveraging some 25 years of nearly uninterrupted observations

Sylvain G. Korzennik*

Solar Stellar and Planetary Science Division, Center for Astrophysics, Harvard & Smithsonian, Cambridge, MA, United States

We have on hand some 25 years of nearly uninterrupted high-quality and high-cadence global helioseismic data. The Global Oscillations Network Group (GONG) project has been producing science quality data since 1995, the Michelson Doppler Imager (MDI) started in 1996, and the Helioseismic and Magnetic Imager (HMI) took over in 2010. Fundamental new constraints have been imposed by helioseismic inferences, yet global helioseismology data processing seems somewhat frozen in time for some of its methodologies. I review and discuss some specific aspects of global helioseismology data analysis, with an emphasis on the issues and challenges presented by mode fitting and inversion techniques. I compare and contrast results derived by different fitting methods, whether using different techniques, different lengths of time series, or different fitting parameters, like leakage matrices or the inclusion or omission of the mode profile asymmetry, leading to our current best handle on the residual systematic errors.

KEYWORDS

Sun, Sun—interior, Sun—helioseismology, mode fitting, inverse modeling

1 Introduction

In 1995, the Global Oscillations Network Group (GONG) project started acquiring nearly continuous helioseismic observations using a network of six nearly identical instruments deployed around the globe in both the northern and southern hemispheres (Harvey et al., 1996). The GONG instrument design is based on a Michelson interferometer called a Fourier tachometer (Beckers and Brown, 1978); while it was initially outfitted with 256-by-242-pixel detectors, the cameras were a little later upgraded to 1,024-by-1,024-pixel detectors (in the 2001–2002 time frame).

Almost a year later, the Michelson Doppler Imager (MDI), flying on board the Solar and Heliospheric Observatory (SOHO), started its nearly continuous helioseismic observations. The MDI instrument used two tunable Michelson interferometers centered around the Ni I photospheric absorption line at 6767.8 Å (Title and Ramsey,

1980). Since SOHO is in a halo orbit around the Sun–Earth L1 Lagrangian point, the available telemetry was limited; therefore, the 1,204-by-1,024-pixel full-disk images had to be binned down on board the spacecraft to 256-by-256-pixel images to fit within the telemetry rate and still provide continuous helioseismic observations, known as the *Structure* program (Scherrer et al., 1995). Contact with the SOHO spacecraft was at some point lost and later recovered; hence, the MDI observations have a high duty cycle along with two long gaps in late 1998 and early 1999.

In 2010, the Helioseismic and Magnetic Imager (HMI), flying on board the Solar Dynamics Observatory (SDO), started its nearly continuous helioseismic observations. The HMI instrument is conceptually very similar to MDI but is centered around the Fe I absorption line at 6173 Å and uses 4,096-by-4,096-pixel detectors and a 45 s observing cadence (Scherrer et al., 2012), while GONG and MDI use a 60 s observing cadence. Since SDO is in an inclined geosynchronous orbit, it allows nearly continuous, high-data-rate telemetry, thus leading to the acquisition of nearly continuous full-disk images. After nearly a year of coeval observations with HMI, the MDI observation program was terminated in 2011.

For historical reasons, the GONG project reduces its observations in 36-day-long chunks and performs mode fitting, based on Anderson et al. (1990), using 108-day-long time series starting every 36 days. The MDI and the HMI observations are reduced in 72-day-long chunks and the mode fitting, based on a very different approach (Schou, 1992; Larson and Schou, 2015), uses 72-day and 360-day-long abutting time series. Hence, direct comparison of mode characteristics between GONG and MDI or HMI, computed by the respective projects, is for the least tricky and neither project has attempted to implement either a more modern and sophisticated fitting methodology or a more systematic approach to fitting longer time series.

The GONG fitting methodology (Anderson et al., 1990) is known to have its pitfalls since it fits what it considers individual modes with not built-in knowledge of the leakage matrix, has no gap filling, and uses a symmetric mode profile. Some of these individual modes are known to be blended modes, and thus, their measurements have known systematic errors, while the mode profile is known to be asymmetric (Duvall et al., 1993; Hill et al., 1998; Howe and Hill, 1998; Howe and Thompson, 1998; Komm et al., 1998). Similarly, the methodology used to fit the MDI and HMI observations does not fit individual modes but uses an expansion in m to characterize the frequency splittings. It does include a model of the leakage matrix, and it includes, or not, an asymmetric mode profile; yet despite the extensive work presented in Larson and Schou (2015), residual systematic errors are still present.

A drastically different approach has been developed in Korzennik (2005) and later improved in Korzennik (2008a)

and Korzennik (2008b), using long time series at first, but also shorter ones, and a logical progression and regular spacing.[1] Expanding on this approach, I present results from fitting the longest to date available time series, namely, a 9,216-day-long time series of GONG observations and results from fitting 4,608-day-long and 2,304-day-long time series of MDI and HMI observations, respectively.

Leveraging the high SNR and low background noise of such long time series, I explore the precision and systematic errors primarily associated with the specifics of the leakage matrix used to carry out the mode fitting, but also the effect of averaging two shorter time series versus fitting a longer time series, or fitting coeval observations using different instruments.

2 Data sets

The data sets used for this work are time series of spherical harmonic coefficients computed by the GONG, MDI, and HMI projects using the respective projects' pipeline processing.

For GONG, these correspond to a weighted mean of spherical harmonic coefficients computed for each coeval dopplergram acquired by the GONG network at any given time. The procedures used by the GONG project to calibrate the raw images and to compute these spherical harmonic coefficients are described in Harvey et al. (1996). To date, these are available on a 1-min cadence for 9,720 days (or 26.6 years), starting on 1995.05.07 and ending on 2022.01.20. While the GONG project does not perform any gap filling, I have adapted the maximum entropy gap filling methodology as coded by Rasmus Munk Larsen to fill these time series. This methodology is used for gap filling in the MDI and HMI time series and is based on Fahlman and Ulrych (1982). This increased the overall GONG fill factor from 85.4% to 95.0%.

For MDI, these correspond to the spatial decomposition of the on-board Gaussian-smoothed 1,024-by-1,024-pixel full-disk images of the *Structure* program. This program took continuous observations at a 1-min cadence, reducing them, on board the spacecraft, to 256- by-256-pixel images to fit within the continuously available telemetry. The procedures used by the MDI project to calibrate the raw images and to compute these spherical harmonic coefficients are described in Scherrer et al. (1995). They are available for 5,472 days, starting on 1996.05.01 and ending on 2011.04.24, with an overall fill factor of 93.6%, after gap filling, with two long gaps when contact with the SOHO spacecraft was lost; the average fill factor, when excluding these gaps, is 96.7%, after gap filling.

---

1   I have fitted time series as short as 36 days and used a factor two progression, fitting time series that are 36, 72, 144, 288 days etc and offsetting them by half their length for those longer than 72 days.

For HMI, these correspond to the spatial decomposition of 4,096-by-4,096-pixel full-disk images acquired nearly continuously at a 45 s cadence. The procedures used by the HMI project to calibrate the raw images and to compute these spherical harmonic coefficients are described in Scherrer et al. (2012). To date, they are available for 4,464 days, starting on 2010.04.30 and ending on 2022.07.20, with a fill factor as high as 99.97%, after gap filling.

As the GONG project reached the milestone of producing its 273rd 36-day-long time series of spherical harmonic coefficients, I was able to fit the longest time series to date of helioseismic observations, namely, a 9,216-day-long time series of 13,271,040 data points spanning some 25.23 years or approximately a full solar cycle.[2]

Such a time span is twice as long as the longest time series I had analyzed before, using either MDI or GONG observations. By contrast, the HMI observations have yet to reach the 4,608-day-long mark (i.e., 64 of the 72-day-long segments), so the longest time series of HMI data I have fitted to date remains a 2,304-day-long one (i.e., 32 consecutive 72-day-long segments).

3 Fitting and inversion methodologies

The analysis of these time series consists in estimating the limit spectrum for each spherical harmonic (ℓ, m), where ℓ is the degree and m is the azimuthal order, with $|m| \leq \ell$, then fitting the individual modes present in these spectra for each singlet (n, ℓ, m), where n is the radial order of the mode.

Using long time series produces estimates of limit spectra with a high SNR and a very low background noise level, but it also averages out any changes of the underlying mode properties. In practice, since the mode frequencies are known to change with time, as they are affected by the solar activity, the observed mode widths will not simply be the convolution of the intrinsic mode width by the window function but will also be widened by the mode frequency changes.

However, fitting shorter time series leads to lower precision on the mode characteristics since the fitted spectra have a lower SNR following the $\sqrt{T}$ law, where T is the length of the time series, but the fitting is also affected by the effective mode amplitude over the given time series, due to stochastic excitation.

This leads, in practice, to mode attrition: the exact same set of modes cannot be successfully fitted each time; hence, computing the average of results from fitting shorter time series—like using the nearly equivalent 85 108-day-long time series fitted by

GONG, the equivalent 128 72-day-long time series fitted by MDI or HMI, or the not quite equivalent 25 360-day-long time series, also fitted by MDI or HMI—can only produce a smaller set of always present modes.

This has been emphasized in Korzennik (2013), while such comparisons are further complicated by the facts that 1) the GONG fitting procedure does not include any information on the leakage matrix and is fitting symmetric mode profiles—although the mode profile asymmetry is now well established (Duvall et al., 1993)—all the while, 2) the MDI and HMI fitting methodologies use an expansion in m to parameterize the mode frequency splitting rather than fitting individual singlets, and 3) GONG has been fitting 108-day-long time series, offset by 36 days, while MDI and HMI have been fitting 72- and 360-day-long time series, offset by 72 or 360 days, respectively (see discussion in Korzennik, 2013).

There is a clear advantage in fitting the longest time series available since 1) the resulting spectra have high SNR and low background noise, and 2) most of the individual modes will have amplitude well above the noise level. The drawback, of course, is that one ends up measuring an average of the mode characteristics over the time span by the time series.

While I present results from fitting the same very long GONG time series with different leakage matrices, I also present results from fitting coeval GONG and MDI or GONG and HMI somewhat shorter time series, using the same methodology and a leakage matrix computed for each instrument using the same formalism.[3] To further assess the effect of any residual systematic errors, I also invert the corresponding rotational splittings to infer an estimate of the mean solar internal rotation rate and contrast such inferences when using rotation splittings derived using, for instance, the same time series and fitting methodology but different leakage matrices.

3.1 Mode fitting

The mode fitting methodology I used has been described in detail in Korzennik (2005), Korzennik (2008a), and Korzennik (2008b). Let me outline here some of its distinctive features. This method uses least squares minimization of an optimal estimate of the limit spectrum at each m, fitting simultaneously a portion of the $2\ell + 1$ multi-tapered spectra, centered around a previous estimate of the mode frequency at each m, for a target mode of a given n and ℓ.

---

2   While the GONG network started observing some 360 days before MDI, I have aligned all the segments I have fitted with the MDI start date (1996.05.01) to allow meaningful comparisons of coeval time series. I have fitted shorter time series of GONG observations starting as early as 1995.05.07, but I choose to fit the longer time series starting on 1996.05.01.

---

3   One cannot use the same leakage matrices from the three instruments since i) the image resolution and the spatial decomposition are different, ii) the instruments' point spread functions are different, and iii) MDI observations are binned down using a Gaussian smoothing.

The mode profile model is a generalized Lorentzian, that is, modified to include asymmetry, and is parameterized by a power amplitude, $A_{n,\ell,m}$; a frequency, $\nu_{n,\ell,m}$; a FWHM, $\Gamma_{n,\ell}$; and an asymmetry, $\alpha_{n,\ell}$—with the FWHM and the asymmetry set to be independent of m. A background term, $B_{n,\ell,m}$, is also added to model the fraction of the fitted spectra.

The limit spectra are estimated using sine multi-tapers to smooth out the realization noise resulting from the stochastic nature of the mode excitation. The number of tapers used when fitting a given multiplet, $\nu_{n,\ell}$, is selected from a preset list to be the largest one that corresponds to a spectral resolution that is at least half the estimated effective mode width, that is, the convolution of the intrinsic mode width by the spectral resolution. For the 9,216-day-long time series, this list is given by $1 + 2^k$ for $k = 1, 8$ (i.e., a 3, 5, 9... 257 progression).

Obviously, the spatial decomposition of the observed dopplergrams can only be carried out using the visible solar disk. It also must be performed with some apodization at the limb to avoid noise magnification and uses the line-of-sight projection of the velocity. The resulting projection on any target spherical harmonic, ℓ, m, is not orthogonal. Consequently, non-target modes leak into the spherical harmonic coefficient estimates; namely, the computed coefficient, $\tilde{C}_{\ell,m}$, is a linear combination of $C_{\ell',m'}$, the true spherical harmonic coefficients, weighted by leakage factors, also known as the elements of the leakage matrix.

The observed images are at first approximation limited to half a sphere; hence, the dominant leaks obey the parity rules of spherical harmonics, and these will correspond to values of ℓ', m' when $\delta\ell + \delta m = \pm 2k$, where $\delta\ell = \ell' - \ell$, $\delta m = m' - m$, and k is a positive integer. The amplitude of the leaks diminishes with $|k|$ but remains significant for $|k| \leq 6$, and the fitting is performed on a section of the spectrum limited to cover at least eight times the FWHM of the target mode in most cases.[4]

The leakage coefficients are further combined to account for the distortion of the eigenvalues by differential rotation, following the formalism in Woodard (1989). The differential rotation is characterized by two coefficients, $b_2 = -0.07712$ and $b_4 = -0.04396$ nHz, while the horizontal to vertical ratio, $\beta_{\ell,n}$, is computed using an adiabatic oscillation code and a standard solar model.

The fitted model for each m includes the (n, ℓ, m) target mode and the (n, ℓ', m') modes that leak into the estimate of the spherical harmonic coefficient $\tilde{C}_{\ell,m}$, with an amplitude attenuated by the leakage matrix. Hence, $2\ell + 1$ individual modes (also known as singlets) are fitted simultaneously for a given (n, ℓ), with as many mode amplitudes, frequencies, and background levels but only one FWHM and one asymmetry parameter. A "sanity check" is carried out throughout the fitting

process, and any singlet whose amplitude is below a set threshold compared to the spectrum's SNR is not fitted.[5]

The procedure also accounts for the cases where modes for (n', ℓ', m'), where $n' \neq n$, leak into the fitting window of a (n, ℓ, m) mode. To keep the fitting procedure tractable, this is included in an iterative fashion, using the estimate for the n' modes from the previous iteration.

The whole process starts with a set of initial guesses, and without this n' "contamination," it is then included, and these two first least squares minimizations are configured to allow the fitted parameters to vary significantly from the initial guesses. Thereafter, the fitting is repeated 20 times to properly account for the mode contamination[6] but with the fitting somewhat restricted so as to "explore" a smaller parameter space for efficiency. For the GONG observations, spherical harmonic coefficients are computed and fitted up to $\ell = 200$, while for MDI and HMI observations, these are available and, thus, fitted up to $\ell = 300$.

Uncertainties on all the fitted parameters were estimated from the covariance matrix of the problem. This covariance matrix was estimated from the Hessian matrix (Press et al., 1986) using numerical estimates of the second derivative of the merit function. The increases appropriate for the estimate of these derivatives were based on the size of the shrunken simplex resulting from the fitting.

It should be noted that the leaks closest in frequency to the target modes are in most cases the ones corresponding to $\delta\ell = 0$ and $\delta m = \pm 2$. These leaks are about 840 nHz away from the target mode since the frequency spacing with respect to m is due to the frequency splitting caused by the solar rotation (or 420 nHz). Therefore, when the mode FWHM is comparable to or larger than frequency separation, the mode and its closest leaks are not resolved. If the leakage matrix is ignored or erroneous, these modes will be the ones most affected by systematic errors. To characterize such systematic errors, I have fitted the same time series using different leakage matrices.

3.2 Rotation inversion techniques

To further characterize systematic errors, I also computed and contrasted inferences of the mean rotation rate using different estimates of the individual frequency splittings. To that effect, I used two inversion techniques that were described in Eff-Darwich and Perez Hernandez (1997), Eff-Darwich et al. (2010), Gregor and Fessler (2015), and Korzennik & Eff-Darwich (in preparation).

---

4   This range is set to be at least some empirical minimum value and not to exceed the mode spacing in n around the target mode.

5   Only modes with a power amplitude three times greater than the RMS of the residuals to the fit are fitted.

6   That number of iterations was deemed adequate based on the changes in the parameters at the last iteration.

The first one is a piece-wise constant regularized least squares method, with the regularization based on smoothing the solution's second derivative. This smoothing is parameterized by a single Lagrange multiplier that acts as a trade-off between spatial resolution and error magnification, while the smoothing is spatially weighted by the relative precision of the inferred profile at each grid point. The second one uses an iterative algorithm[7] with either a global second-derivative smoothing factor or a local first-derivative smoothing factor, set to span 3% of the radius and co-latitude range, and a single control parameter that again acts as a trade-off between spatial resolution and error magnification. The solution grid, that is, the grid used to derive the piece-wise constant rotation profile, was spaced out according to the solution's spatial resolution.

For both techniques, the corresponding averaging kernels are computed at each grid point as well as some of their properties, namely, the center of gravity of the main peak, a proxy for the effective location of the inferred solution at that grid point, as well as the width of that main peak, an estimate of the solution's spatial resolution at that location.

4 Results

The longest GONG time series, that is, the 9,216-day-long one, has been fitted with the exact same procedure but using five different estimates of the leakage matrix. One is the estimate computed by Schou (private communication) using the MDI/HMI formalism but using the specific set of parameters used to spatially decompose the GONG images, which are different from MDI or HMI spatial decomposition.

The other four leakage matrices were computed using my own formalism, where I compute the line of sight of a velocity signal for a single spherical harmonic on a grid that matches the observed images and then spatially decompose these images using the same parameters used to decompose the GONG images. One leakage matrix was computed using a value of $B_o = 0$, where B_o is the heliographic latitude at disk center, which varies as the Earth rotates around the Sun. Hence, for a time series much longer than a year, it averages out to zero. Vorontsov (private communication) has argued that the mean leakage matrix should instead be computed using the quadratic mean of the variation of B_o. Thus, I also computed a leakage matrix for $B_o = 7.155\sqrt{0.5} \equiv a_2$.

For these two values of B_o, I computed two additional leakage matrices where I also included a model of the mean point spread function (PSF) of the GONG instrument (Williams et al., 1995), a model based on the observed sharpness of the limb profile.

---

7 Also known in tomography as the simultaneous iterative reconstruction technique of regularized weighted least squares problems.

For additional comparisons, I also present results from fitting the 9,216-day-long GONG time series using a symmetric mode profile by "simply" modifying my fitting procedure to set $\alpha_{n,\ell}$ to zero as an initial guess and removing it from the list of adjusted parameters. I also present results from fitting two abutting 4,608-day-long time series of GONG observations and averaging them, fitting the 4,608-day-long MDI time series coeval with the GONG time series, and comparing these results to fitting the 4,608-day-long time GONG series. Finally, since the longest time series of HMI observations I have fitted to date is still "only" 2,304 days long, I also present comparisons of fitting coeval GONG and HMI 2,304-day-long time series. For these comparisons, I used the leakage matrix computed using the same formalism but with parameters adjusted to match the way each instrument's images were processed.
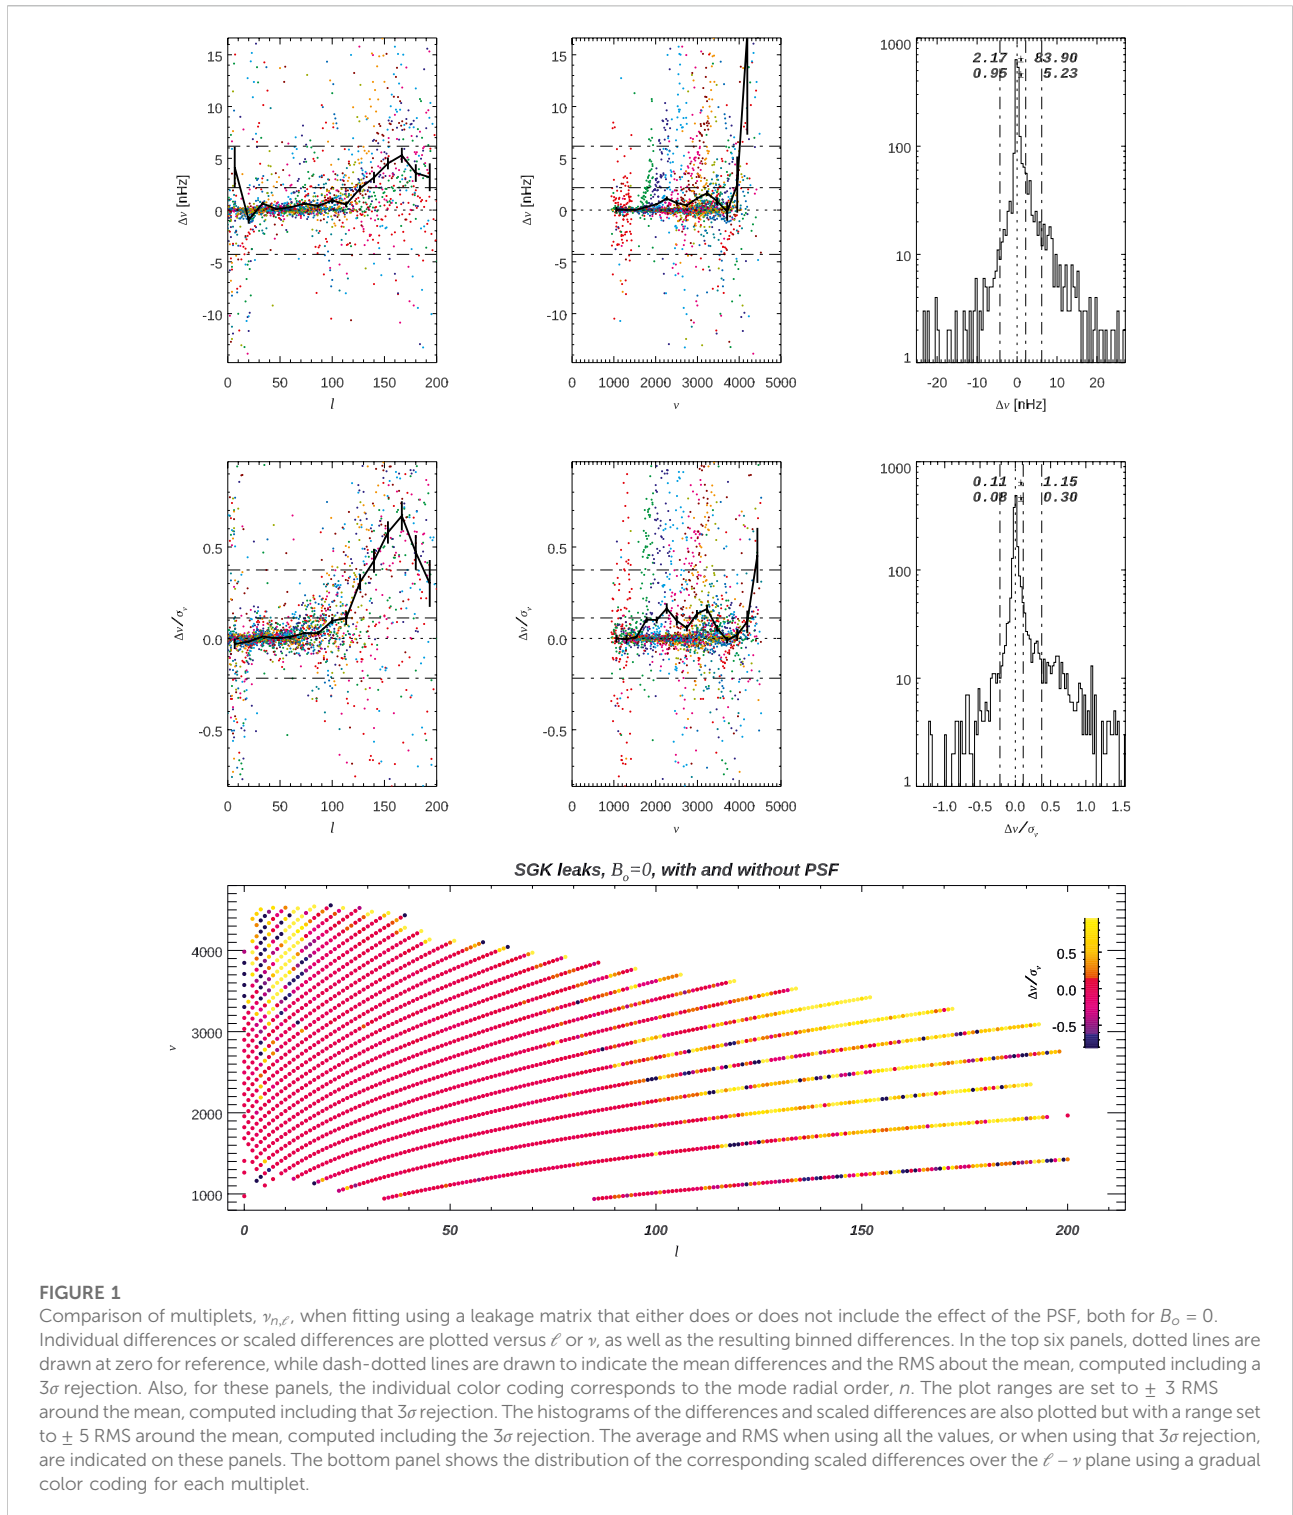
4.1 Mode parameter comparisons

Figure 1 compares mode frequencies with or without a model of the GONG mean PSF in the computation of the leakage matrix, both using the $B_o = 0$ value. The plots in that figure compare the mode multiplets, namely, $\nu_{n,\ell}$ estimated from the singlets and $\nu_{n,\ell,m}$ after fitting a ninth-order Clebsch–Gordan expansion in m/ℓ to them.

The raw differences and the scaled differences, that is, the differences divided by their respective uncertainties, show systematic differences at higher degrees, namely, for $\ell \geq 100$, but also some increased scatter at very low degrees for the high frequencies. For those low-ℓ high-ν modes, the mode width is large, and thus, the modes are not well resolved; hence, the sensitivity of the fitting to the leakage matrix is the greatest. For the high-ℓ, the effect of the PSF is to attenuate the mode visibility; hence, it skews the leaks with respect to $\delta\ell$. On average, the difference is 2.17 ± 83.90 nHz (or 2.61 ± 51.5 nHz when comparing singlets), while the binned difference, when binning in ℓ, increases to 5 nHz and spikes to 16 nHz at high frequency, when binned in ν. The mean and RMS of the scaled differences are 0.11 ± 1.15 (or 0.018 ± 0.15 when comparing singlets), but when binned in ℓ, it increases up to 0.7σ and to 0.45σ when binned in ν, with quite a large scatter at the low and high degrees. Some 8.7%, 1.9%, and 1.1% of the multiplets show differences that exceed the 1, 3, and 5σ thresholds, respectively. This is shown in the histogram tails, while the representation in a $\ell - \nu$ diagram clearly shows specific locations for the larger differences and that they are not distributed randomly.

Figure 2 compares other mode characteristics when including or not including the PSF, both also using the $B_o = 0$ value: estimates of the frequency uncertainty, the FWHM, the asymmetry, and the mean power amplitude and their ratios or differences. The σ_ν and Γ ratios show departure from 1.0 at the 4% and 0.5% levels, respectively, while the binned values, with respect to ν or ℓ, suggest systematic patterns. The comparison of

FIGURE 1
Comparison of multiplets, $\nu_{n,\ell}$, when fitting using a leakage matrix that either does or does not include the effect of the PSF, both for $B_o = 0$. Individual differences or scaled differences are plotted versus ℓ or ν, as well as the resulting binned differences. In the top six panels, dotted lines are drawn at zero for reference, while dash-dotted lines are drawn to indicate the mean differences and the RMS about the mean, computed including a 3σ rejection. Also, for these panels, the individual color coding corresponds to the mode radial order, n. The plot ranges are set to $\pm$ 3 RMS around the mean, computed including that 3σ rejection. The histograms of the differences and scaled differences are also plotted but with a range set to $\pm$ 5 RMS around the mean, computed including the 3σ rejection. The average and RMS when using all the values, or when using that 3σ rejection, are indicated on these panels. The bottom panel shows the distribution of the corresponding scaled differences over the $\ell - \nu$ plane using a gradual color coding for each multiplet.

the asymmetries shows very good agreement, except where the measured asymmetry is somewhat suspect. Indeed, since one fits a finite set of modes, there are higher radial degree modes whose amplitude is too small to fit successfully. Therefore, these modes leak onto the fitted modes with the highest n, skewing some of their characteristics, especially the

measured asymmetry. Finally, the ratio of the mean power amplitudes of the multiplets, that is, the mean over m of the singlet power amplitude, shows some offset at the 2% level. This is to be expected since including the PSF or not directly affects the amplitude of the leaks and, thus, the measured mode amplitude.
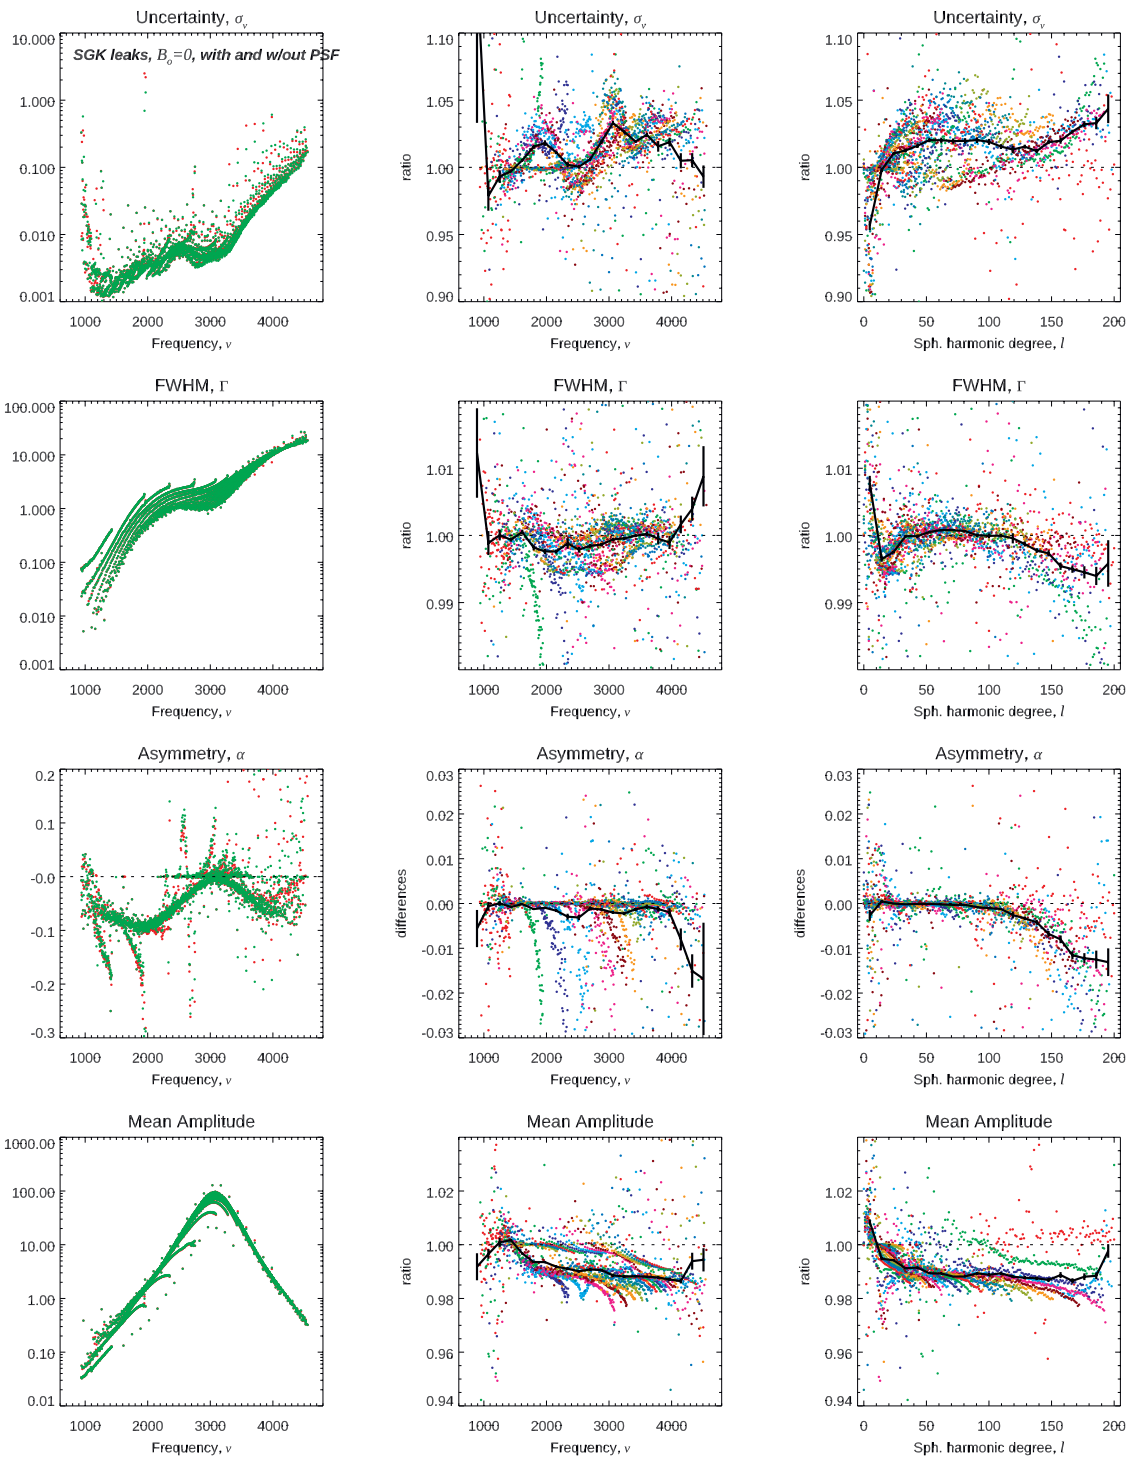
FIGURE 2
Comparison of multiplets' characteristics when fitting using a leakage matrix that either include or does not include the effect of the PSF. From top to bottom, the uncertainty on the frequency, FWHM, the asymmetry, and the mean amplitude are shown, color-coded on the leftmost column for the two cases, and as either ratios or differences on the two rightmost columns, namely, versus ν or versus ℓ. The thick black lines show the resulting binned values, while the color coding corresponds to the mode radial order, n.
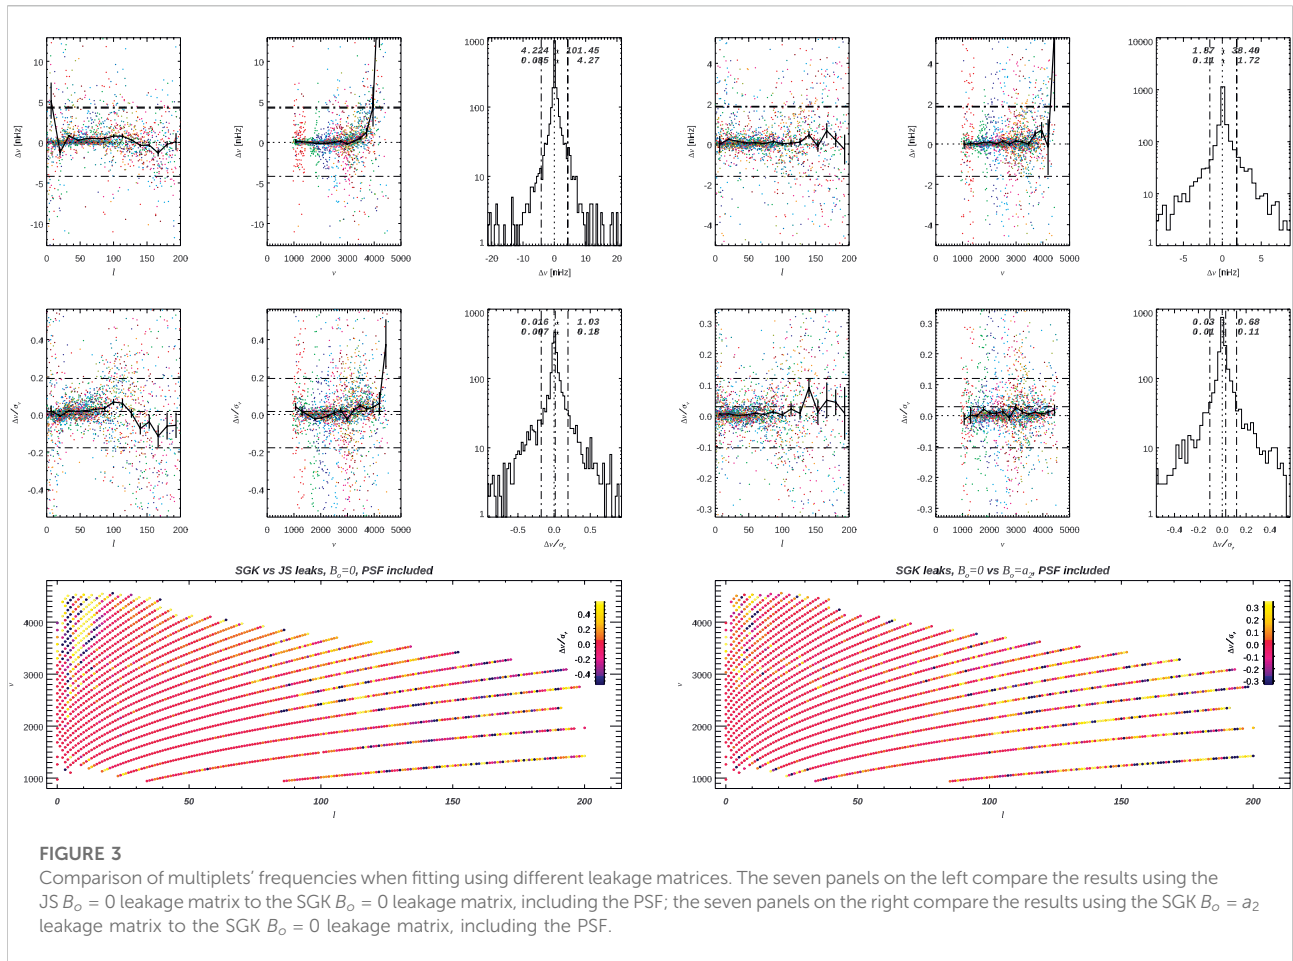
FIGURE 3
Comparison of multiplets' frequencies when fitting using different leakage matrices. The seven panels on the left compare the results using the JS $B_o = 0$ leakage matrix to the SGK $B_o = 0$ leakage matrix, including the PSF; the seven panels on the right compare the results using the SGK $B_o = a_2$ leakage matrix to the SGK $B_o = 0$ leakage matrix, including the PSF.

Figure 3 shows similar comparisons for the multiplet frequencies when using either 1) the leakage matrix computed by Schou (JS) or 2) the leakage matrix I computed for $B_o = a_2$ with the one I computed for $B_o = 0$, including in all cases the effect of a PSF. In those two cases, the difference and scaled differences are small, with the binned differences showing much smaller trends with ℓ or ν, although one can notice a larger scatter at higher degrees and systematic differences for the high-frequency low-degree modes. The comparisons of the mode characteristics (Figure 4) show scatter and systematic differences in the first case (JS *vs.* SGK leaks) and barely any in the second case ($B_o = 0$ *vs.* $B_o = a_2$).

By contrast, comparisons when using a symmetric or an asymmetric mode profile show much larger and significant differences, as shown in Figures 5, 6. Not surprisingly, the frequency differences show a clear pattern with a frequency that mimics the frequency dependence of the asymmetry at the 10–20σ levels.

Having also fitted shorter time series, I can compare the averaged values computed from fitting two abutting 4,608-day-long time series to the corresponding 9,216-day-long time series. This comparison is shown in Figures 7, 8, and it shows puzzling systematic differences with clear trends in ℓ and ν, skewed distributions, and scaled differences at the few σ levels. The differences in asymmetry, although small, trend with ℓ and, thus, are likely to be the main reason for the differences in frequencies.

In order to compare fitting GONG observations to fitting MDI or HMI, I used shorter coeval time series, namely, a 4,608-day-long time series for the MDI *vs.* GONG comparison, starting on 1996.05.01, and a 2,304-day-long time series for the HMI *vs.* GONG comparison, starting on 2012.02.07.

Figures 9, 10 present the comparison of MDI *vs.* GONG, showing again trends with ℓ and especially with ν when comparing frequencies, although they remain below the 1σ level. Estimates of the FWHM and the asymmetry show discrepancies at few percentages with a stronger trend with ℓ, while the measured mean power amplitudes are quite different. Figures 11, 12 present the HMI *vs.* GONG comparison, showing similar trends with ℓ and ν when comparing frequencies to the MDI *vs.* GONG comparison. However, comparisons of Γ and α present different trends.

FIGURE 4
Comparison of multiplets' characteristics when fitting using different leakage matrices. The eight panels on the left compare the results using the JS $B_o = 0$ leakage matrix to the SGK $B_o = 0$ leakage matrix, including the PSF; the eight panels on the right compare the results using the SGK $B_o = a_2$ leakage matrix to the SGK $B_o = 0$ leakage matrix, including the PSF.
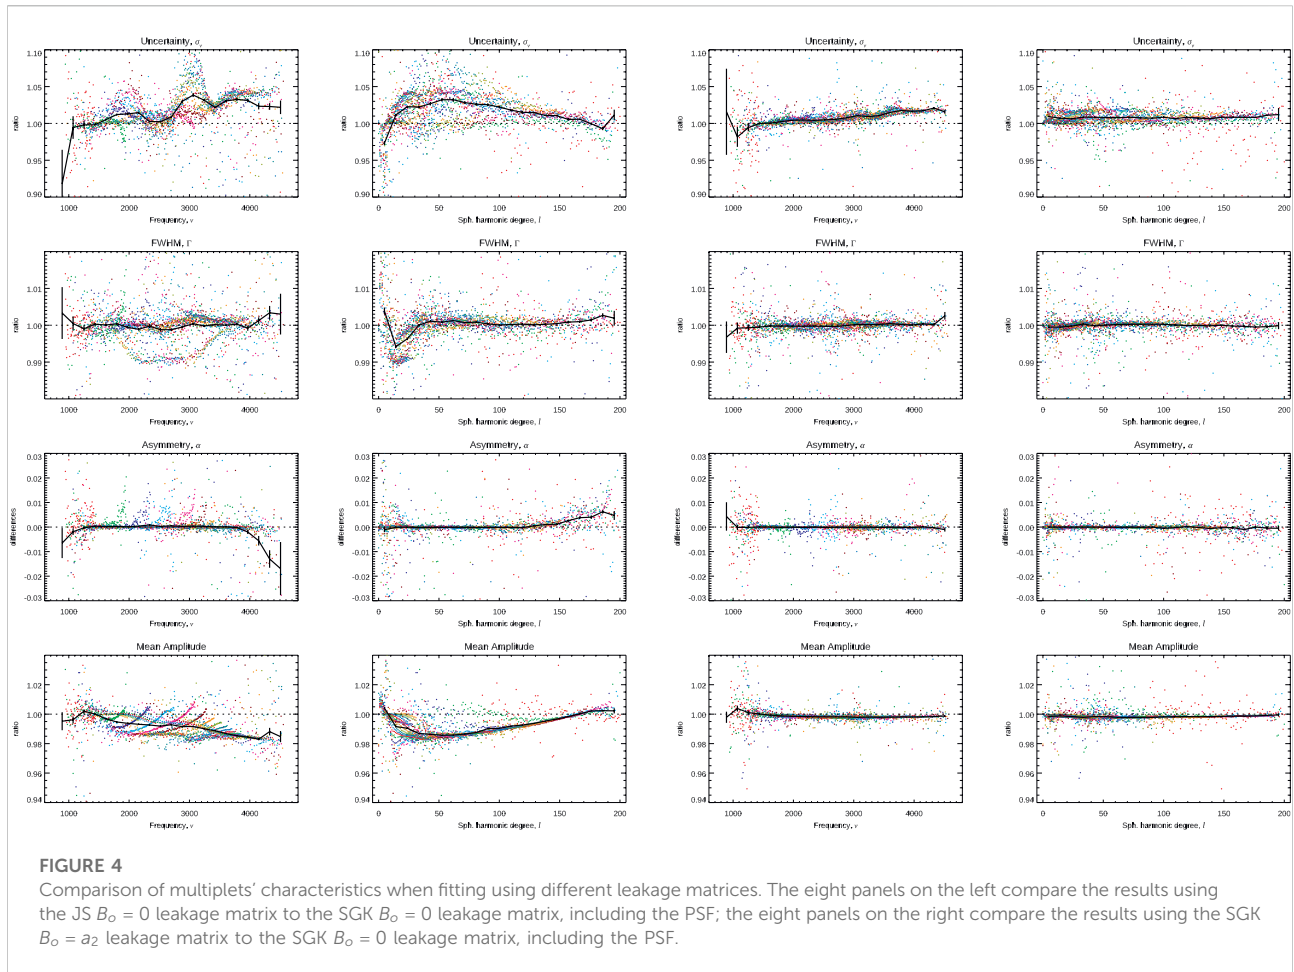
Table 1 summarizes in a numerical form, eight comparisons, where the mean and RMS of the differences and mean and RMS of the scaled differences are tabulated for both singlets, $\nu_{n,\ell,m}$, and multiplets, $\nu_{n,\ell}$. It also tabulates the fraction of modes whose differences exceed 1, 3, or 5σ. Similarly, Table 2 summarizes the scaled differences for some mode characteristics, namely, the FWHM and the asymmetry, for the same eight comparisons, and again the fraction of modes whose differences exceed 1, 3, or 5σ.

Finally, Figure 13 shows the comparison of the uncertainties on the multiplet frequencies, σ_ν, resulting from fitting time series of different lengths. As expected, the resulting precision overall scales with the square root of the length of the time series, although one can observe deviations at low and high frequencies.

4.2 Rotation inversion comparisons

While the previous section presents comparisons of mode characteristics, more importantly, one should consider whether inferences when using results from different fittings lead to significantly different estimates of solar properties.

Figure 14 presents the mean rotation rate resulting from inverting individual rotation splittings when estimated using different leakage matrices and obtained using three different inversion methodologies. Figure 15 shows the differences in the inferred profiles when using a different inversion methodology and the same set of rotation splittings, while Figure 16 shows the differences when using different splittings but the same methodology.

At first glance, the inverted rotation profiles are quite similar, although differences in inferences near the center are evident when different methods were used. This should be contrasted with the difference in the inferred thickness of the tachocline, and both differences are the result of the different ways smoothness is included in the inversion techniques. It is quite reasonable to assume that the variations near the center of some solutions are the results of noise propagation (this has been confirmed when inverting simulated data, Christensen-Dalsgaard, in preparation), while the resulting trade-off is less smooth near the tachocline.
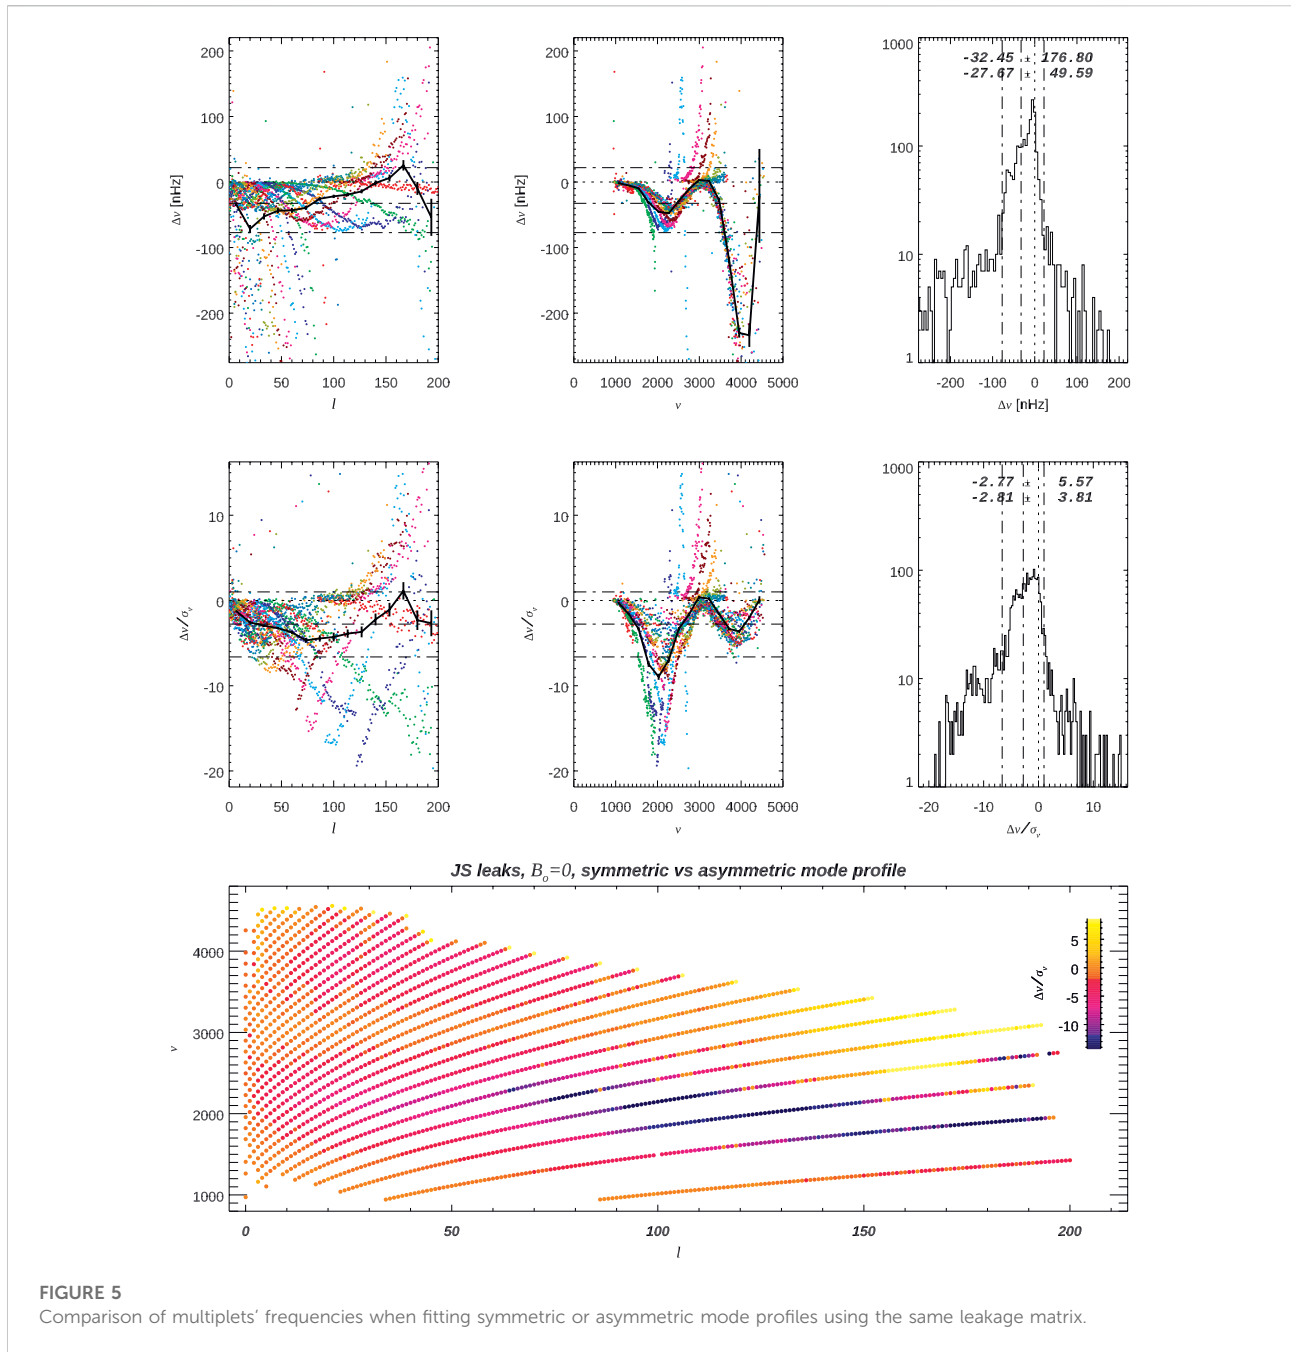
FIGURE 5
Comparison of multiplets' frequencies when fitting symmetric or asymmetric mode profiles using the same leakage matrix.

More interesting, in the context of this study, are the differences resulting from using different splittings. The differences in the top two rows of Figure 16 are quite similar and do not vary much with the inversion methodology (columns). These two correspond to splittings computed by either changing B_o or changing the leakage computation methodology. The bottom row shows differences that also do not vary much with the inversion methodology and are confined to the near center when using or not using a model of PSF.

5 Discussion

The inclusion of PSF in the leakage matrix causes the largest changes in derived mode frequencies, with all other things remaining constant. The effects of using a different B_o or a different formalism to compute the leakage matrix are much smaller. This should not come as a surprise since the inclusion of the PSF skews the ℓ dependence of the leakage, and since GONG is a ground-based set of
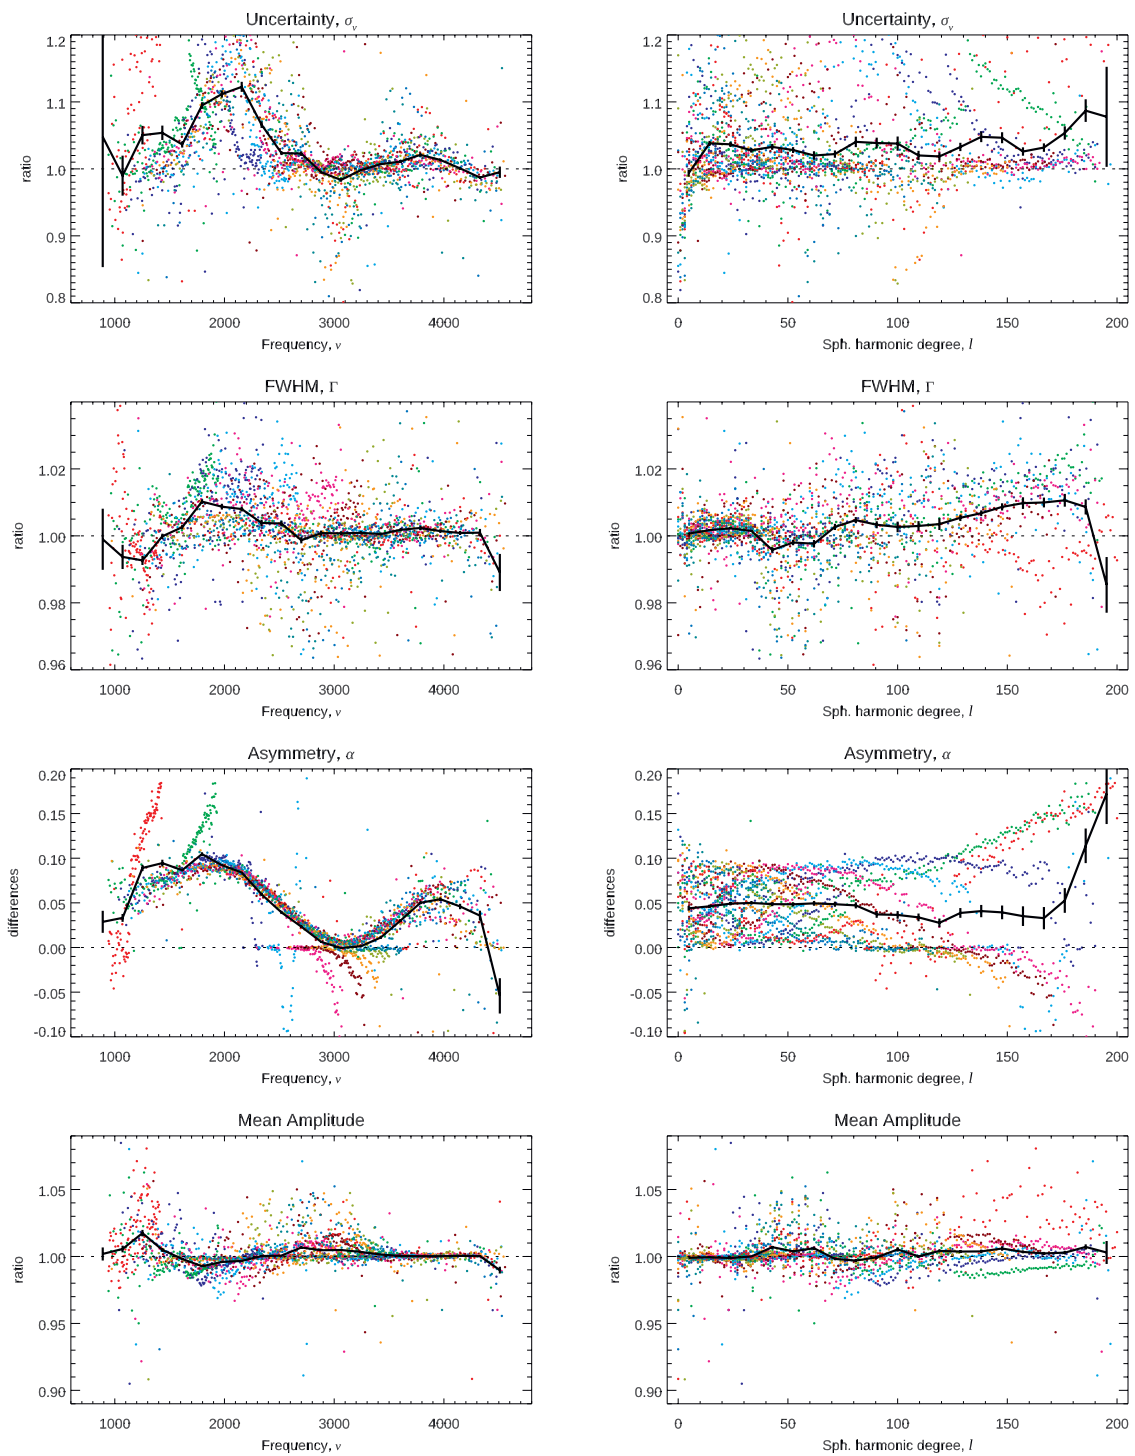
FIGURE 6
Comparison of mode characteristics, σ_ν, Γ, α, and $\bar{A}$ ratios or differences, when fitting symmetric or asymmetric mode profiles, using the same leakage matrix.
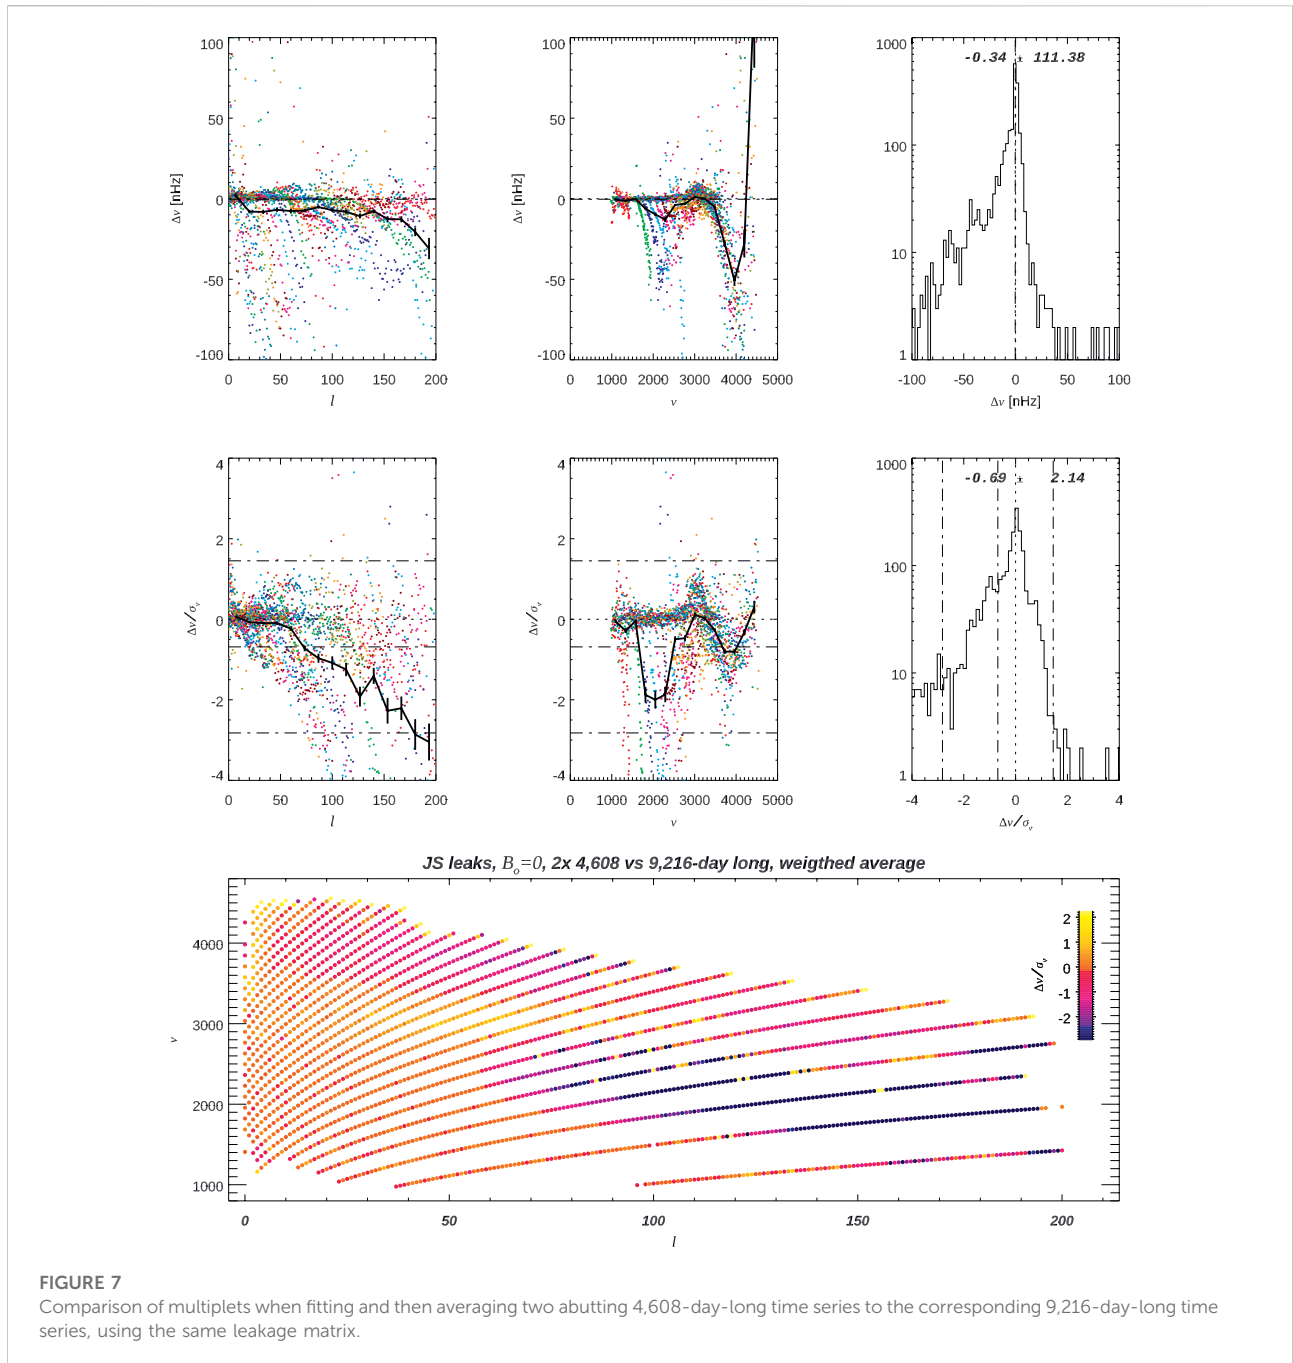
FIGURE 7
Comparison of multiplets when fitting and then averaging two abutting 4,608-day-long time series to the corresponding 9,216-day-long time series, using the same leakage matrix.

instruments, the effect of PSF on the images is non-negligible. However, other changes in the computation of the leakage matrix lead to small but systematic changes. Not surprisingly, ignoring the mode profile asymmetry leads to large differences with a systematic variation in frequency. By contrast, rotation inference differences are more pronounced when using a different B_0 or a different formalism and are dominant deeper into the interior when neglecting PSF.

More counter-intuitive is the comparison of the weighted mean frequencies[8] resulting from fitting two consecutive shorter time series to the results from fitting the corresponding longer one. It suggests that, besides the fundamental problem of mode

---

8 The weights used for those weighted means are the frequency error bars, although using an unweighted mean produces very similar results.
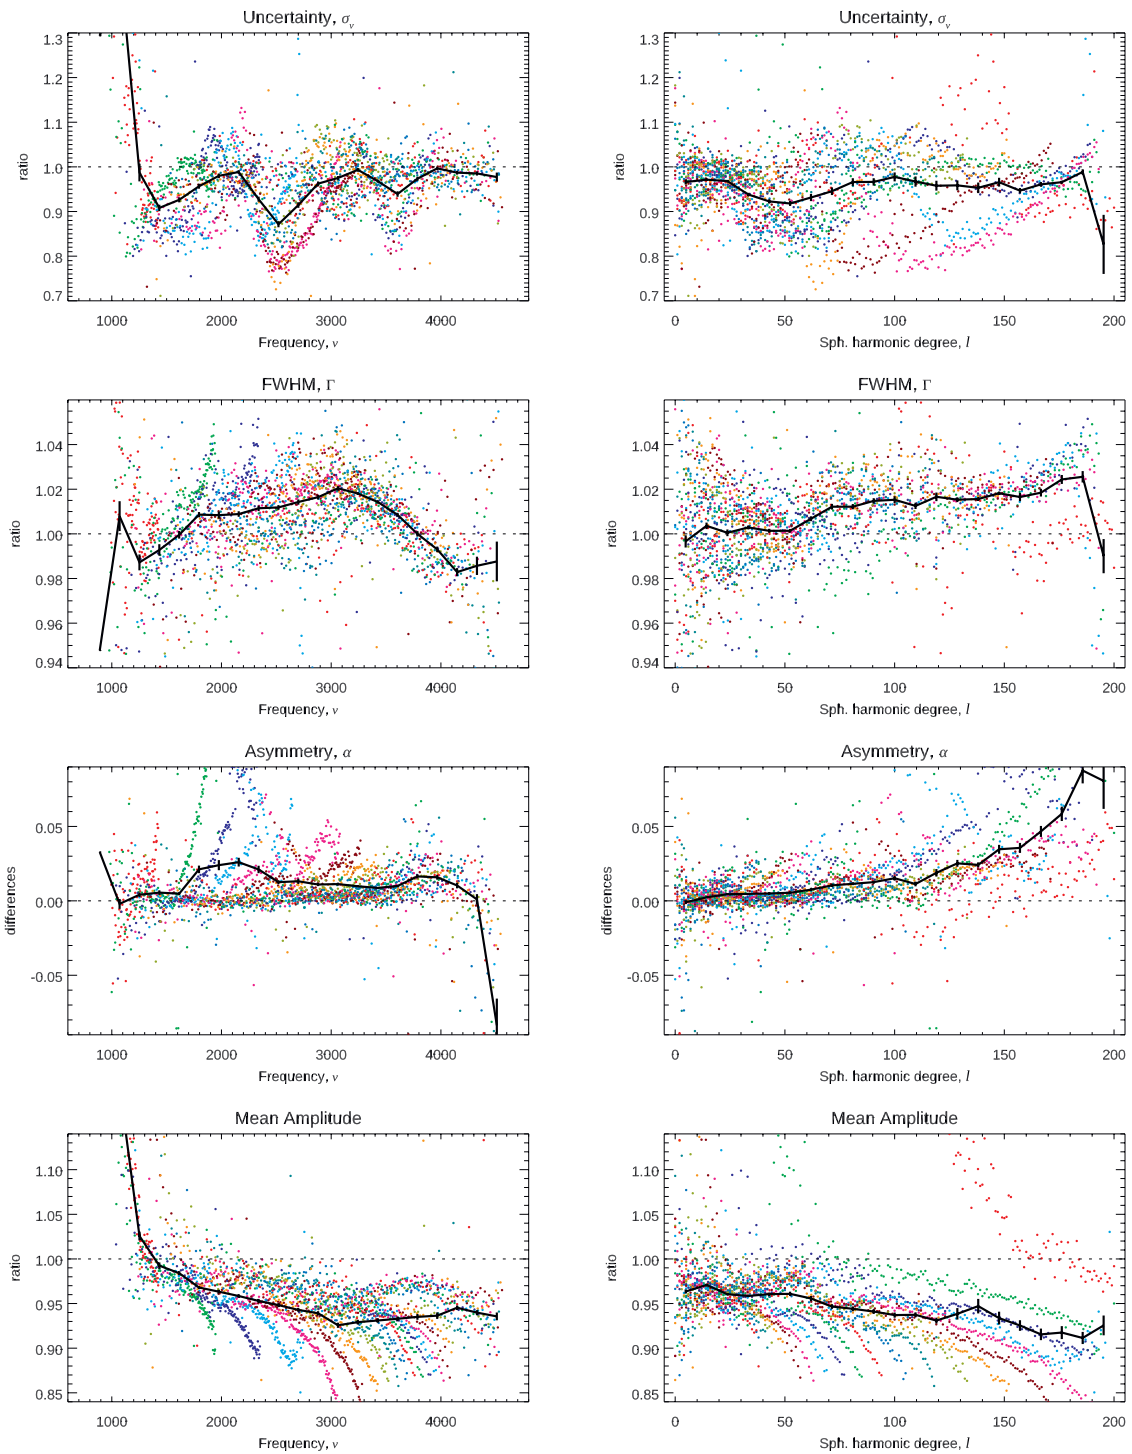
FIGURE 8
Comparison of mode characteristics, σ_ν, Γ, α, and $\bar{A}$ ratios or differences, when fitting and then averaging two abutting 4,608-day-long time series to the corresponding 9,216-day-long time series, using the same leakage matrix.
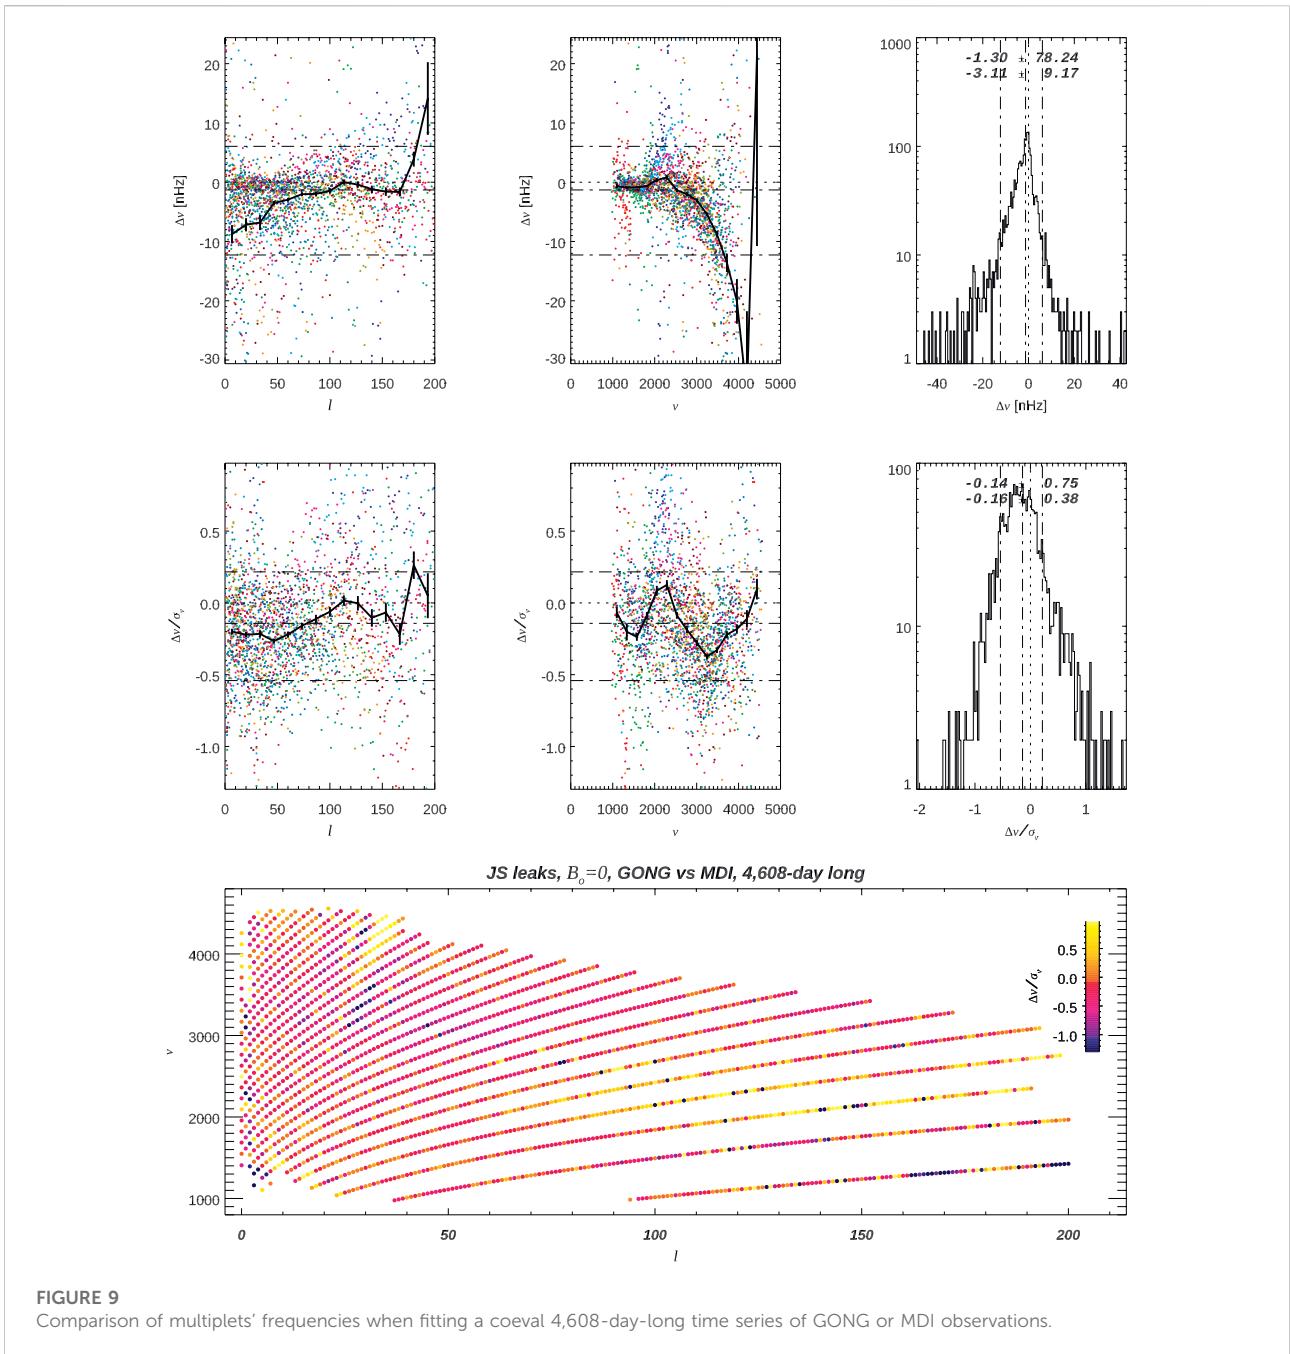
FIGURE 9
Comparison of multiplets' frequencies when fitting a coeval 4,608-day-long time series of GONG or MDI observations.

attrition, one gains precision when leveraging longer time series versus simply averaging results from shorter time series.

Less surprising are the comparisons of results from fitting coeval long time series of GONG and MDI or GONG and HMI observations. Although the leakage matrices were computed using the same formalism, these are quite different and thus the prime suspects for the cause of the differences, not to mention that the PSF and image distortion of the MDI instrument are not perfectly known, while using a mean PSF based only on the limb measurements for six similar but distinct GONG instruments is most likely an oversimplification.
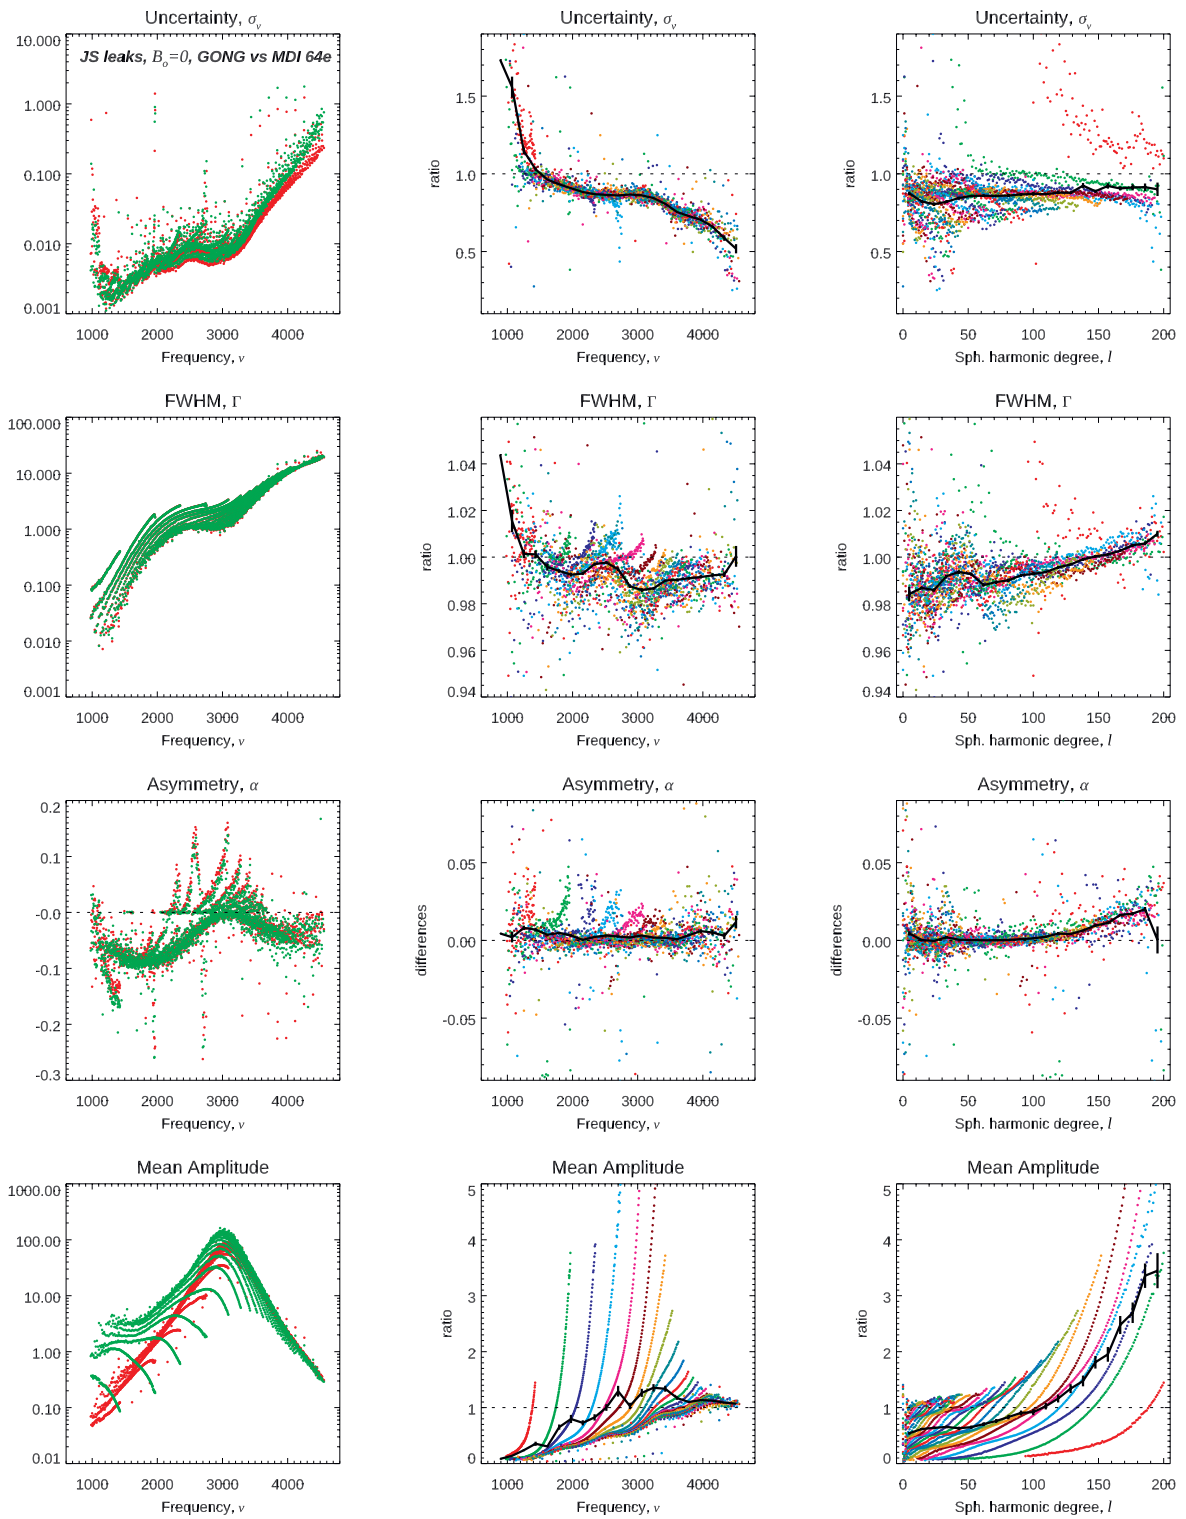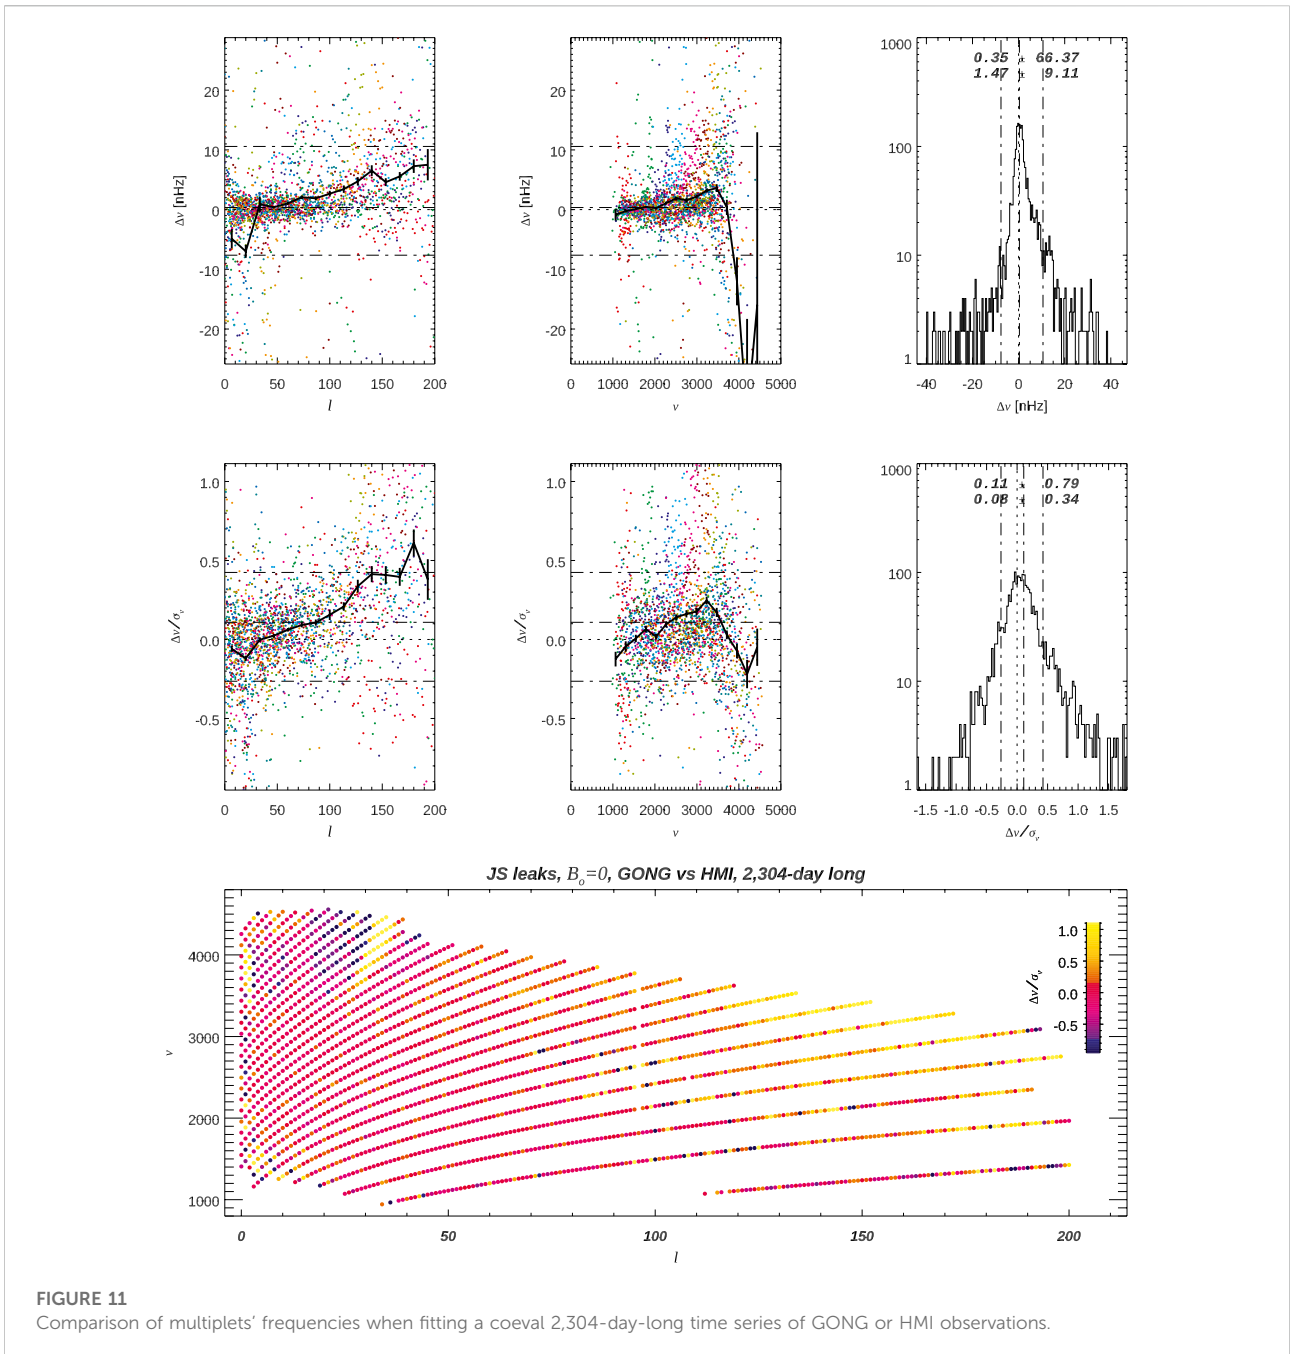
FIGURE 10
Comparison of mode characteristics, σ_ν, Γ, α, and $\bar{A}$ ratios or differences, when fitting a coeval 4,608-day-long time series of GONG or MDI observations.

FIGURE 11
Comparison of multiplets' frequencies when fitting a coeval 2,304-day-long time series of GONG or HMI observations.

Finally, comparing inverted rotation profiles also leads to systematic differences, independent of the inversion methodology. One main conclusion from these comparisons is that one must always keep in mind that despite the exquisite precision of modern helioseismic measurements, there remain systematic errors that are not included in the formal uncertainties and are too often ignored when interpreting helioseismic inferences.

Another challenge worth mentioning is how one could combine MDI and HMI observations in a single, very long time series of spherical harmonics. Firstly, the observing cadence is different, 60 s for MDI and 45 s for HMI, but
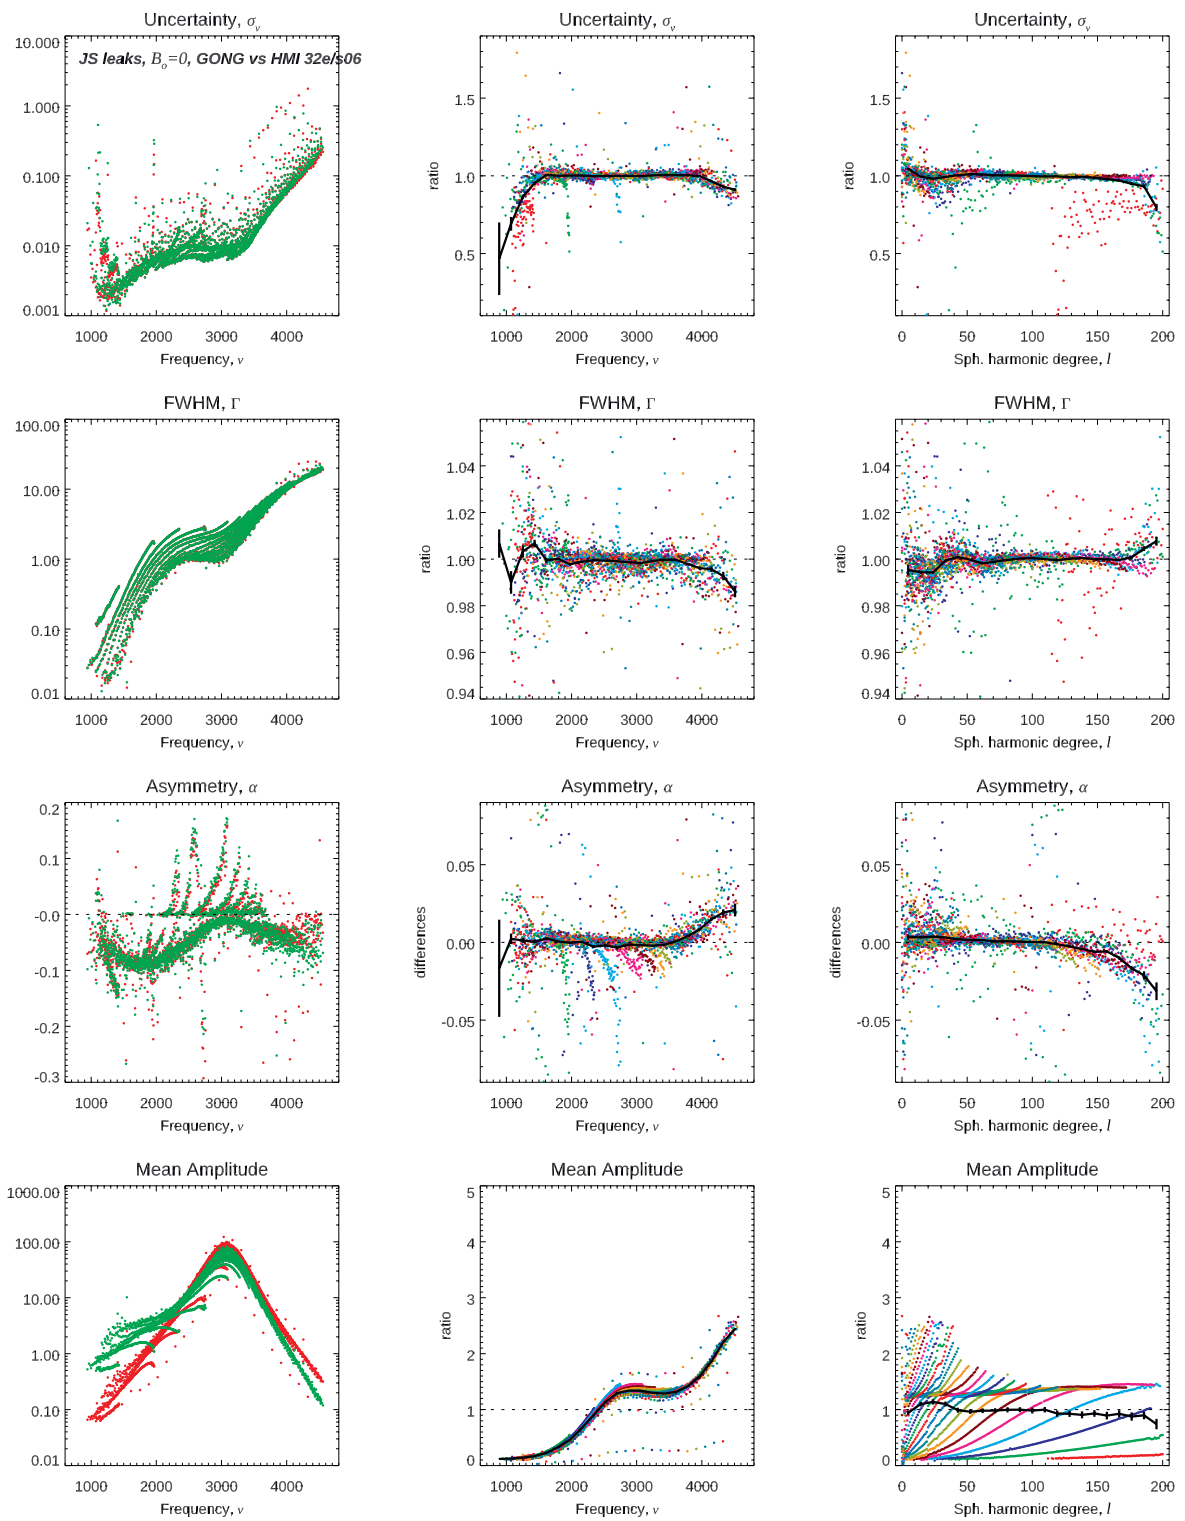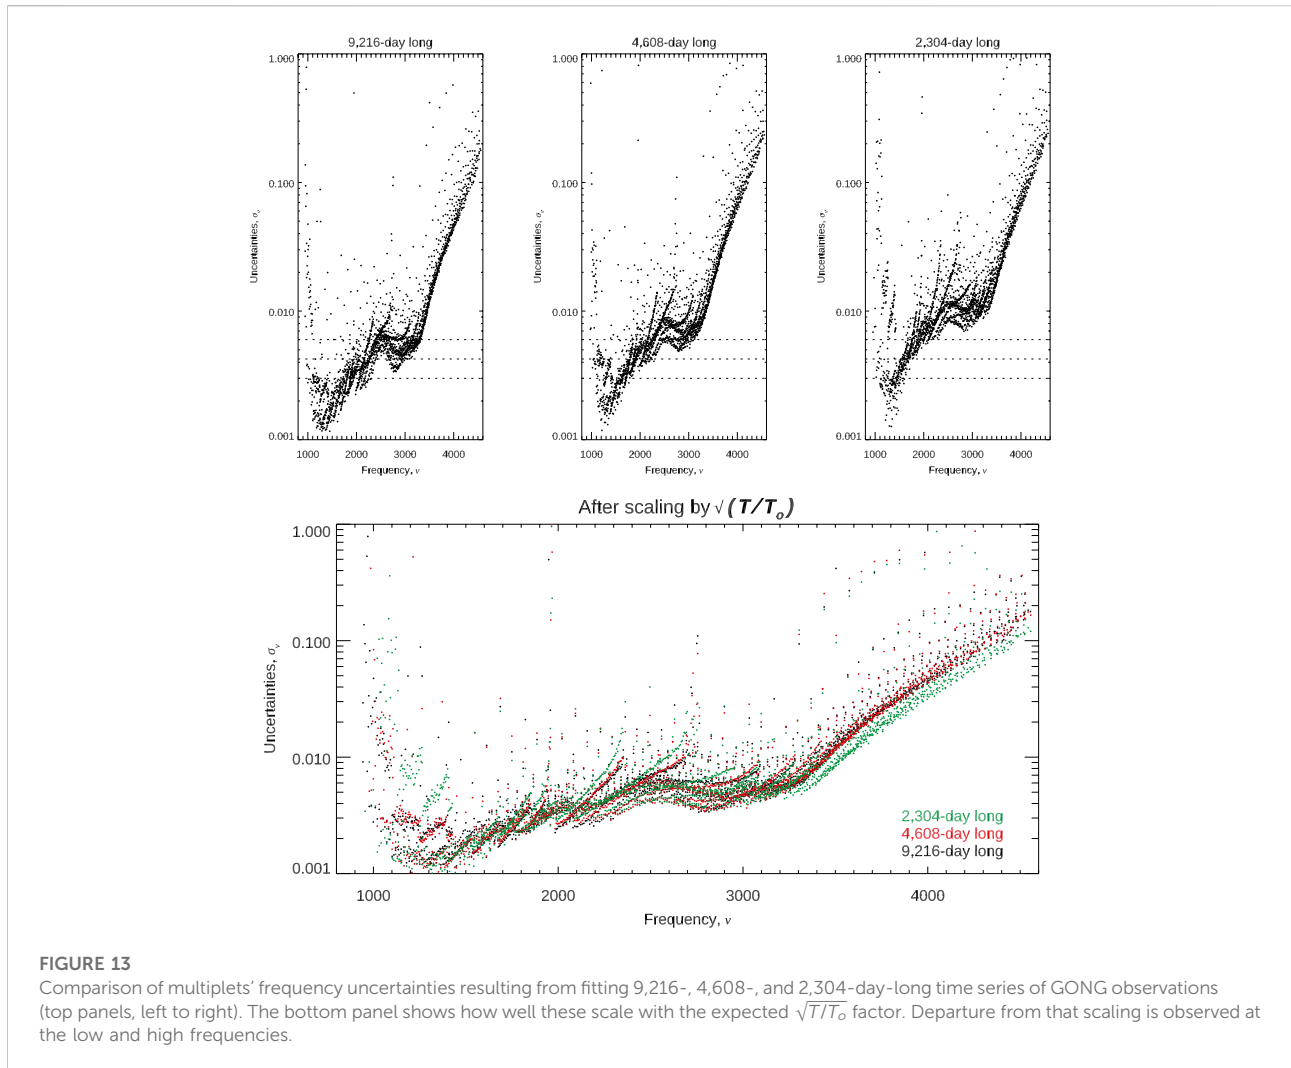
FIGURE 12
Comparison of mode characteristics, σ_ν, Γ, α, and $\bar{A}$ ratios or differences, when fitting a coeval 2,304-day-long time series of GONG or HMI observations.

FIGURE 13
Comparison of multiplets' frequency uncertainties resulting from fitting 9,216-, 4,608-, and 2,304-day-long time series of GONG observations (top panels, left to right). The bottom panel shows how well these scale with the expected $\sqrt{T/T_o}$ factor. Departure from that scaling is observed at the low and high frequencies.

secondly, the instruments measure the surface Doppler shift using a different absorption line (the Ni I photospheric absorption line at 6767.8 Å *vs*. the Fe I absorption line at 6173 Å); hence, the observed mode amplitude distribution is different. Finally, the spatial resolution of the dopplergrams is quite different, resulting in a very different spatial leakage. Either one could try to use a weighted leakage matrix or one would need to degrade the HMI dopplergrams to match the on-board Gaussian smoothing of the MDI observations, although the unknown PSF of MDI and its image distortion would be an additional difficult issue that would also need to be addressed.

Mode fitting remains challenging when aiming for the best possible precision and accuracy, especially when using long and very long time series, since uncertainties become small and residual systematics become significant. In addition to the contribution of the leakage matrix and the impact of including or not the asymmetry of the mode profile, other subtle effects inject their own biases. However, I remain convinced that several venues can

be pursued to improve mode fitting techniques. On one hand, a comprehensive analysis of our estimates of the mode fitting error bars remains warranted, and such an error bar validation can be achieved using Monte Carlo simulations. On the other hand, alternate approaches to fitting time series of various lengths ought to be explored, using techniques that reduce the realization noise. For instance, as suggested by Kosovichev (private communication), "time shifting" might beat down the realization noise, that is, using several time series of a set length but offset by just a small fraction of that length and then combined. Systematic errors set aside, it is ultimately the realization noise that limits the mode fitting precision, and after over two decades of observations, we can no longer count on the $\sqrt{T}$ effect to improve our precision.

Hence, a thorough review of every step of the global helioseismology data analysis pipeline, while not very glamorous, is still needed to fully exploit the potential of the observations we have on hand.
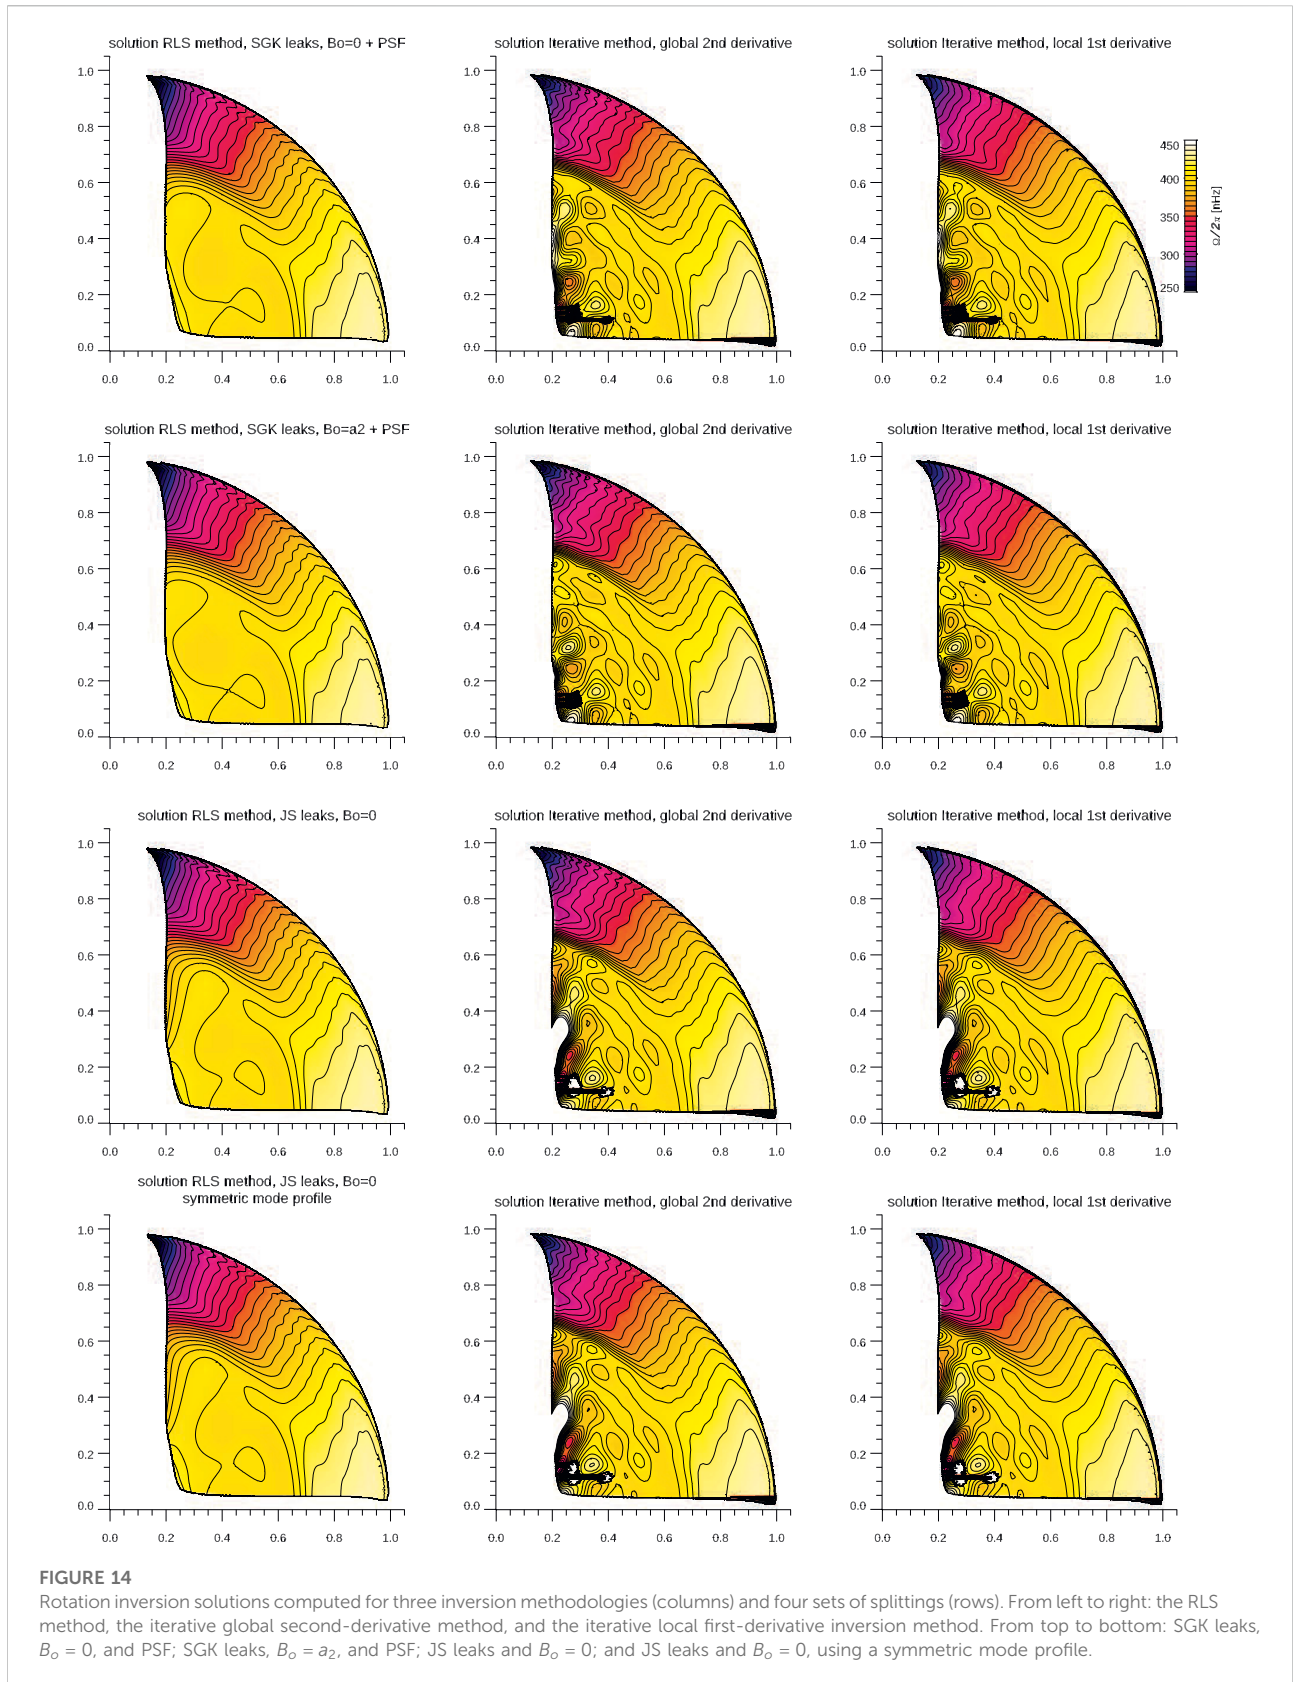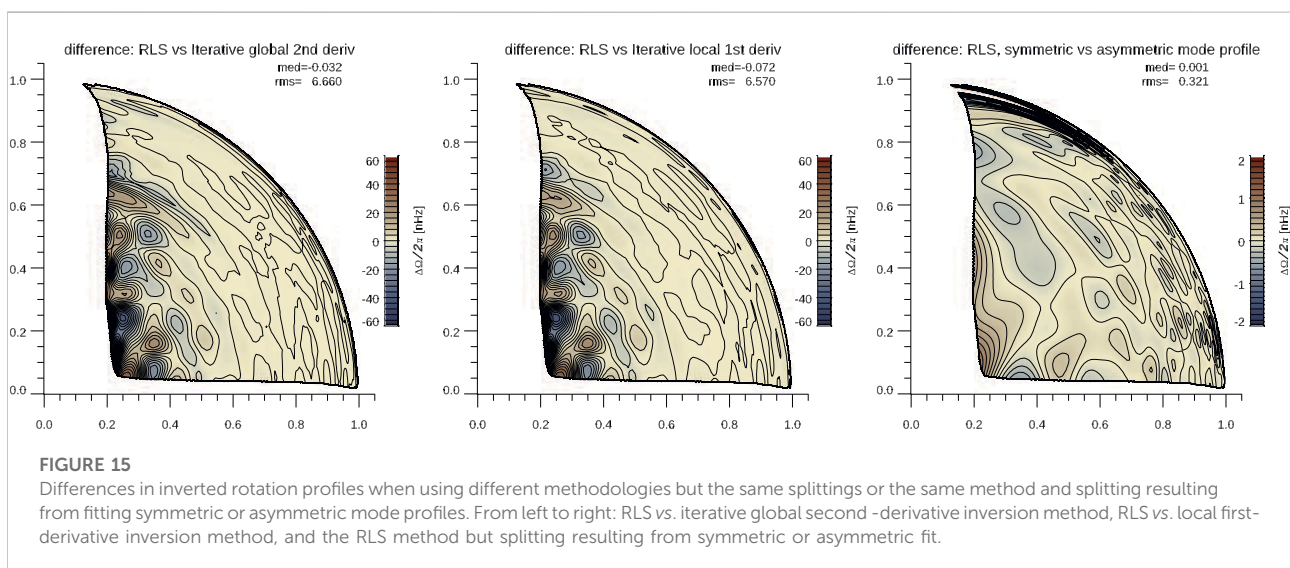
FIGURE 14
Rotation inversion solutions computed for three inversion methodologies (columns) and four sets of splittings (rows). From left to right: the RLS method, the iterative global second-derivative method, and the iterative local first-derivative inversion method. From top to bottom: SGK leaks, $B_o = 0$, and PSF; SGK leaks, $B_o = a_2$, and PSF; JS leaks and $B_o = 0$; and JS leaks and $B_o = 0$, using a symmetric mode profile.

TABLE 1 Overall statistics when comparing frequencies derived using different fitting leakage matrices of time series.

	Difference [nHz]	Scaled difference	#	%> 1σ	%> 3σ	%> 5σ
SGK leaks, and $B_o = 0$, with and without PSF						
Singlets	2.605 ± 51.5	0.018 ± 0.148	325156	0.5	0.0	0.0
Multiplets	2.165 ± 83.9	0.112 ± 1.146	2354	8.7	1.9	1.1
SGK *vs.* JS leaks, $B_o = 0$, and PSF included						
Singlets	0.794 ± 58.6	−0.004 ± 0.139	325142	0.4	0.0	0.0
Multiplets	4.224 ± 101.7	0.016 ± 1.029	2352	4.6	1.6	1.2
SGK leaks, $B_o = 0$ *vs.* $B_o = a_2$, and PSF included						
Singlets	0.840 ± 27.1	0.004 ± 0.093	325358	0.1	0.0	0.0
Multiplets	1.871 ± 38.4	0.029 ± 0.674	2355	2.5	0.8	0.3
JS leaks, $B_o = 0$, and symmetric *vs.* asymmetric mode profiles						
Singlets	−20.827 ± 115.6	−0.308 ± 0.682	322453	13.1	0.1	0.0
Multiplets	−32.450 ± 176.8	−2.773 ± 5.574	2351	77.3	45.9	25.4
JS leaks, $B_o = 0$, and weighted average of two consecutive 4,608-day-long time series *vs.* one 9,216-day-long time series						
Singlets	−7.672 ± 66.0	−0.101 ± 0.282	318557	1.4	0.0	0.0
Multiplets	−0.342 ± 111.4	−0.686 ± 2.137	2334	28.1	10.4	5.8
JS leaks, $B_o = 0$, GONG *vs.* MDI, and 4,608-day-long time series						
Singlets	−1.822 ± 71.8	−0.017 ± 0.246	320437	0.3	0.0	0.0
Multiplets	−1.304 ± 78.2	−0.143 ± 0.750	2342	6.5	1.2	0.4
JS leaks, $B_o = 0$, GONG *vs.* HMI, and 2,304-day-long time series						
Singlets	2.988 ± 55.4	0.017 ± 0.176	312610	0.1	0.0	0.0
Multiplets	0.345 ± 66.4	0.109 ± 0.787	2321	7.7	1.8	0.4

FIGURE 15
Differences in inverted rotation profiles when using different methodologies but the same splittings or the same method and splitting resulting from fitting symmetric or asymmetric mode profiles. From left to right: RLS *vs.* iterative global second-derivative inversion method, RLS *vs.* local first-derivative inversion method, and the RLS method but splitting resulting from symmetric or asymmetric fit.
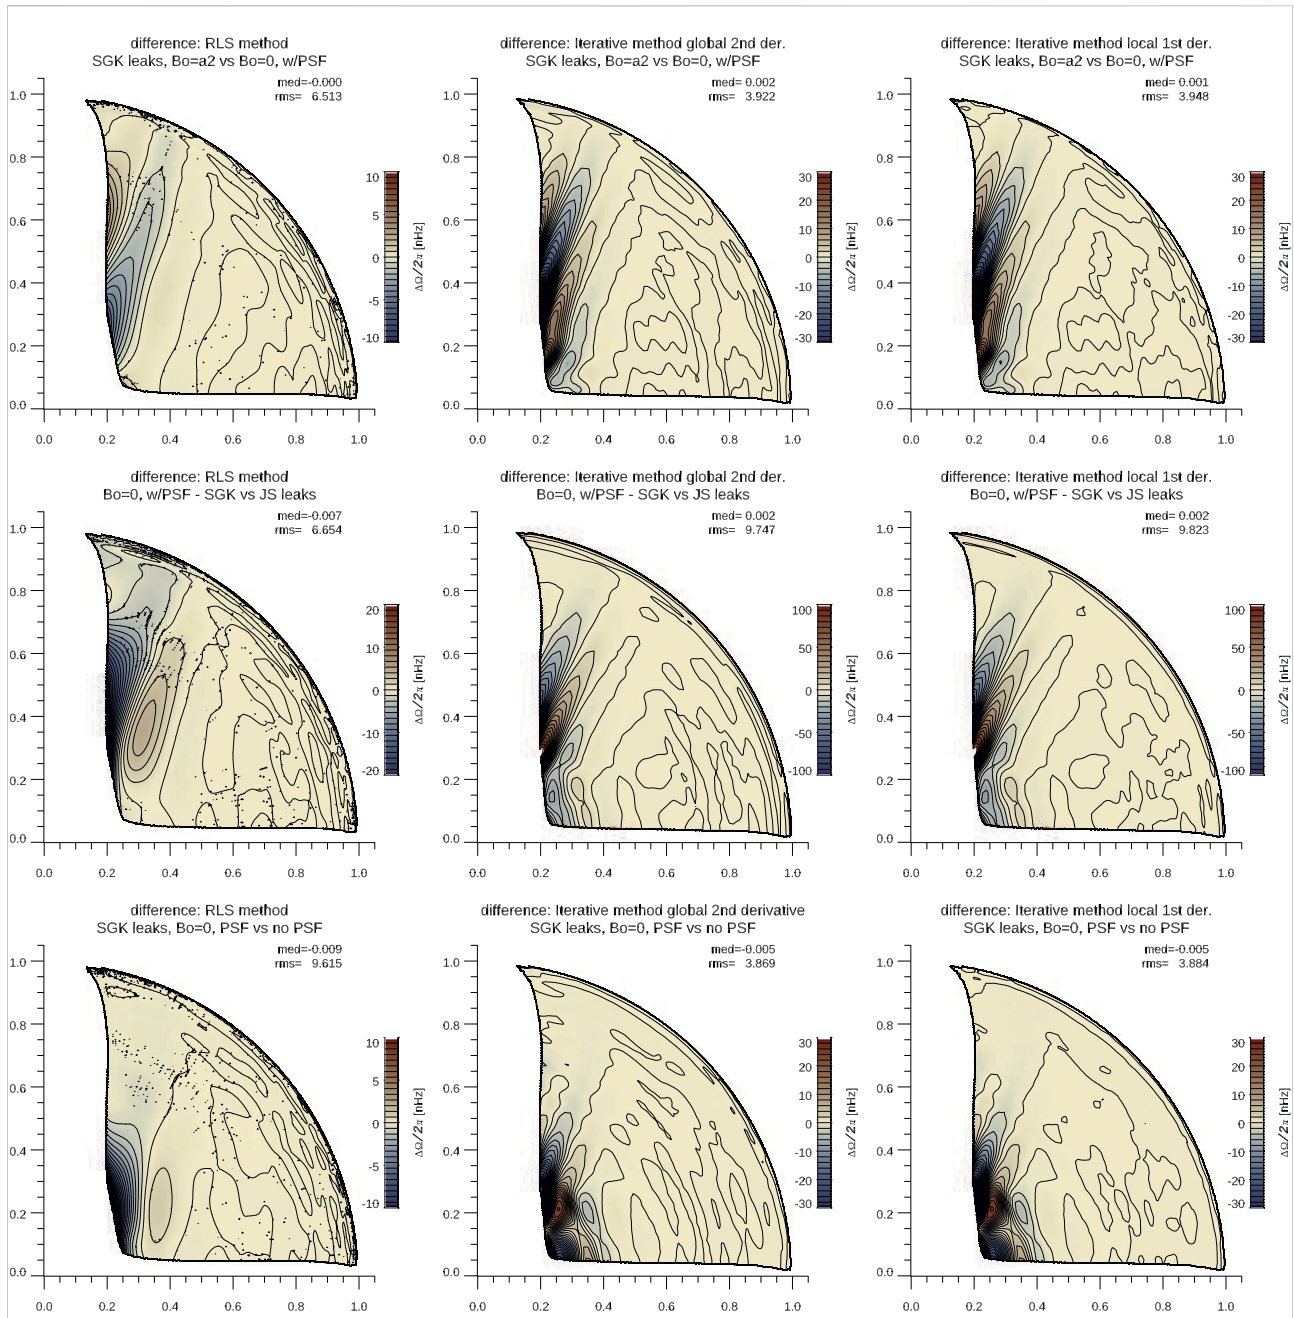
FIGURE 16
Differences in inverted rotation profiles when using different methods (rows) and different splittings (columns). Left to right: the RLS method, the iterative global second-derivative method, and the iterative local first-derivative method. Top to bottom: SGK leaks, including the PSF but with $B_o = a_2$ or $B_o = 0$; SGK vs. JS leaks, including the PSF and with $B_o = 0$; and SGK leaks, $B_o = 0$ but with or without including the PSF.

TABLE 2 Overall statistics when comparing mode characteristics derived using different fitting leakage matrices of time series.

	Scaled difference	%> 1σ	%> 3σ	%> 5σ
SGK leaks, $B_o = 0$, and with and without PSF				
FWHM	-0.056 ± 0.374	3.40	0.08	<0.01
Asymmetry	-0.185 ± 1.324	10.50	2.42	1.36
SGK vs. JS leaks, $B_o = 0$, and PSF included				
FWHM	0.000 ± 0.321	1.66	0.04	<0.01
Asymmetry	0.017 ± 1.143	5.23	1.66	1.02
SGK leaks, $B_o = 0$ vs. $B_o = a_2$, and PSF included				
FWHM	0.001 ± 0.166	0.47	<0.01	<0.01
Asymmetry	-0.030 ± 0.694	1.74	0.72	0.26
JS leaks, $B_o = 0$, and symmetric vs. asymmetric mode profiles				
FWHM	0.118 ± 0.773	13.87	0.85	0.04
Asymmetry	n/a			
JS leaks, $B_o = 0$, and weighted average of two consecutive 4,608-day-long time series vs. one 9,216-day-long time series				
FWHM	0.495 ± 0.923	23.56	0.69	0.04
Asymmetry	1.138 ± 2.300	51.50	9.08	5.53
JS leaks, $B_o = 0$, GONG vs. MDI, and 4,608-day-long time series				
FWHM	-0.344 ± 0.477	4.82	<0.01	<0.01
Asymmetry	0.181 ± 0.805	6.83	1.37	0.51
JS leaks, $B_o = 0$, GONG vs. HMI, and 2,304-day-long time series				
FWHM	-0.033 ± 0.222	0.60	<0.01	<0.01
Asymmetry	-0.007 ± 0.765	5.90	1.21	0.52

Data availability statement

The data sets analyzed in this study can be found at the GONG Data Archive at the National Solar Observatory and at the Joint Science Operations Center at Stanford University. The results of this analysis are available at the author's website and the author's Dataverse repository. The leakage matrices are available upon request. Further inquiries can be directed to the corresponding author.

Author contributions

SGK contributed solely to the design, implementation, and execution of the mode fitting. The spherical harmonic coefficients used in this work were produced by the GONG, MDI, and HMI projects, while Jesper Schou provided some of the leakage matrices used in this work. The design of the inversion methods is a collaborative effort with Antonio Eff-Darwich, while the implementation and execution of the inversions is solely the work of SGK.

Funding

Acknowledgments

Conflict of interest

Publisher's note

References

Anderson, E. R., Duvall, J., Thomas, L., and Jefferies, S. M. (1990). Modeling of solar oscillation power spectra. *Astrophysical J.* 364, 699. doi:10.1086/169452

Beckers, J. M., and Brown, T. M. (1978). *Oss. Mem. D. Oss. Astrofis. D. Arcetri* 106, 189.

Duvall, T. L., Jr., Jefferies, S. M., Harvey, J. W., Osaki, Y., and Pomerantz, M. A. (1993). Asymmetries of solar oscillation line profiles. *Astrophysical J.* 410, 829. doi:10.1086/172800

Eff-Darwich, A., Korzennik, S. G., and García, R. A. (2010). Advances in solar rotation rate inferences: Unstructured grid inversions and improved rotational splittings. *Astron. Nachrichten* 331, 890–895. doi:10.1002/asna.201011420

Eff-Darwich, A., and Perez Hernandez, F. (1997). A new strategy for helioseismic inversions. *Astron. Astrophys. Suppl. Ser.* 125, 391–398. doi:10.1051/aas:1997229

Fahlman, G. G., and Ulrych, T. J. (1982). A new method for estimating the power spectrum of gapped data. *Mon. Not. R. Astron. Soc.* 199, 53–65. doi:10.1093/mnras/199.1.53

Gregor, J., and Fessler, J. A. (2015). Comparison of SIRT and SQS for regularized weighted least squares image reconstruction. *IEEE Trans. Comput. imaging* 1, 44–55. doi:10.1109/TCI.2015.2442511

Harvey, J. W., Hill, F., Hubbard, R. P., Kennedy, J. R., Leibacher, J. W., Pintar, J. A., et al. (1996). The global oscillation network Group (GONG) project. *Science* 272, 1284–1286. doi:10.1126/science.272.5266.1284

Hill, F., Anderson, E., Howe, R., Jefferies, S. M., Komm, R., and Toner, C. (1998). "Estimated mode parameters from the fitting of GONG spectra," in *Structure and Dynamics of the interior of the Sun and sun-like stars. Vol. 418 of ESA special publication*. Editor S. Korzennik, 231.

Howe, R., and Hill, F. (1998). "Estimating low-degree mode parameters from GONG data using the leakage matrix," in *Structure and Dynamics of the interior of the Sun and sun-like stars. Vol. 418 of ESA special publication*. Editor S. Korzennik, 237.

Howe, R., and Thompson, M. J. (1998). A strategy for fitting partially blended ridges in GONG solar p-mode spectra. *Astronomy Astrophysics Suppl.* 131, 539–548. doi:10.1051/aas:1998285

Komm, R. W., Anderson, E., Hill, F., Howe, R., Kosovichev, A. G., Scherrer, P. H., et al. (1998). "Comparison of SOHO-SOI/MDI and GONG spectra," in *Structure and Dynamics of the interior of the Sun and sun-like stars. Vol. 418 of ESA special publication*. Editor S. Korzennik, 253.

Korzennik, S. G. (2005). A mode-fitting methodology optimized for very long helioseismic time series. *Astrophysical J.* 626, 585–615. doi:10.1086/429748

Korzennik, S. G. (2013). "Mode frequencies from GONG, MDI, and HMI data," in *Fifty years of seismology of the Sun and stars. Vol. 478 of astronomical society of the pacific conference series*. Editors K. Jain, S. C. Tripathy, F. Hill, J. W. Leibacher, and A. A. Pevtsov, 137.

Korzennik, S. G. (2008a). YAOPBM - yet another peak bagging method. *Astron. Nachrichten* 329, 453–460. doi:10.1002/asna.200710979

Korzennik, S. G. (2008b). YAOPBM—II: Extension to higher degrees and to shorter time series. *J. Phys. Conf. Ser.* 118, 012082. doi:10.1088/1742-6596/118/1/012082

Larson, T. P., and Schou, J. (2015). Improved helioseismic analysis of medium-R data from the Michelson Doppler imager. *Sol. Phys.* 290, 3221–3256. doi:10.1007/s11207-015-0792-y

Press, W. H., Flannery, B. P., and Teukolsky, S. A. (1986). *Numerical recipes. The art of scientific computing.*

Scherrer, P. H., Bogart, R. S., Bush, R. I., Hoeksema, J. T., Kosovichev, A. G., Schou, J., et al. (1995). The solar oscillations investigation - Michelson Doppler imager. *Sol. Phys.* 162, 129–188. doi:10.1007/BF00733429

Scherrer, P. H., Schou, J., Bush, R. I., Kosovichev, A. G., Bogart, R. S., Hoeksema, J. T., et al. (2012). The helioseismic and magnetic imager (HMI) investigation for the solar Dynamics observatory (SDO). *Sol. Phys.* 275, 207–227. doi:10.1007/s11207-011-9834-2

Schou, J. (1992). *On the analysis of helioseismic data*. Ph.D. thesis (Aarhus, Denmark: Aarhus University).

Title, A. M., and Ramsey, H. E. (1980). Improvements in birefringent filters. 6: Analog birefringent elements. *Appl. Opt.* 19, 2046–2058. doi:10.1364/AO.19.002046

Williams, W. E., Toner, C., and Hill, F. (1995). "Implementation of an mtf-based merging algorithm for GONG image data," in *Helioseismology. Vol. 376 of ESA special publication*. Editors J. T. Hoeksema, V. Domingo, B. Fleck, and B. Battrick, 185.

Woodard, M. F. (1989). Distortion of high-degree solar p-mode eigenfunctions by latitudinal differential rotation. *Astrophysical J.* 347, 1176. doi:10.1086/168206