Balmer Filaments in Tycho's Supernova Remnant: An Interplay between Cosmic-ray and Broad-neutral Precursors

Sladjana Knežević[1,10] , Ronald Läsker[2], Glenn van de Ven[3], Joan Font[4], John C. Raymond[5] , Coryn A. L. Bailer-Jones[3], John Beckman[4], Giovanni Morlino[6] , Parviz Ghavamian[7] , John P. Hughes[8], and Kevin Heng[9]

[1] Department of Particle Physics and Astrophysics, Faculty of Physics, The Weizmann Institute of Science, P.O. Box 26, Rehovot 76100, Israel
sladjana.knezevic@weizmann.ac.il
[2] Finnish Centre for Astronomy with ESO (FINCA), University of Turku, Väisäläntie 20, FI-21500 Kaarina, Finland
[3] Max Planck Institute for Astronomy, Königstuhl 17, D-69117, Heidelberg, Germany
[4] Instituto de Astrofísica de Canarias, Vía Láctea, La Laguna, Tenerife, Spain
[5] Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA
[6] INFN Gran Sasso Science Institute, viale F. Crispi 7, I-67100 L'Aquila, Italy
[7] Department of Physics, Astronomy and Geosciences Towson University, Towson, MD 21252, USA
[8] Department of Physics and Astronomy, Rutgers University, 136 Frelinghuysen Road, Piscataway, NJ 08854, USA
[9] University of Bern, Center for Space and Habitability, Sidlerstrasse 5, CH-3012, Bern, Switzerland

Abstract

We present Hα spectroscopic observations and detailed modeling of the Balmer filaments in the supernova remnant (SNR) Tycho (SN 1572). We used GHαFaS (Galaxy Hα Fabry–Pérot Spectrometer) on the William Herschel Telescope with a $3\farcm4 \times 3\farcm4$ field of view, $0\farcs2$ pixel scale, and $\sigma_{\mathrm{instr}} = 8.1\,\mathrm{km\,s^{-1}}$ resolution at $1''$ seeing for ∼10 hr, resulting in 82 spatial−spectral bins that resolve the narrow Hα line in the entire SN 1572 northeastern rim. For the first time, we can therefore mitigate artificial line broadening from unresolved differential motion and probe Hα emission parameters in varying shock and ambient medium conditions. Broad Hα line remains unresolved within spectral coverage of $392\,\mathrm{km\,s^{-1}}$. We employed Bayesian inference to obtain reliable parameter confidence intervals and to quantify the evidence for models with multiple line components. The median Hα narrow-line (NL) FWHM of all bins and models is $W_{\mathrm{NL}} = (54.8 \pm 1.8)\,\mathrm{km\,s^{-1}}$ at the 95% confidence level, varying within $[35, 72]\,\mathrm{km\,s^{-1}}$ between bins and clearly broadened compared to the intrinsic (thermal) ≈$20\,\mathrm{km\,s^{-1}}$. Possible line splits are accounted for, significant in ≈18% of the filament, and presumably due to remaining projection effects. We also find widespread evidence for intermediate-line emission of a broad-neutral precursor, with a median $W_{\mathrm{IL}} = (180 \pm 14)\,\mathrm{km\,s^{-1}}$ (95% confidence). Finally, we present a measurement of the remnant's systemic velocity, $V_{\mathrm{LSR}} = -34\,\mathrm{km\,s^{-1}}$, and map differential line-of-sight motions. Our results confirm the existence and interplay of shock precursors in Tycho's remnant. In particular, we show that suprathermal NL emission is near-universal in SN 1572, and that, in the absence of an alternative explanation, collisionless SNR shocks constitute a viable acceleration source for Galactic TeV cosmic-ray protons.

Key words: cosmic rays – ISM: individual objects (Tycho) – ISM: supernova remnants – shock waves

1. Introduction

Supernova remnant (SNR) shocks are suspected to be the long-sought Galactic cosmic-ray (CR) sources. Observational evidence for particle acceleration at work in SNRs has been seen in synchrotron emission from electrons, gamma radiation, signatures of amplified magnetic fields, higher compression ratios than predicted by jump conditions, and optical line profiles from shocks propagating in a partially ionized plasma (for a review, see Helder et al. 2012). Since most of the CRs that we detect on Earth consist of protons, finding evidence for accelerated CR protons in SNRs with energies up to the "knee" (∼1 PeV) is thus of particular importance. Insight into CR protons is possible through either the gamma-ray spectrum as a result of neutral pion decay (Ackermann et al. 2013) or the optical wavelength window by carefully analyzing Hα-line profiles (Helder et al. 2009, 2010; Nikolić et al. 2013). The latter shock emission, usually also referred to as optical emission from Balmer-dominated shocks (BDSs), is the subject of the study of this paper.

The spectra of BDSs, typically observed around SNRs that originate from SN Ia explosions, show the presence of strong

two-component hydrogen lines (Heng 2010). When a shock wave encounters partly ionized interstellar medium (ISM), the cold pre-shock hydrogen atoms overrun by the shock can either be excited by hot post-shock gas, resulting in the narrow Hα-component emission, or enter a charge exchange (CE) process with the hot post-shock plasma, producing hot neutrals whose collisional excitation then gives rise to the broad Hα component. The two-component Hα-line parameters provide valuable information on the existence of CR precursors in the shocks. A narrow line broadened beyond $10–20\,\mathrm{km\,s^{-1}}$ gives direct evidence of the presence of non-thermal particles in the shock precursor (Morlino et al. 2013). This is due to the fact that hydrogen atoms are ionized at temperatures larger than ≈10,000 K, but also because the lifetime of the neutral hydrogen in the post-shock region is too short for the collisional interaction to broaden the line profiles (Smith et al. 1994). The CRs will heat the cold neutrals in the ISM, resulting in the broadening of the narrow Hα line, but the CRs will also reduce the broad Hα-line width by removing energy from the protons in the post-shock region.

Several authors proposed that narrow-line broadening can also arise from a broad-neutral (BN) precursor: hot neutrals created in CE processes between hot protons and slow neutrals

streaming to the pre-shock region (Hester et al. 1994; Smith et al. 1994). Recent theoretical studies (Blasi et al. 2012; Morlino et al. 2012) show that a BN precursor does not broaden the narrow component, but rather introduces a third intermediate component with the FWHM of around $\sim 150\,\mathrm{km\,s^{-1}}$ and depends on the shock speed. The reason is that only a small number of incoming neutrals interacts with ions in the BN precursor because its extent, which corresponds to the interaction length of the returning neutrals, is much smaller than the CE interaction length of the incoming neutrals. Therefore, Balmer lines can be used to study the microphysics of collisionless shocks and are currently the only means that give insight into the collisionless shocks.

We observed Tycho's SNR, which has already been well studied across all wavelength ranges (Reynolds & Ellison 1992; Bamba et al. 2005; Lee et al. 2004; Stroman & Pohl 2009; Katsuda et al. 2010; Acciari et al. 2011; Tian & Leahy 2011; Giordano et al. 2012), and evidence for particle acceleration in the shocks of Tycho's SNR, including acceleration up to the knee in the CR spectrum, was found (Warren et al. 2005; Cassam-Chenaï et al. 2008; Eriksen et al. 2011; Slane et al. 2014). In 1572, the star exploded as an SN Ia, leaving a remnant at an estimated heliocentric distance of 2.3 ± 0.5 kpc (Chevalier et al. 1980).[11] At that distance, the remnant's diameter of $8'$ corresponds to ≈ 5 pc. Density gradients in the medium around the remnant (Williams et al. 2013) modified the evolution of the shock, which in turn resulted in the asymmetric remnant. The lower shock velocity inferred in the northeastern (NE) part suggests that the shock interacted with a dense ambient medium, namely a diffuse cloud (Reynolds & Keohane 1999; Lee et al. 2004).

Previous optical studies of Tycho's SNR have shown indications for CRs (e.g., Ghavamian et al. 2000; Lee et al. 2007, 2010). However, these studies focused on the Hα-bright, but very complex, "knot g," where multiple or distorted shock fronts can contribute to the measured narrow-line (NL) broadening and thus partially mimic the effect that CR acceleration would have. Using the Fabry–Pérot instrument GHαFaS (Galaxy Hα Fabry–Pérot System) on the William Herschel Telescope (WHT), we observed a great portion of the shock front in the NE region of the remnant. The high spatial and spectral resolution, together with the large field of view (FOV) of the instrument, allow us to measure the narrow Hα-line width across individual parts of the shocks simultaneously, and thereby study the indicators of CR presence in a large variety of shock front conditions. In particular, the spatial resolution allows us to distinguish intrinsic line broadening from line broadening originating in geometric distortions and differential kinematics. Moreover, our observational setup provides us with unique insight into the existence of an intermediate component along the entire filament previously only reported for the bright "knot g" (Ghavamian et al. 2000; Lee et al. 2007). Apart from vastly enhancing the amount and quality of the spectroscopic data available for the NE filament, our study also improves the analysis: instead of fitting line models, we employ Bayesian inference to obtain full information and realistic uncertainties on the line parameters, as well as quantitative, reliable evidence for the presence of an intermediate line (IL) originating in a BN precursor and multiple

---

[11] Several optical studies based on modeling the observed Hα-line spectra give a distance of 2–3 kpc. Radio, X-ray, and gamma-ray observations prefer larger distances of 3–5 kpc (for a review, see Hayato et al. 2010).

shock fronts. The interpretation of the results is based on predictions of the state-of-the-art shock models that include the effects of BN and CR precursors on the observed Hα profiles (Morlino et al. 2012, 2013). With the GHαFaS spectral coverage of approximately $400\,\mathrm{km\,s^{-1}}$, we are not able to resolve the broad Hα component that was found to be about $\approx 2000\,\mathrm{km\,s^{-1}}$ in previous studies (Chevalier et al. 1980; Ghavamian et al. 2001).

2. Observations and Data Reduction

In order to resolve the narrow Hα lines along the rim of Tycho's SNR, we used the instrument GHαFaS mounted on the Nasmyth focus of the 4.2 m WHT (Hernandez et al. 2008), which operates at the Observatory del Roque de Los Muchachos in La Palma, Canary Islands. GHαFaS is a Fabry–Pérot interferometer–spectrometer with an FOV of $3\farcm4 \times 3\farcm4$. Its detector is an Image Photon Counting System (IPCS) for which the absence of readout noise is an advantage for observations of diffuse emission from extended objects. IPCS cameras are almost insensitive to CRs and thus do not require CR rejection. We used a high-resolution mode, acquiring data on 1024×1024 pixels^2 with $R \sim 21{,}000$ resolving power and a pixel scale of nearly $0\farcs2\,\mathrm{pixel^{-1}}$. The free spectral range (FSR) of the etalon was $8.56\,\mathrm{\mathring{A}}$ or $392\,\mathrm{km\,s^{-1}}$ centered at $6561\,\mathrm{\mathring{A}}$ and split into 48 channels that differ in their central wavelength, leading to a sampling velocity resolution of $8.16\,\mathrm{km\,s^{-1}}$. The instrument response function is well approximated by a Gaussian with FWHM of $19\,\mathrm{km\,s^{-1}}$ (Blasco-Herrera et al. 2010).

The observations were conducted on 2012 November 15–19 under FWHM $\simeq 1''$ seeing conditions. Successive exposures differ by one channel, and thus 48 successive exposures complete one cycle that covers the full spectral range. The total integration time was ≈ 9.6 hr, comprising 72 cycles and 3456 10 s exposures. Observations conducted over several cycles provide homogeneous airmass and atmospheric conditions for all channels. We reduced the data (see Figure 1) by first applying the phase correction to all exposures individually, where we follow the standard procedure for GHαFaS data described in Hernandez et al. (2008). The phase correction is a process of designating the photons' positions for the interference rings and assigning the corresponding wavelength to each position (x,y) on the image, $\lambda_i = \lambda_i(x, y)$. As a result, from each exposure $D_i(x, y)$, we build a data-subcube $D_i(x, y, \lambda)$ with 48 monochromatic images. In order to use the largest possible GHαFaS FOV, we did not use the optical derotator. Therefore, we have to align and derotate the observed data subcubes before co-adding them. To determine the exposures' relative pointing and orientation, we measure the centroid positions of bright point sources (stars) on a stack that combines each exposure with the 2×4 exposures that precede and succeed it. Thus, we are assured that at least three bright point sources are detected with enough flux for a <1 pixel centroid precision. This requirement means that we have to discard the first four and the last four D_i of each "run" of consecutive observations, but this concerns only 57 exposures (2%). Along with removing 13 cycles that suffer from reflected light or other defects that we noticed in visual inspection of the data, and also excluding seven cycles that lack reference sources for derotation, we retain 2439 exposures in 52 cycles. The final data cube $D(x, y, \lambda)$ results from summing all data
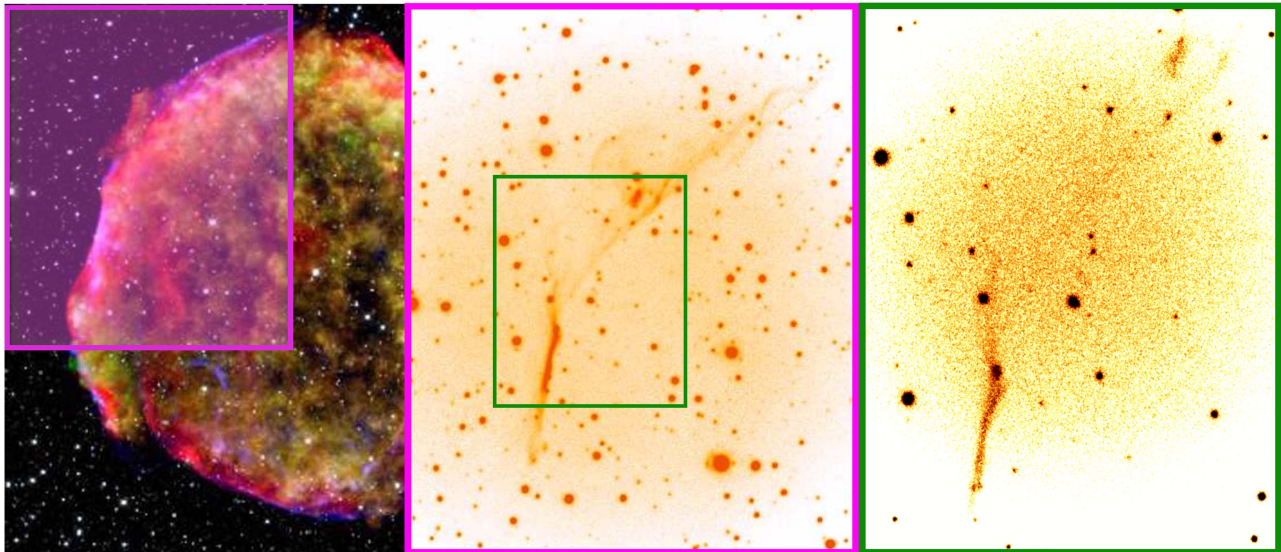
Figure 1. Left panel shows a composite image of the remnant ($\sim 8'$ in diameter) of Tycho Brahe's 1572 supernova, combining data from the *Chandra X-ray Observatory* (yellow, green, blue; NASA/CXC/SAO), *Spitzer Space Telescope* (red; NASA/JPL-Caltech), and the Calar Alto Observatory (white stars; Krause et al.). The transparent magenta box indicates the pointing of the ACAM (Auxiliary-port CAMera) on the Cassegrain focus of the WHT with an FOV of $4' \times 4'$. The center panel shows a zoom-in on the ACAM FOV. Using the same pointing as ACAM, we covered the same region with the GHαFaS Fabry–Pérot interferometer with an FOV of $3'.4 \times 3'.4$. The green box marks the region that is zoomed-in in the right panel to show our reduced and integrated GHαFaS Hα image.

subcubes and consists of 48 calibrated constant-wavelength slices.

Besides the data cube, that is, the stack of all aligned data subcubes, we produce a background cube and a flatfield cube. To this end, we model the background flux in individual exposures as well as the flatfield image (the position-dependent throughput of the optical system), and subsequently process the individual background and flatfield frames in exactly the same manner as the corresponding data frames. By constructing the co-added background and flatfield cubes and including them in our parametric models of the observed shock emission, as opposed to subtracting the background from the individual exposures and dividing them by the flatfield image before calibration and co-addition, we preserve the photon (Poisson) statistics in the data cube and simultaneously account for the variable effective exposure time in the stack of aligned subcubes. Details on the modeling of the flatfield and background frames are given in Appendix A.

3. Analysis

Our analysis is guided by the following goals and principles:

1. To determine the Hα NL width across a maximal area of the observed shock.
2. To achieve maximum spatial resolution, implying minimal binning (bin size) and signal per bin.
3. To still extract line parameters reliably and, in particular, characterize their uncertainties accurately.
4. To do so even when including up to two additional lines (10 model parameters).
5. To compare single-NL and multi-line models and to quantify their relative evidence.

In order to achieve these aims, we decide to perform parameter estimation and model comparison using Bayesian inference instead of traditional (maximum-likelihood, minimum-χ^2) fitting routines. We also account for the Poisson statistics of the data, as opposed to the often tacitly applied

Gaussian approximation. The details of our method can be found in Section 3.4, and Appendices E and F.

3.1. Motivation for a Multi-line Analysis

Narrow Hα lines in non-radiative shocks around SNRs are conventionally modeled by a single Gaussian, which theoretically has a width (FWHM) of $W_{NL} \sim 15 \, \text{km s}^{-1}$ in accordance with the pre-shock temperature of $\sim 10^4 \, \text{K}$ expected for the warm ISM. However, visual inspection of our data indicates that this basic model may not be descriptive; theoretical considerations also justify investigation of more complex models for the shock emission spectrum. Since the spatial elements (bins) cover a small but finite part of the shock, and moreover the observed filament is the projection of an extended shock section along the line of sight (LOS), a single-NL model is only suitable if one assumes that within each projected resolution element (bin) the following conditions are realized:

1. the pre-shock ambient medium is homogeneous with constant temperature,
2. the velocity distribution is uniform, without differential bulk velocity components along the LOS, and
3. there are no precursors (classical BDS).

If the three conditions above are all satisfied, the projection effects cannot modify the NL width, because both the width and the centroid are the same everywhere, corresponding to the upstream plasma temperature and bulk speed, respectively. Unresolved or projected inhomogeneous pre-shock temperature causes the superposition of Gaussians of different width. The presence of a CR precursor can alter the NL in two different ways: the cold neutrals in the ISM will be heated, resulting in the NL being broadened beyond the normal 10–20 km s^{-1} gas dispersion (Morlino et al. 2013), and they acquire a bulk speed up to a few percent of the shock speed. Therefore, inhomogeneous CR emission can be one reason for a non-Gaussian NL. CRs also transfer momentum to the pre-shock neutrals and potentially introduce a Doppler shift between the

gas in the precursor and the pre-shock gas not affected by the precursor if the shock is not moving strictly perpendicular to the LOS (Lee et al. 2007), in addition to shifts from any pre-existing differential bulk motion in the ambient medium. These projected or spatially unresolved shifts will also alter the line shape, and if the effect is pronounced enough, even lead to a split in the NL. One way to account for those distortions of the Gaussian line shape is by allowing generalized Gaussians with non-zero third- and fourth-order moments, for example using a Gauss−Hermite polynomial. However, we visually identify splits in the narrow line and choose to represent a non-Gaussian NL by a sum of two Gaussian components.

Apart from CRs there is another possible precursor—a BN precursor, which introduces a new line component: the so-called IL that can be described by a Gaussian[12] with FWHM in the range 100–350 km s^{-1} for a typical shock velocity in the range of [1500, 3500] km s^{-1} (Morlino et al. 2012). We therefore also investigate models that include this IL.

Previous observations focused on the brightest Hα knot, where the spectrum was measured by combining all pixels across the so-called "knot g" (Ghavamian et al. 2000; Lee et al. 2007) which, as *HST* imaging shows, has a complex structure (Lee et al. 2010). It was then used to estimate the NL and IL line widths of "knot g," with best fits of 44 ± 4 km s^{-1} and 150 km s^{-1} (Ghavamian et al. 2000), or 45.3 ± 9.0 km s^{-1} and 108 ± 4 km s^{-1} (Lee et al. 2007).

We set out to check if these results hold when "knot g" is spatially resolved and its parts analyzed individually, or if spatially averaging the spectrum introduced an artificially large W_{NL} and an IL that may have mimicked the effects of a CR or BN precursor. We also vastly extend the areal coverage of those earlier studies to include the lower-surface-brightness parts of the filament, and there, too, exploit the high spatial resolution provided by the GHαFaS instrument. In this way, we can investigate whether the signature of the existence of CR and BN precursors is also present in regions less complex than that of "knot g" and looking for possible differences in the physical properties of the shock.

We find that the contribution of the emission arising in the photoionization precursor (PIP), previously suggested and measured by Ghavamian et al. (2000) and Lee et al. (2007), is negligible (see Appendix B), so that we do not need to account for it in our filament flux models.

3.2. Definition of Models

For each location (bin), we consider several parametrized models (S) to characterize the shock Hα emission. Regardless of the parametrization (type of model) or specific parameter values (θ), we factor in the local flatfield spectrum (F) and add to it the observed background spectrum (B) before comparing the model with the data. This has the advantage of preserving the correct photon statistics, and contrasts with the common approach of subtracting the background from the data and dividing by the flatfield before modeling. Hence, for a given location (bin), the full model is represented by $M(\theta) = S(\theta) \times F + B$, while by model we mostly refer just to S, the intrinsic or "source" component. Note that in this expression, F and B are the result of binning the flatfield cube

and the background cube in the same way as the data, and therefore they only depend on wavelength (λ) for the bin that M describes. F and B are constructed separately from the data and inserted into the model without free parameters. More information on how we established F and B can be found in Section 2 and Appendix A.

Every S consists of an NL with a Gaussian profile as well as a constant component (c) that accounts for the sum of the continuum level and the broad Hα line (BL), which is several times wider than our spectral range. Depending on its type, S may further include one or two Gaussian components that represent an IL or an additional NL. Overall, in each bin we therefore have four different models S to compare with the data spectrum:

1. NL—constant plus single narrow line.
2. NLNL—constant plus two narrow lines.
3. NLIL—constant plus one narrow and one IL.
4. NLNLIL—constant plus two narrow lines and one IL.

In connection with calculating the models' relative evidence, we also consider a "no-line model" (0L; i.e., with the constant spectrum as the only component), which gives an auxiliary baseline for the relevance of at least the NL being present in the data. A possible choice of model parameters is the component fluxes ($f_c = f_{\text{constant}}$ and $f_i = f_{\{\text{NL, NL1, NL2, IL}\}}$), line centroids ($\mu_i$), and the lines' FWHM ($W_i$). The general form of this model is

$$S(\lambda) = f_c + (2\pi)^{-1/2}\sum_i f_i \exp\left(-(\lambda - \mu_i)^2/(2\sigma_i^2)\right)/\sigma_i , \quad (1)$$

where $\sigma_i = \sqrt{W_i^2 + W_{\text{instr}}^2}/\sqrt{4\ln 4}$ is the observed (instrumentally broadened) Gaussian dispersion and W_{instr} the FWHM of the instrumental response.

In practice, our models employ a transformed version of those parameters, which has the advantage of more direct interpretation, and a simpler functional form of the desired parameter priors (see the next paragraph). For example, instead of the two NL centroids, μ_{NL1} and μ_{NL2}, the NLNL and NLNLIL models use the NL centroid mean ($\langle\mu_{NL}\rangle$) and the separation between the NLs ($\Delta\mu_{NL}$). Similarly, the IL centroid is specified by its offset from the NL centroid or NLNL centroid mean, $\Delta\mu_{IL} = \mu_{IL} - \mu_{NL}$. We replace the component fluxes by the total flux and the components' flux fractions, and use the logarithm of the total flux and line widths as they are strictly positive quantities. A more detailed account of the definition of model parameters can be found in Appendix D.

The parameter priors, $P(\theta)$, are an integral part of the model definition: they encapsulate what we know (or assume) about the relative probabilities of the parameter values a priori, before considering the data. In particular, they can impose parameter boundaries by way of being zero outside of those boundaries. The line centroid parameters are effectively restricted to our spectral window. In our models, the FWHM parameters W_{NL} and W_{IL} are limited to [15, 100] km s^{-1} and [100, 350] km s^{-1}, respectively. The lower boundary of W_{NL} reflects the lower limit of the pre-shock temperature (≈5000 K), while the upper boundary is based on theoretical models that include the effects of the CR precursor (Morlino et al. 2013). The W_{IL} range is the theoretical expectation for shock velocities around 2000 km s^{-1} and a range of shock parameters (see Figure 10 in Morlino et al. 2012).

---

[12] Strictly speaking, the IL is not a perfect Gaussian, because it results from the population of neutrals undergoing CE in the BN precursor and they do not have enough time to thermalize to a single temperature (Morlino et al. 2012).
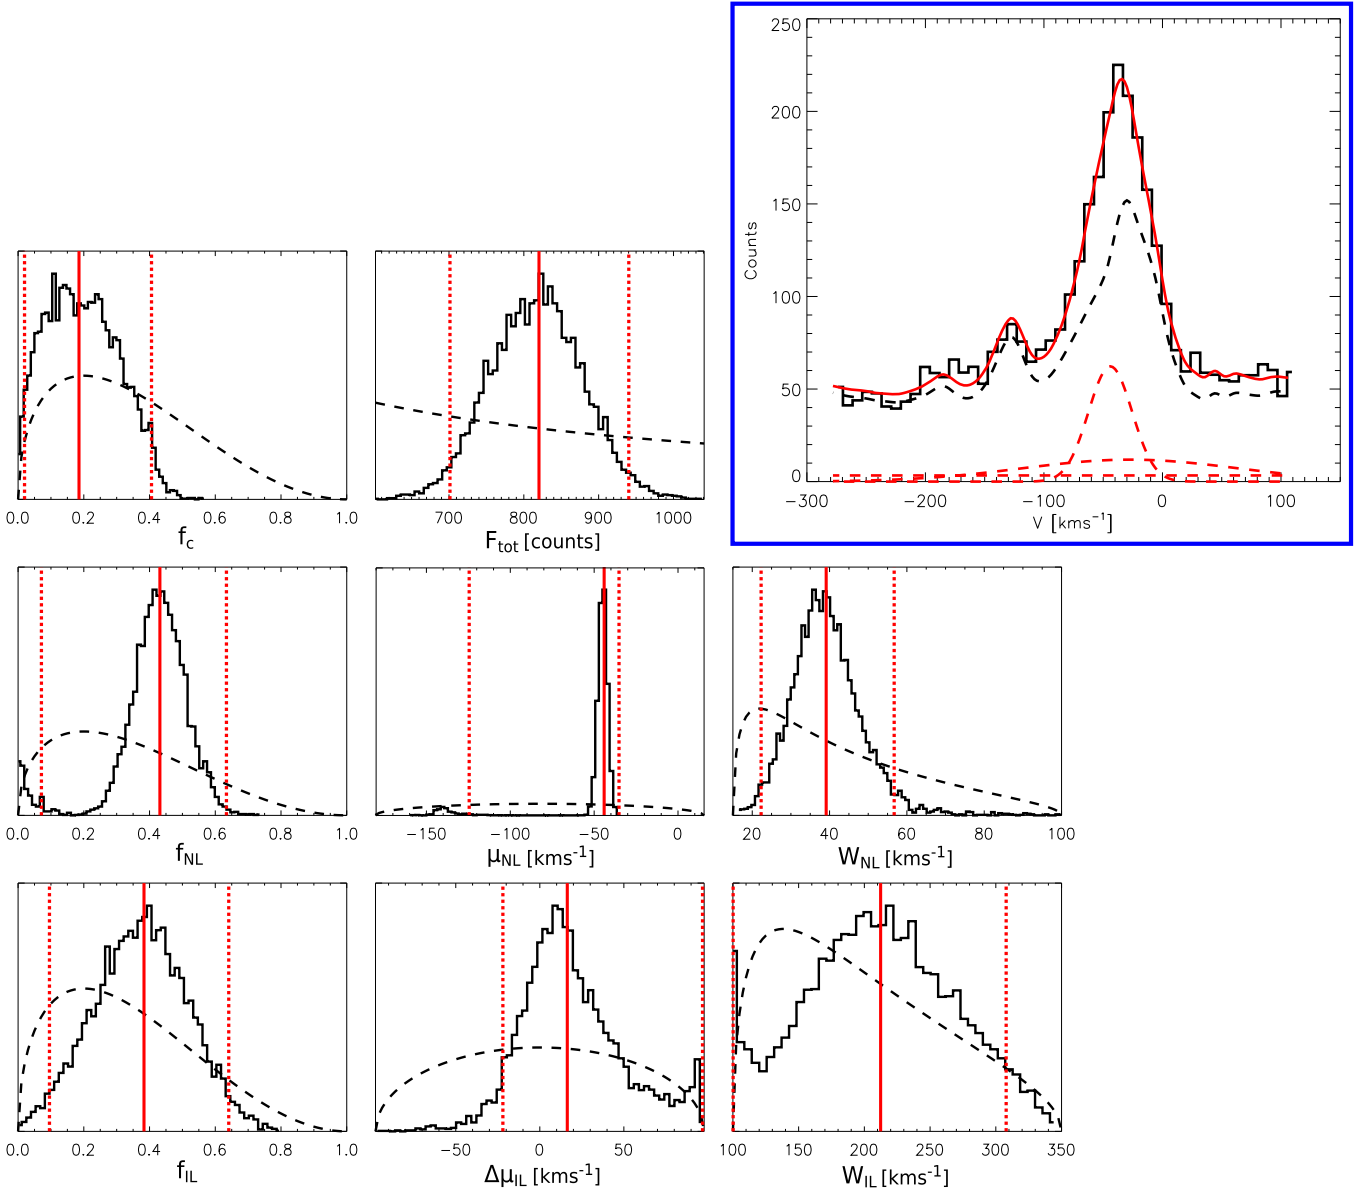
Figure 2. Parameter estimation via Bayesian inference for a bin in the NE filament of Tycho's SNR. The top-right blue panel shows the observed spectrum (solid black line), the background model (dashed black line), and the components of the intrinsic median NLIL model (dashed red lines). The median model is overplotted with the solid red line. The remaining eight panels are the 1D-marginalized posteriors over model parameters (solid black lines): total flux (F_{tot} in counts), flux fractions in the continuum (f_c) and lines (f_{NL} and f_{IL}), NL centroid (μ_{NL}), IL offset from the NL centroid ($\Delta\mu_{IL}$), and intrinsic line widths (W_{NL} and W_{IL}), with the latter three quantities given in km s^{-1}. The dashed black lines are the prior distributions, and the vertical red lines are the estimated parameters of the median model, i.e., the median values (solid red) and the boundaries enclosing the highest density 95% confidence intervals (dashed red).

Inside of the parameter boundaries, we desire to not strongly favor any particular parameter values a priori. However, on physical grounds, we prefer a smooth transition of the prior to zero for parameter values approaching the boundaries (see Figure 2). We therefore employ shifted and scaled Beta distributions with $\alpha = \beta = 1.5$ for the centroid parameters and logarithmic line widths, and a Dirichlet distribution with $\alpha = 1.5$ for the component flux fractions, which have the constraint of being summed up to unity. For comparison, the special case of "flat" Beta and Dirichlet distributions would be realized by setting $\alpha(=\beta) = 1.0$. Thus, our choice slightly favors the center of the allowed parameter range. The logarithm of the total flux has a flat unbounded prior. The model parameters and their priors are summarized in the Table 1, and detailed definitions are presented in Appendix D.

3.3. Binning

We analyze two shock filaments, one in the more eastern part of the NE rim which contains "knot g," and the other in the more northern part (Figure 3). We use the Weighted Voronoi Tessellation (Diehl & Statler 2006) with the adaptive bin size to spatially bin the pixels and obtain a signal-to-noise ratio (S/N) of $\gtrsim 10$ (Appendix C) in the wavelength-integrated signal that remains after subtraction of the background. This implies an average minimum S/N of 1.4 per spectral element. Due to the seeing of 1″, we require at least 5 pixels across, so that for a round bin this implies a minimum of 19 pixels. We exclude bins that would require >400 pixels for our target S/N, so that unaccounted-for residual background variations, which we estimate to be at most ~2% of the background level, do not significantly effect our measurements. Following these criteria,

<div align="center">

Table 1
Model Parameters and Their Prior PDFs

</div>

Parameters	Meaning	Prior
$\ln(F_{tot})$	Natural log-based total flux	flat prior
f_i, $i = [c, NL, NL1, NL2, IL]$	Flux fractions	Dirichlet prior: $\prod_i f_i^{\alpha-1}$, $\alpha = 1.5$; $\sum_i f_i = 1$, $f_i \in (0, 1)$
μ'_{NL}, $\langle \mu'_{NL} \rangle$	NL centroid, NL centroid mean	Beta prior: $x^{\alpha-1}(1-x)^{\beta-1}$, $\alpha = \beta = 1.5$; $x \in (0, 1)$
$\Delta\mu'_{NL}$	Separation between the two NLs	
$\Delta\mu'_{IL}$	IL centroid offset from NL centroid (mean)	
$w'_{NL(IL)}$	NL (IL) natural log-width	

Model	Model Parameters
NL	$\ln(F_{tot})$, f_{NL}, f_c, μ'_{NL}, w'_{NL}
NLNL	$\ln(F_{tot})$, f_{NL1}, f_{NL2}, f_c, $\langle \mu'_{NL} \rangle$, $\Delta\mu'_{NL}$, w'_{NL1}, w'_{NL2}
NLIL	$\ln(F_{tot})$, f_{NL}, f_{IL}, f_c, μ'_{NL}, $\Delta\mu'_{IL}$, w'_{NL}, w'_{IL}
NLNLIL	$\ln(F_{tot})$, f_{NL1}, f_{NL2}, f_{IL}, f_c, $\langle \mu'_{NL} \rangle$, $\Delta\mu'_{NL}$, $\Delta\mu'_{IL}$, w'_{NL1}, w'_{NL2}, w'_{IL}

Note. All parameters apart from $\ln(F_{tot})$ are defined in the (0,1) range (see Appendix D for notation).

we study 73 Voronoi bins in the eastern and 9 Voronoi bins in the northern filament.

3.4. Parameter Estimation and Model Comparison

For each Voronoi bin, we want to find which of the four models, M (see Section 3.2), and which vector of model parameters, θ, best explain its data (spectrum). Since the data are noisy, model comparison and parameter estimation are inherently probabilistic. For both tasks and the reasons discussed below, we use Bayesian inference.

The "standard" approach to estimate parameters is the maximum-likelihood (or minimum-χ^2) method. It relies on a well-defined likelihood maximum (mode) and the convergence of the optimizer deteriorates when high noise or multiple modes are present. Both conditions are met in our study: we desire high spatial resolution and therefore small bin size, implying low S/N. Our models are nonlinear and comprise up to 10 parameters, implying generally multiple modes. We wish to characterize all those modes, not just the "main" (global) maximum, toward reliable, high-confidence level parameter uncertainties ("errors"), instead of the minimal but common 68% confidence ("1 − σ") error. Often, and in our models with their non-trivial likelihood function, error propagation over a large parameter range is cumbersome or impossible. Finally, the maximum of the likelihood does not provide a quantitative, well-defined measure for the relative probability of the different models with different parametrizations (model comparison). These circumstances make maximum-likelihood or other fitting methods insufficient for our purposes.

Bayesian inference, by contrast, provides the full, multivariate parameter probability distribution function (PDF), the so-called *posterior* $P(\theta|D, M)$, as well as its integral (marginalization) over all parameters, $P(D|M)$—the *evidence*. Evidences of models are the relative model probabilities, without reference to any specific (e.g., best-fit) parameter value. Bayes' theorem states that the posterior as a function of θ is

$$P(\theta|D, M) = \frac{P(D|\theta, M)P(\theta|M)}{P(D|M)}. \quad (2)$$

It is proportional to the product of the likelihood, $L = P(D|\theta, M)$, and prior, $P(\theta|M)$. L reflects the model and is the probability of the data for given model parameters and

measurement errors. Our IPCS instrument counts photons, which are described by a Poisson distribution that therefore represents the measurement error, with expectation value and variance equal to the flux predicted by the spectral model. The prior is the parameter PDF that we know or assume before taking into account the data at hand. Apart from making these assumptions or knowledge explicit (fitting methods implicitly assume a flat prior in the chosen parameters), it has the advantage of naturally facilitating self-consistent parameter changes.

We represent the N-dimensional posterior PDF of the N model parameters as a sample, which we obtain using the Markov Chain Monte Carlo (MCMC) method. For details on the sampling algorithm, see Appendix E. The posterior can be summarized in many ways. One is the posterior maximum; we do not emphasize this as it is a relatively noisy estimator and its computation is not unique. Instead, we provide the median and the highest density (shortest) 95% confidence intervals of the one-dimensional marginalized distributions that result from integrating $P(\theta|D)$ over all but one parameter. We deem 95% to be the minimal confidence level worth quoting, and more reliable than the frequently employed 68% ("1 − σ") level, which carries a high probability (32%) of not including the optimal parameter values.

For the parameter estimation, we are interested only in the posterior's shape (relative parameter probabilities), which does not necessitate normalization by the evidence. However, the Bayes factor, i.e., the evidence ratio of two models, is a probabilistically well-defined, quantitative measure for comparing models. For a given M, the evidence is defined as $P(D|M) = \int_\theta P(D|\theta, M)P(\theta|M)d\theta$. It is the probability of observing the data when assuming that the model is "true" but the parameters are not specified. In practice, the evidence integral is often high dimensional and therefore computationally intensive. In order to approximate it numerically, we use the cross-validation (CV) likelihood (Bailer-Jones 2012), particularly the leave-one-out (LOO) CV likelihood: $P(D|M) = L_{LOO-CV}(D|M)$; see also Appendix F. To compute it, samples are drawn from the data partition posteriors instead of the prior as in "standard" evidence integrals, and it has the advantage that it depends on the prior only to second order. Each bin spectrum has 48 elements and the LOO-CV is applied to these elements by
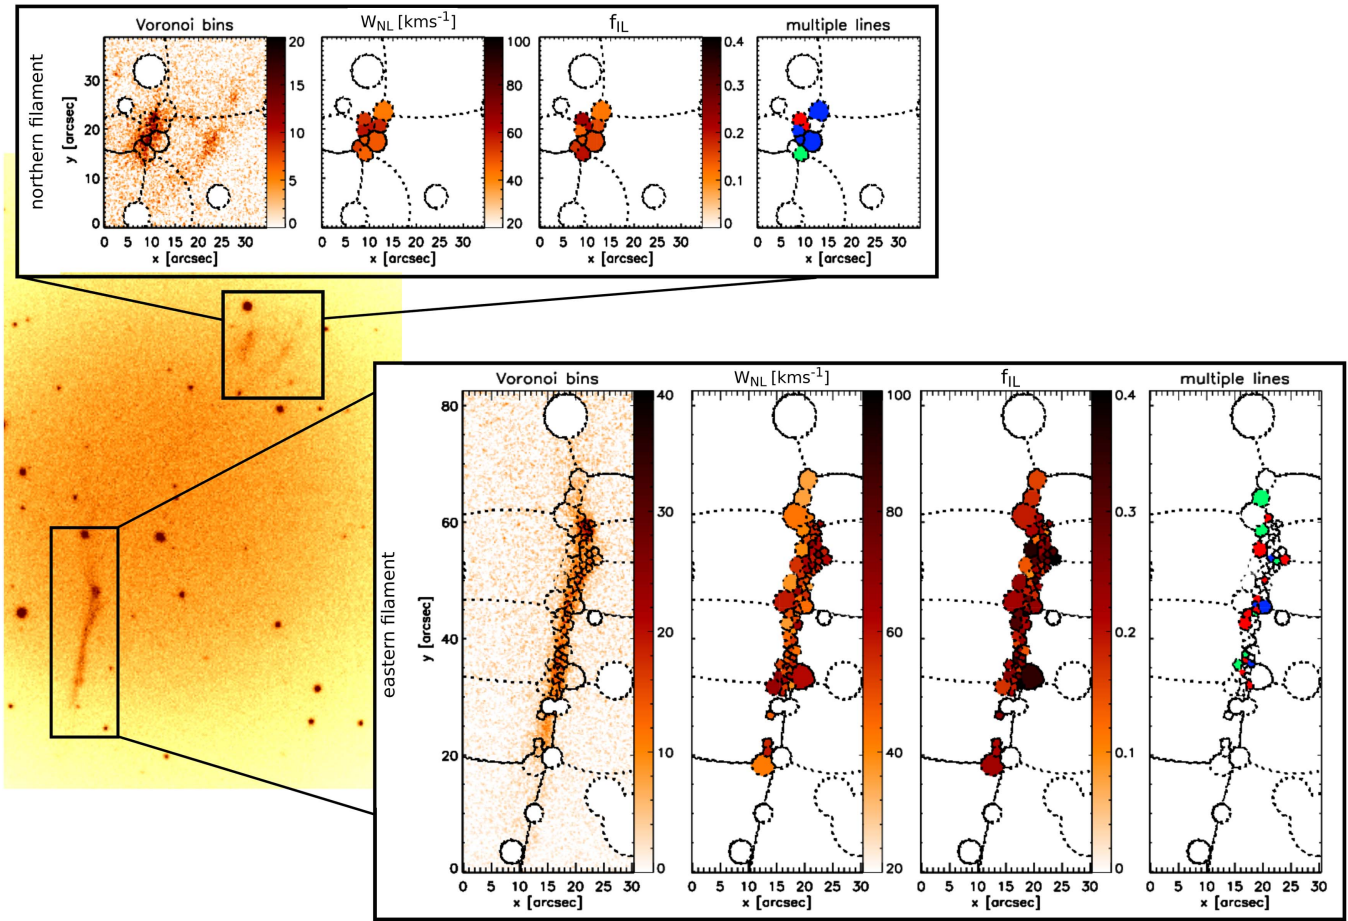
Figure 3. NE filament of Tycho's SNR. The two boxes on the GHαFaS Hα image show the northern and eastern shock filaments. The four panels in each box represent the bin contours overplotted on the background-subtracted cube, spatial variation of the median values of evidence-weighted NL width (in km s^{-1}) and IL flux fraction posteriors, and bins that show the necessity of IL (red), second NL (blue), and both the second NL and IL (green). In this last (rightmost) panel, white indicates either bins with too little flux, or those where no line is required in addition to the single NL to describe the data.

predicting each of them from the remaining 47 elements under the model M (marginalized over the parameters).

Since $P(D|M)$ can be a very large (or small) number, its absolute value is meaningless, and only relative values are needed for different models, we express it as the base-10 logarithm of the ratio with some reference model (Figure 4, Tables 2 and 3). As a matter of choice, we consider 0.5 dex log evidence differences as "significant" to clearly prefer one model over another. This choice is somewhat conservative; testing on simulated data reveals that we start to distinguish the correct models at 0.2 dex, while our numerical precision is around 0.05 dex.

With the 0.5 dex criterion per bin, we do not rule out the other respective models, but rather indicate that significant evidence exists that either an IL or double-NL (or both) is present in the data as a "population" in the filament overall, or conversely, that such an additional line emission is most likely not present if the evidence for the NL model relative to other models is larger than 0.5 dex.

The Bayes factors, and in turn the fraction of bins that show significant evidence for a double-NL (e.g. bin in Figure 5) or an additional IL (e.g. bin in Figure 2), depend on the line width used to distinguish an NL from an IL. Our choice of FWHM = 100 km s^{-1} as the lower IL width limit corresponds to shock speeds as low as 1500 km s^{-1} while previous results show the shock speed of Tycho's SNR to be at least 2000 km s^{-1} (Ghavamian et al. 2001). $W_{IL} < 100$ km s^{-1} requires unrealistic

full electron–proton temperature equilibrium (Caprioli 2015) or shock speeds lower than 1500 km s^{-1} and zero equilibration, i.e., $T_e/T_p = m_e/m_p$. The latter case was already debated by Morlino et al. (2012), and furthermore, $T_e/T_p \ll 0.01$ was never measured in any of the remnants (Ghavamian et al. 2013). On the other hand, NL widths larger than 60 km s^{-1} have not been observed before, and >100 km s^{-1} would require a CR acceleration efficiency $>40\%$ (Morlino et al. 2013), compared to the more realistic 10%–20% efficiency in SNR shocks (as was also found in Tycho's SNR by Morlino & Caprioli 2012). Therefore, our NL –IL separating line width limit robustly distinguishes the NL and IL that arise from different processes (cold neutral excitation and BN precursor).

4. Results

Following the analysis described above, we calculated posterior parameter distributions for every model and Voronoi bin, which we then summarized by the median of the 1D-marginalized posteriors as a central parameter estimator, and the boundaries of the shortest (highest-density) 95% confidence interval. We also compared models using the Bayes factors (evidence ratios) calculated using the CV likelihood method. Our focus is on the measured NL width, evidence for and the magnitude of a split in the NL, evidence for an IL and its

strength and width, the variation of LOS velocities across the filament, and possible correlations among the line parameters.

In Figure 2, we illustrate the application of our analysis to one of the bins from our data set. The observed spectrum (solid black line) is shown in the top blue panel, and the median NLIL model in solid red, while the model components are shown with dashed lines: the nonparametric background model (dashed black) and the parametrized "source" (SNR emission) model components—a constant component, one NL, and one IL—in dashed red. The background spectrum, which has been derived and fixed independently (see Appendix A) shows geocoronal and Galactic Hα emission, including the Galactic [N II] line at around $-130 \, \mathrm{km \, s^{-1}}$ LOS velocity. The 1D-marginalized posterior parameter distributions (solid black) and prior parameter distributions (dashed black) are shown in the remaining eight panels. The prior parameter distributions are overplotted for comparison with the posterior. The priors are chosen to not strongly prefer any parameter values over the allowed parameter range, but taper off smoothly toward zero at the range limits. Most posteriors are significantly different from the priors, and thus all of the parameters are well-constrained by the data. The parameters of the median model are denoted with red solid vertical lines. In this example, the median NL FWHM width is $W_{NL} \approx 40 \, \mathrm{km \, s^{-1}}$. The IL is $W_{IL} \approx 210 \, \mathrm{km \, s^{-1}}$ wide and comprises $\approx 40\%$ of the total flux. With the red dashed vertical lines, we marked the boundaries of the 95% confidence interval. Compared to all other model parameters, the posterior shape of W_{IL} is more sensitive to the choice of prior (Appendix F). Still, the median and 95% confidence intervals of W_{IL} of a flat and Beta prior agree within $\approx 10\%$. This relatively benign difference occurs because by coincidence the peak of the likelihood is close to the center of the prior range, and because the 95% confidence interval boundaries are similar to the minimum/maximum limits of the prior range.

To summarize the parameter estimation of one bin, we consider the evidence-weighted 1D-posterior (see Figure 6, middle row), to which any of the four models that feature the parameter of interest contribute their marginalized posterior in proportion to their evidence. We then give the median of the evidence-weighted posterior, as well as the boundaries of its 95% confidence interval.

To summarize the results of all bins, we take three different routes.

1. We consider the distribution of all 82 evidence-weighted posteriors' medians, as well as their 2×82 lower and upper 95% confidence interval boundaries (Figure 6, top panels). This yields information on the variability of line parameters across the filament.
2. Second, we combine (average) the evidence-weighted posteriors of all bins, providing a representation of the information that we typically find in one individual bin (Figure 6, middle row of panels, black curves). We also combine model-specific posteriors separately to illustrate the relative contribution of different models and how parameter estimates depend on the model choice (colored curves).
3. Finally, we evaluate the parameter constraints as imposed by all data (bins) combined. For that purpose, we sample from all bins' evidence-weighted posteriors, each time computing the median value over all bins. That is, we define one new parameter for each model parameter: the

cross-bin median. Its posterior is plotted in the bottom row of Figure 6, and again summarized by its median and 95% confidence interval boundaries. We emphasize that this measure is decidedly distinct from modeling the spectrum of all bins combined, because it is still based on models constrained by all bins individually and independently, and, in particular, allows for local shifts in the line centroids to avoid artificial line broadening.

Apart from estimating model parameters, we compare the probabilities of models (different numbers of emission lines) against each other. In each bin, we define a model as favored if it has the highest evidence *and* a >0.5 dex logarithmic evidence ratio (>3:1 probability) over the NL model (see Figure 3, right panel). This threshold is a matter of choice; it is twice the 0.2 dex by which the correct model is typically favored in our tests on simulated data and reflects our approximate notion of the minimum for a "significant" probability. If no model satisfies the 0.5 dex requirement, we consider the fiducial single-NL model as favored. In this scheme, the data in one bin may favor an IL (NLIL or NLNLIL model), a double-NL (NLNL or NLNLIL model), both (NLNLIL), or none of them (NL). Using this criterion for each bin separately, we ascertain the fraction of bins in which the evidence indicates an IL, as well as the fraction of double-NL occurrence.

4.1. Narrow-line Width

The example in Figure 2 is by no means the only one where the NL width, W_{NL}, is much larger than the maximally allowed thermal NL broadening ($20 \, \mathrm{km \, s^{-1}}$). On the contrary, our central estimator, the median of the evidence-weighted W_{NL} posterior, is never lower than $35 \, \mathrm{km \, s^{-1}}$, in any bin. The spatial distribution of the estimated W_{NL} is shown in panel 2 of Figure 3. We do not recognize a strong spatial pattern of W_{NL}; most of the bin-to-bin variations appear to be randomly distributed. However, in the northern part of the eastern filament, there appears to be a trend of lower (higher) W_{NL} on the pre-shock (post-shock) side, and a generally higher W_{NL} in the southern part. The spatial variation of W_{NL} (Figure 3) probably indicates variations in the amount of neutrals in the ambient medium: more neutrals imply a more efficient ion−neutral damping of magnetic waves excited by CRs, thus resulting in a pre-shock gas heated to larger temperatures. The histogram of the median and 95% confidence interval boundaries can be found in the top-left panel of Figure 6. As measured by the cross-bin median, the global W_{NL} is $(54.8 \pm 1.8) \, \mathrm{km \, s^{-1}}$ with 95% confidence. Suprathermal NL widths are required even when only models with double-NL or an additional IL are considered; their average W_{NL} is only a few $\mathrm{km \, s^{-1}}$ lower (see the next two subsections).

4.2. Evidence for a Split in the NL

In 18% of the Voronoi bins (15 of 82), we find significant evidence for an NLNL(IL) model, i.e., for a split in the NL. For one of these bins (bin 5 in Table 3), we show posteriors in Figure 5. In this example, NL centroids and widths are well-determined, while the flux of the second (right, high-velocity) NL closely follows the prior distribution. The two NLs are separated by $\approx 40 \, \mathrm{km \, s^{-1}}$ and have median widths of ≈ 52 and $70 \, \mathrm{km \, s^{-1}}$. That is, despite using two Gaussians, both are much wider than the thermal $20 \, \mathrm{km \, s^{-1}}$, indicating that the
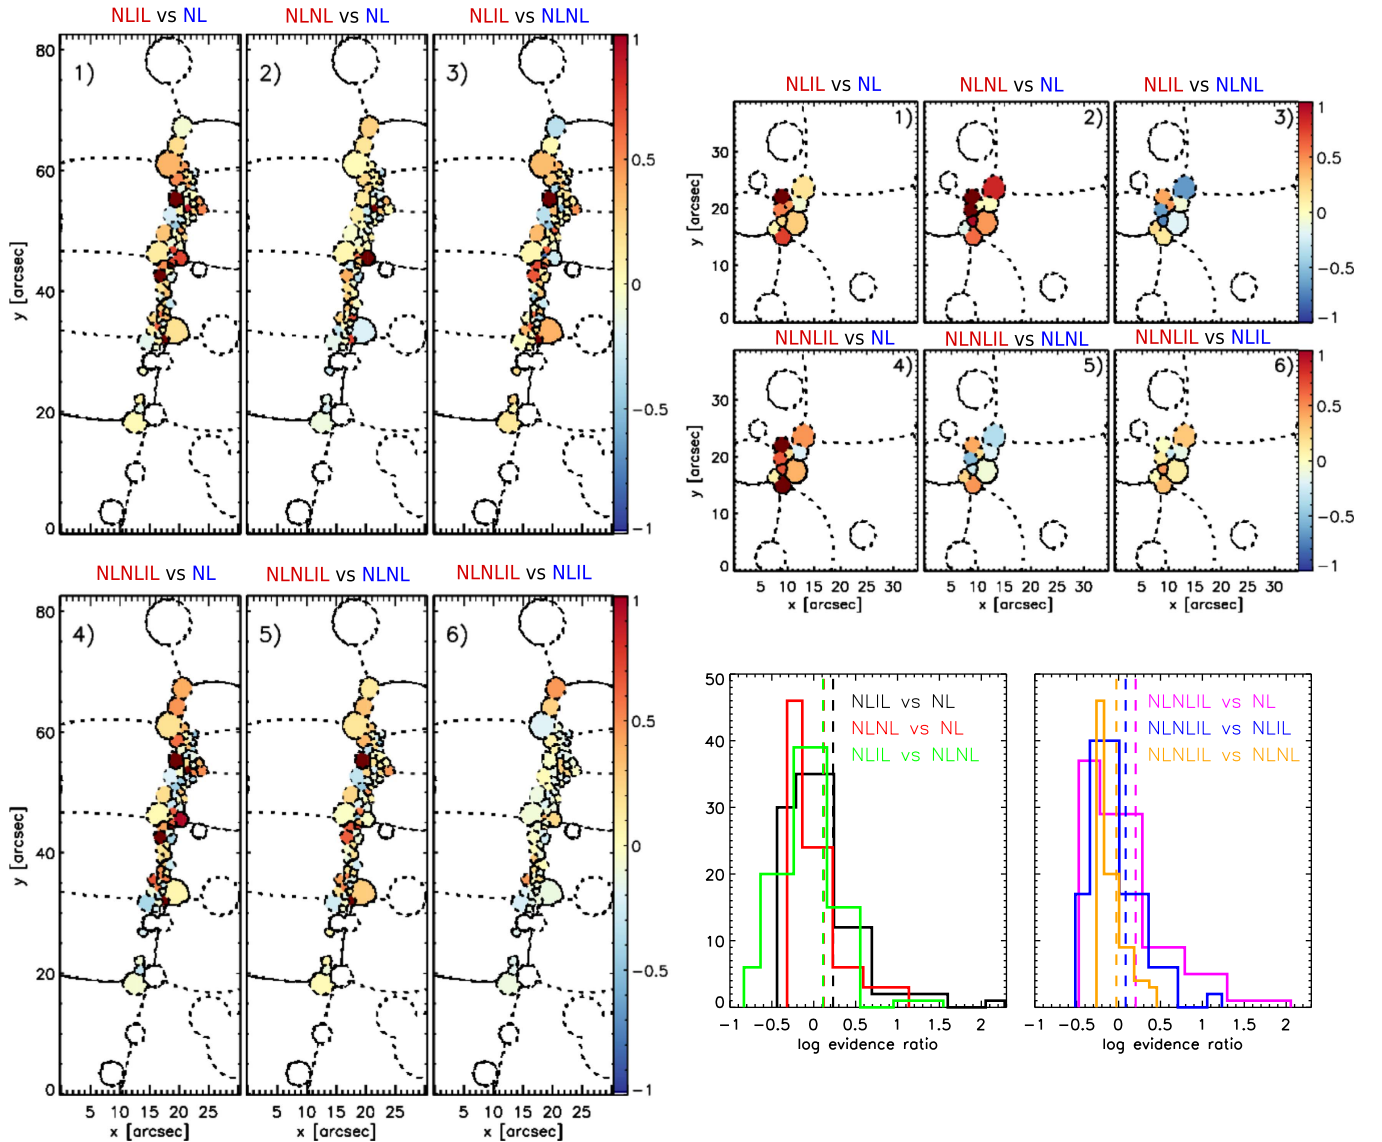
Figure 4. Spatial variation and histogram representation of logarithmic (base-10) evidence ratios across the NE rim in Tycho's remnant. The logarithmic evidence ratios are shown in the range ±1 dex in the spatial maps, while the histograms show the entire range of log evidence ratio values. Bins with log evidence ratio of around +1 dex ratio appear as dark red, while bins with log evidence ratios of around −1 dex appear as dark blue. For example, the dark red bins in panel (2) of the eastern filament favor the NLNL model when compared to the NL model, with the log evidence ratio between the two models of >1 dex. Model comparison is also quantitatively presented in the histogram visualization of the evidence ratios where we also plot the mean values of the distributions (vertical dashed lines).

broadening is mostly not an artifact of unresolved differential LOS motion. Not only is the log evidence ratio of the NLNL to NL model larger than 1 dex in favor of the NLNL model, but it is also more than three times ($\gtrsim 0.5$ dex) more probable than the NLIL or NLNLIL model. This implies highly significant evidence for the presence of a second (split) NL, and in this particular case, also significant evidence against an IL.

The example in Figure 5 is not untypical: *even in locations where double-NL models are favored*, the average measured $W_{\rm NL}$ is 49 km s^{-1}, again much larger than the upper limit of the intrinsic thermal NL width. The median parameters and 95% confidence intervals for the favored model of each bin are listed in Tables 4 and 5. The cross-bin median NL−centroid separation constrained from all of the bins in the NE rim is $\Delta\mu_{\rm NL} = 38.5 \pm 5.1$ km s^{-1} (95% confidence).

In the case of a perfectly spherical shock and homogeneous ambient ISM, the parts of the filament closest to the upstream are seen edge-on, and we would expect to see a single NL,

while the parts of the filament closer to the downstream are inclined to the LOS and should exhibit a split in the NL. Evidence for a double NL in the eastern filament is found in a few inner bins (blue bins in Figure 3), and also in some outer (green) bins. Given the small number of double-NL occurrence and its scattered locations, any determination of the shell geometry would be vague. However, the detection of double NLs with $W_{\rm NL} \gg 20$ km s^{-1} clearly points toward heating and momentum transfer in the CR precursor.

4.3. Intermediate-line Evidence and Parameters

As quantified by the Bayes factors, we find that 34% of the Voronoi bins (28 bins out of 82) are significantly better explained when a line is added to the fiducial NL-only model. In 74% of those, either the NLIL or NLNLIL model is also preferred over an NLNL model, which means 24% of the bins overall. This is illustrated in Figure 4, where the logarithmic

(base-10) ratios of all models versus one another are shown in separate panels, and each panel shows the spatial variation of the corresponding ratio. These values are visualized by the histograms and also tabulated in the Tables 2 and 3. Apart from just a few bins, the NLNLIL model is not clearly favored over the NLIL model (panel 6). The situation changes when it comes to the comparison between NLNLIL and NLNL (panel 5), further confirming the necessity of an IL component. One example (bin 17 in Table 2) is illustrated in Figure 2, where the logarithmic evidence ratio of the NLIL model relative to the NL and NLNL models is larger than 1 dex. This means that an IL in addition to a single NL explains our data more than 10 times better than a simple NL model or NLNL model, irrespective of any particular parameter values. Both NLIL and NLNLIL models are equally likely for this bin, which implies no need for the second NL component. For this particular bin, we present posteriors for all other models in Appendix D.

Prominent ILs are seen across the entire filament (panel 3 in Figure 3), with an IL flux fraction of up to 42%. If we only consider bins that favor the NLIL or NLNLIL model, we get an IL flux fraction of 28% on average, and the intermediate-to-narrow flux fraction $f_{\rm IL}/f_{\rm NL}$ is estimated at 0.61 with a 95% confidence interval of (0.01–1.87).

In contrast to the IL flux, $W_{\rm IL}$ is not well-constrained in individual bins, as indicated by the similarity of the prior and average bin-specific posterior (Figure 6, middle panels). However, several bins with high S/N and/or strong IL emission (as the bin in Figure 2) have good constraints on the IL width. Moreover, the combination of the data in all bins provides even better information: the cross-bin median is $W_{\rm IL} \in [166.77, 194.15]$ km s^{-1} (95% confidence; see the bottom of Figure 6). The global 95% confidence interval of $f_{\rm IL}/f_{\rm NL}$ is [0.34, 0.47]; however, we caution that it varies considerably between bins, with median $\log(f_{\rm IL}/f_{\rm NL}) \in [-0.79, 0.05]$.

4.4. Line-of-sight Velocity

We present the observed LOS velocity, i.e., the NL centroid $\mu_{\rm NL}$, in Figure 7. The map of the bin-specific $\mu_{\rm NL}$ is shown in the left panels. In the top-right panel, we show the corresponding histogram (green), as well as the distribution of the 95% confidence interval boundaries. Across the entire NE rim, we find median LOS velocities in the range $[-67.9, -28.6]$ km s^{-1}. Notably, the northern filament moves with respect to the eastern filament, $\mu_{\rm NL} = -32.9$ km s^{-1} and -40.5 km s^{-1} on average. The bin-to-bin variations in $\mu_{\rm NL}$ are most noticeable in the eastern filament, indicating inhomogeneities in the shock and, in turn, in the ISM density. The bulk (median) LOS velocity of the filament as a whole is $\mu_{\rm NL} = -38.3 \pm 1.5$) km s^{-1} at the 95% confidence level (see the middle-right panel). Converting to the local standard of rest (LSR), we obtain $V_{\rm LSR} \approx -34$ km s^{-1}. We also check for the correlation between the LOS velocities and surface brightness (bottom-right panel), but find them to be uncorrelated.

4.5. "Knot g"

This location in the eastern filament has been the target of previous observational Hα studies, but it was not spatially resolved with spectroscopic data. We spatially resolve this "knot g," and the entire NE filament, with our spectroscopic data

for the first time. More precisely, the centroids of six of our bins drop within the Lee et al. (2007) slit. Averaging over these bins, we find $W_{\rm NL} = 49 \pm 15$ km s^{-1} (68% confidence), which is in agreement with Lee et al.'s reported values of 44 ± 4 km s^{-1}. The mean LOS velocity is -39.7 km s^{-1} enclosed by the 68% confidence interval of $[-50.2, -34.7]$ km s^{-1}. The mean LSR velocity in "knot g" is found to be ≈ -36 km s^{-1}, in agreement with the previous estimate by Lee et al., but with a blueshifted offset of nearly 6 km s^{-1}. Typical measured parameters and their 68% confidence intervals in "knot g" are $W_{\rm IL} = 188$ [105, 229] km s^{-1}, $f_{\rm IL}/f_{\rm NL} = 0.35$ [0.05, 0.54], and $\Delta\mu_{\rm NL} = 32.4$ [8.1, 46.7] km s^{-1}.

4.6. Correlation in Parameters

We investigate the correlation between various line parameters and surface brightness, as well as correlation among model parameters, using Spearman's rank correlation coefficient (ρ). It is sensitive to any monotonic relationship, even if it is not linear as for Pearson's correlation coefficient, and does not require normally distributed data. The result indicates a strong correlation (or anticorrelation) for values close to ± 1.

In the top row in Figure 8, we plotted the median of evidence-weighted $W_{\rm NL}$, $f_{\rm IL}/f_{\rm NL}$, $W_{\rm IL}$, and $\Delta\mu_{\rm NL}$ posteriors against the surface brightness (SB; total flux of the intrinsic model divided by the bin area). We separately test the correlation for bins for which a single-NL model (NL or NLIL, marked with squares) has the highest evidence, and those bins for which a double-NL model (NLNL or NLNLIL, marked with circles) has the highest evidence of all models. We do not find any strong correlation: parameters of single-NL models are uncorrelated with surface brightness, while double-NL models show weak correlations with SB. The highest correlation is with $\Delta\mu_{\rm NL}$ ($\rho = 0.57$), meaning that with increasing SB, the two NL centroids tend to be more separated.

In the bottom-row panels of the same figure, we estimated the correlation among the parameters. Although we see a hint of a positive correlation between the NL width and NL separation ($\rho = 0.47$ and $\rho = 0.49$; see panel 8), we do not find a clear correlation between any of the parameters. We get very similar results when we apply Pearson's correlation: its coefficient is either very close to Spearman's rank coefficient or closer to zero. In addition, the typical uncertainty of all of these points (plotted in the top-right corner of each panel and given as the average shortest 68% confidence interval) propagates into the correlation rank coefficient uncertainty.

5. Discussion

Previous Hα observations of Tycho's "knot g" (Lee et al. 2007) were modeled by Wagner et al. (2009) by computing a series of time-dependent numerical simulations of CR-modified shocks. Assuming a distance of 2.1 kpc to the remnant, they found the CR diffusion coefficient of $\kappa = 2 \times 10^{24}$ cm^2 s^{-1} and the lower limit of the injection parameter $\xi_{\rm inj} = 4.2 \times 10^{-3}$ to be in good agreement with the observations, suggesting that CR acceleration in the shock is efficient. Diffuse emission $1''$ ($\sim 10^{16}$ cm) ahead of the eastern filament was also detected by Lee et al. (2010) and interpreted as emission from the CR precursor with $T \sim 80{,}000$–$100{,}000$ K.

In what follows, we will summarize the main results of our paper and look for theoretical explanations for our findings.

Table 2
Model Comparison for the 73 Spatial Bins in the Eastern Shock Filament of Tycho's SNR

Bin (1)	x ["] (2)	y ["] (3)	Pix (4)	S/N (5)	0L (6)	NL (7)	NLNL (8)	NLIL (9)	NLNLIL (10)
1	22.0	59.2	20	14.40	27.073	0.108	0	0.077	0.234
2	21.4	58.4	26	14.21	31.122	0.307	0.273	0	0.098
3	22.4	58.2	21	14.96	40.522	0.245	0	0.236	0.447
4	22.8	59.7	26	12.52	13.405	0	0.318	0.220	0.476
5	21.6	57.4	22	12.33	27.723	0.435	0.343	0	0.110
6	21.0	60.7	46	13.08	31.293	0.459	0.391	0	0.091
7	19.6	64.2	188	12.04	39.630	0.483	0.325	0.270	0
8	21.1	59.5	26	15.19	30.855	0.176	0.113	0	0.160
9	18.3	61.1	378	12.84	37.771	0.376	0.348	0	0.185
10	20.5	67.3	234	11.43	33.046	0.410	0.146	0.473	0
11	20.4	57.1	51	12.56	25.860	0.236	0	0.370	0.295
12	21.3	56.3	32	11.07	25.137	0	0.217	0.074	0.266
13	22.4	57.2	22	11.97	17.165	0	0.161	0.246	0.356
14	19.7	58.5	104	13.81	42.871	0.594	0.366	0.088	0
15	22.5	55.9	48	10.68	14.121	0	0.182	0	0.141
16	21.3	55.0	40	10.47	19.664	0.188	0.178	0.272	0
17	19.5	55.3	137	13.05	30.044	1.456	1.307	0	0.027
18	22.4	54.5	27	9.69	15.533	0	0.203	0.013	0.121
19	21.4	53.8	37	10.44	12.876	1.387	0	0.475	0.203
20	20.2	53.4	51	11.33	17.289	0.094	0.360	0	0.126
21	23.4	55.0	40	8.49	9.724	0.436	0.084	0	0.092
22	21.3	52.6	31	10.51	20.602	0	0.109	0.050	0.207
23	22.5	53.3	30	10.98	19.092	0.518	0.212	0.232	0
24	18.5	52.7	138	11.15	23.991	0.107	0	0.333	0.303
25	22.3	52.1	27	9.17	4.052	0	0.249	0.132	0.333
26	20.4	51.8	38	11.57	10.662	0	0.267	0.281	0.434
27	21.6	51.4	23	9.86	18.144	0.016	0	0.081	0.119
28	19.4	51.0	48	11.10	21.513	0.018	0	0.456	0.444
29	20.8	50.8	19	10.63	20.281	0.166	0.223	0	0.146
30	23.9	53.6	67	9.73	9.985	0.524	0.488	0	0.024
31	19.3	49.6	36	10.09	23.207	0	0.131	0.041	0.254
32	20.3	50.0	28	11.16	18.678	0.459	0.628	0	0.152
33	20.9	49.0	32	10.58	16.697	0.218	0.105	0.010	0
34	19.9	48.7	26	11.08	21.270	0.081	0.098	0	0.099
35	18.7	48.1	47	12.01	23.668	0	0.085	0.084	0.265
36	20.2	47.5	48	11.38	21.533	0	0.008	0.052	0.165
37	17.5	49.6	165	11.17	19.644	0.316	0.352	0	0.112
38	19.0	46.9	37	11.92	25.193	0.730	0.622	0	0.070
39	18.6	46.0	21	12.10	23.762	0.603	0	0.834	0.452
40	20.3	45.5	115	12.72	21.540	1.005	0	0.258	0.033
41	17.3	44.3	59	11.28	23.175	0.503	0.713	0	0.072
42	18.9	44.9	35	11.55	23.385	0.731	0.432	0	0.084
43	18.5	44.2	16	9.10	11.179	1.116	0.790	0.065	0
44	16.4	46.4	279	12.59	27.058	0.125	0.066	0	0.082
45	18.4	43.4	31	11.30	21.451	0.033	0.072	0	0.173
46	16.8	42.6	98	12.74	22.116	1.082	0.673	0	0.020
47	18.7	42.3	52	11.63	20.054	0	0.303	0.172	0.404
48	16.9	40.6	58	12.11	26.149	0.207	0	0.051	0.041
49	17.9	41.3	29	11.62	18.944	0.195	0.282	0	0.118
50	18.0	39.7	42	12.24	20.337	0.324	0.289	0	0.112
51	15.5	35.6	58	11.36	28.214	0.685	0.432	0.459	0
52	17.4	38.1	21	10.47	17.933	0.324	0	0.113	0.045
53	16.9	39.0	38	12.27	22.999	0.265	0	0.404	0.107
54	18.5	38.0	58	10.56	12.943	0	0.022	0.292	0.268
55	16.9	37.3	29	12.04	18.049	0.466	0.152	0.195	0
56	17.9	36.7	29	11.74	18.795	0.198	0.316	0	0.151
57	16.8	36.2	23	12.53	16.550	0.589	0.619	0	0.083
58	19.5	33.6	344	11.89	16.533	0.187	0.395	0	0.117
59	17.8	35.7	34	12.40	30.850	0.695	0	0.361	0.150
60	16.6	35.2	19	11.81	16.934	0.091	0	0.096	0.119
61	17.2	33.5	25	11.15	17.640	0.033	0.203	0	0.118
62	17.3	34.7	28	10.86	19.534	0	0.038	0.172	0.264
63	16.4	34.2	25	11.63	26.490	0.648	0.649	0	0.110

Table 2
(Continued)

Bin (1)	x ["] (2)	y ["] (3)	Pix (4)	S/N (5)	0L (6)	NL (7)	NLNL (8)	NLIL (9)	NLNLIL (10)
64	17.6	32.1	42	12.86	26.277	2.415	1.670	0	0.098
65	16.5	32.8	29	13.09	22.555	0	0.123	0.128	0.278
66	16.5	31.7	34	12.12	27.519	0.086	0.126	0	0.154
67	16.0	30.7	32	12.02	17.226	0	0.131	0.088	0.203
68	14.4	31.7	155	12.61	32.667	0	0.148	0.165	0.393
69	13.2	22.1	59	9.33	18.851	0.095	0.241	0	0.171
70	12.3	18.4	285	10.04	15.999	0.047	0.152	0	0.112
71	13.8	26.8	47	8.78	8.551	0	0.239	0.030	0.194
72	15.3	33.6	61	11.76	19.728	0.090	0.377	0	0.185
73	13.3	20.6	59	9.43	11.914	0	0.208	0.138	0.300

Note. Columns 1–5: number of the (Voronoi) bin, x and y coordinates of the bin centroid, number of combined pixels, and signal-to-noise ratio. Columns 6–10: relative log CV likelihoods of the favored model (denoted with 0) to other models.

Table 3
Same as Table 2 just for the 9 bins in the Northern Filament

Bin (1)	x ["] (2)	y ["] (3)	Pix (4)	S/N (5)	0L (6)	NL (7)	NLNL (8)	NLIL (9)	NLNLIL (10)
1	13.0	23.7	351	14.30	33.932	0.846	0	0.667	0.357
2	10.4	20.6	62	15.31	43.600	0.483	0.479	0	0.264
3	8.9	22.1	169	15.36	53.524	1.645	0.423	0	0.009
4	12.3	20.8	169	14.77	41.332	0.003	0	0.068	0.205
5	8.6	19.7	129	15.50	49.913	1.194	0	0.644	0.517
6	11.7	17.3	355	16.72	34.586	0.470	0	0.186	0.077
7	9.0	17.8	75	16.31	42.528	0.959	0	0.741	0.269
8	9.1	15.0	190	14.65	40.034	1.090	0.479	0.329	0
9	7.4	16.5	104	14.73	37.961	0.086	0.246	0.025	0

Presence of a CR precursor. The NL width is much broader than $20\,\mathrm{km\,s^{-1}}$ ($\approx 55\,\mathrm{km\,s^{-1}}$ on average) in the entire NE rim regardless of whether the shock emission is described with a single or double NL. In other words, even when differential velocities (double NL) are present and accounted for, the NL is still significantly broadened. This clearly points toward gas heating in a CR precursor (Morlino et al. 2013). Furthermore, momentum transfer in a CR precursor might result in a split in the NL if we have two inclined shocks projected on our LOS. This is something that we observe—more precisely, we find significant Bayesian evidence for it in 18% of the data. Apart from the separation between the two NLs being $38\,\mathrm{km\,s^{-1}}$ on average, at the same time we find that their intrinsic widths are around $49\,\mathrm{km\,s^{-1}}$. If we assume that we have the contribution of two shocks inclined with the same angle to the LOS and that neutrals in the CR precursor acquired 10% of the previously estimated shock velocity of $2500\,\mathrm{km\,s^{-1}}$ (Ghavamian et al. 2001), we find that the shock normal inclination of $85°$–$86°$ explains the centroid separation of $38\,\mathrm{km\,s^{-1}}$. Obviously, for the same shock speed and smaller acquired bulk velocity, we would need the shocks slightly more inclined, i.e., with angles $<85°$. However, since only few bins show evidence for a split in the NL, we do not have enough information to construct the geometry of the shock.

Presence of a BN precursor. The main signature of the BN precursor is an IL (Morlino et al. 2012); 24% of the bins demand an additional line being specifically IL. We find a median value for the intrinsic IL width of $180\,\mathrm{km\,s^{-1}}$, comprising on average a 41% intermediate-to-narrow flux

ratio. Moreover, the observed high pre-shock neutral fraction of 0.9 (Ghavamian et al. 2001) in combination with the shock velocity of $\sim 2500\,\mathrm{km\,s^{-1}}$ supports the BN presence since a large neutral fraction and the specified shock velocity contribute to efficient CE and the larger number of created BNs play an essential role in the formation of and heating in the BN precursor. BNs are ionized inside the BN precursor almost immediately when they cross the shock. The newly created protons move with a bulk speed larger than the Alfvén speed, hence they can trigger the streaming instability (and possibly other kind of instabilities), resulting in an increase of the ion temperature in the BN precursor. In turn, the CE between the pre-shock neutrals and warm ions creates warm neutrals that produce the IL.

For illustrative purposes, in Figure 9 we report the temperature profile of the pre-shock gas for specific values of shock parameters as calculated in Morlino et al. (2013) where the CR and BN precursors can be clearly distinguished. The BN precursor acts on scales of 10^{16}–$10^{17}\,\mathrm{cm}$ in the immediate pre-shock region and can heat the gas to a temperature of $\sim 10^6\,\mathrm{K}$. The CR precursor length is much larger, depends on the maximum energy of the accelerated particles, and extends over several $10^{17}\,\mathrm{cm}$. The expected gas temperature in the CR precursor is $\sim 10^5\,\mathrm{K}$.

Collisionless shock model prediction in partially ionized medium. Following the results of Morlino et al. (2013) and their Figure 9, 40–$50\,\mathrm{km\,s^{-1}}$ NL widths require efficient turbulent heating η_{TH} in the CR precursor, but also accelerated particles with the maximum momentum of $p_{\max} = 40\,\mathrm{TeV/c}$
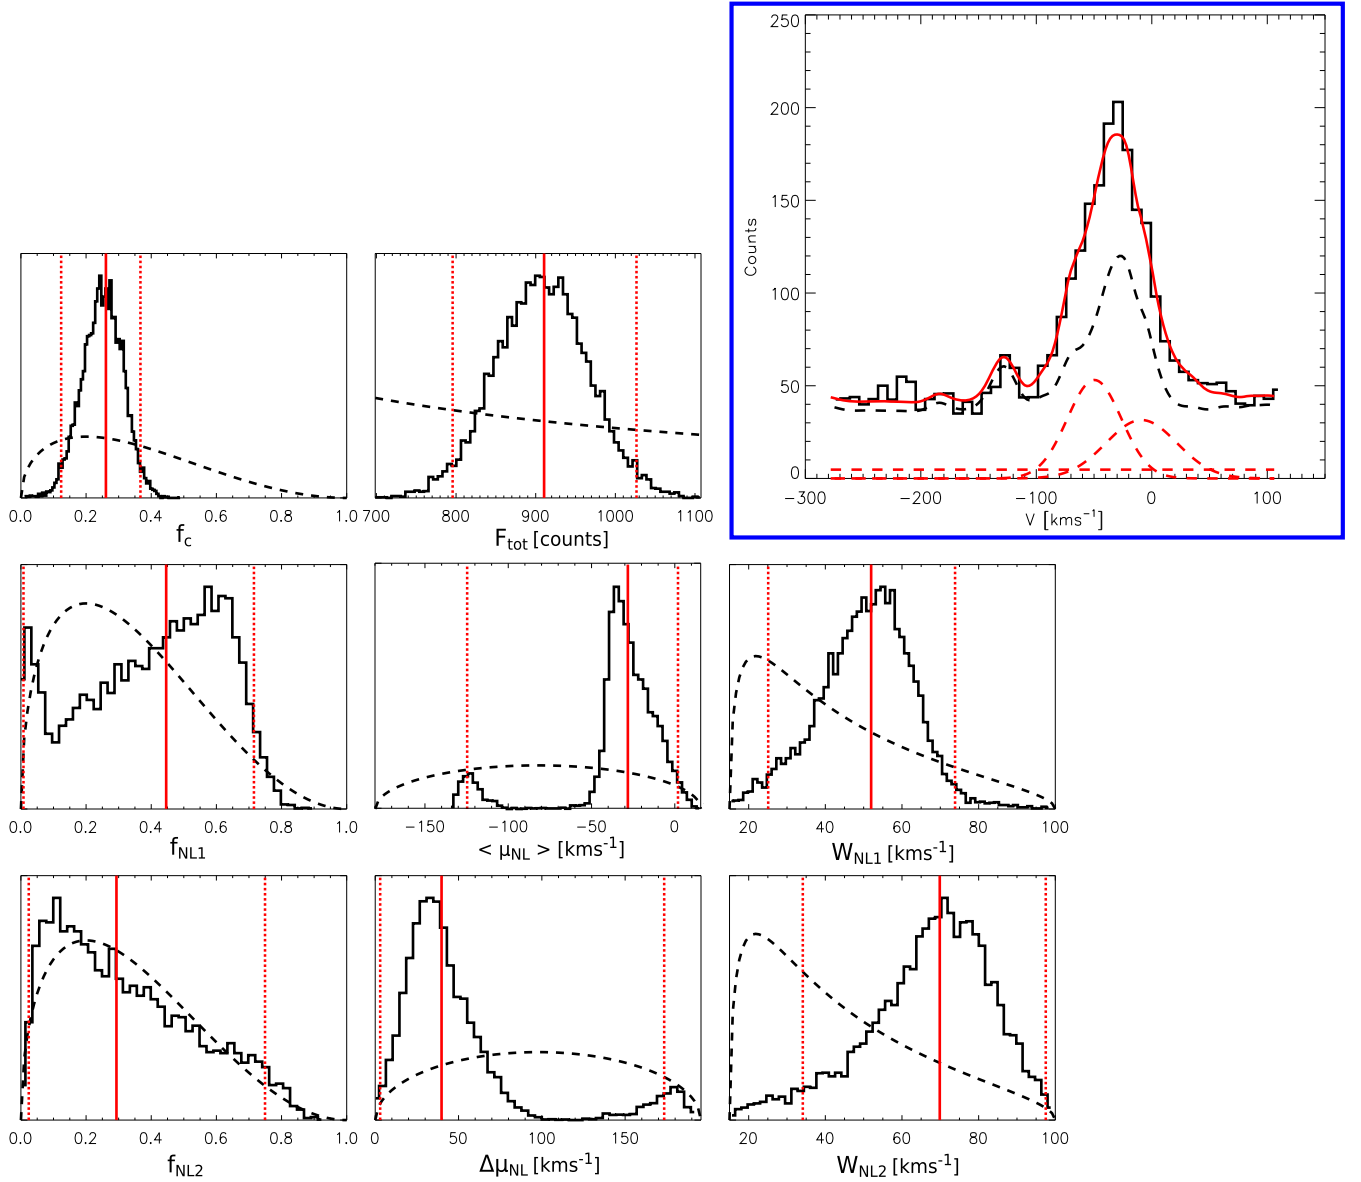
Figure 5. Parameter estimation via Bayesian inference for a bin in the NE filament of Tycho's SNR that requires a second NL. The top-right blue panel shows the observed spectrum (solid black line), the background model (dashed black line), and the components of the intrinsic median NLNL model (dashed red lines). The median model is overplotted with the solid red line. The remaining eight panels are the 1D-marginalized posteriors over model parameters (solid black lines): total flux (F_{tot} in counts), flux fractions in the continuum (f_c) and two narrow lines (f_{NL1}, f_{NL2}), NL centroid mean $\langle \mu_{NL} \rangle$, the separation between the two NLs $\Delta\mu_{NL}$, and intrinsic NL widths (W_{NL1}, W_{NL2}), with the latter four all given in km s^{-1}. The dashed black lines are the prior distributions and the vertical red lines are the estimated parameters of the median model, i.e., the median values (solid red), and the boundaries enclosing the highest density 95% confidence intervals (dashed red).

or higher. However, this result is dependent on the shock velocity assumption, which is 4000 km s^{-1} in the cited paper. Furthermore, Figure 10 in the same paper shows the IL width as a function of the CR acceleration efficiency ϵ for a fixed shock speed $V_{sh} = 4000$ km s^{-1}, $\eta_{TH} = 0.5$, and $p_{max} = 50$ TeV/c, and two values of downstream electron-to-proton temperature ratios β_{down} being either 0.01 or 1. Interestingly, our measured IL width of 180 km s^{-1} on average can be explained with the mentioned shock parameters and acceleration efficiency of 15%–25%, although any further constraints are difficult since we do not know β_{down}. The latter value, as well as the shock velocity, is something that can be constrained from the observations of the broad Hα component.

In order to obtain the measured $W_{NL} \approx 40$ km s^{-1}, one requires a combination of shock parameters (and CR acceleration included): $p_{max} = 10$ TeV/c, $V_{sh} = 2500$ km s^{-1}, and β_{down} of up to 0.1 (Morlino et al. 2013). At the same time, this configuration predicts an IL width of around 300 km s^{-1} and intermediate-to-narrow flux ratio of 1.6 on average for various values of η_{TH} and ϵ. Both of these IL parameters are much higher than the median values we infer from our data and analysis. We speculate that somewhat higher β_{down} than 0.1 (see Figure 13 in Morlino et al. 2012) and a shock velocity of around 3000 km s^{-1} would possibly be able to explain the observed $W_{IL} \approx 180$ km s^{-1}, $f_{IL}/f_{NL} \approx 0.41$, and $W_{NL} \approx 55$ km s^{-1}. Finally, we also notice that keeping η_{TH} constant and

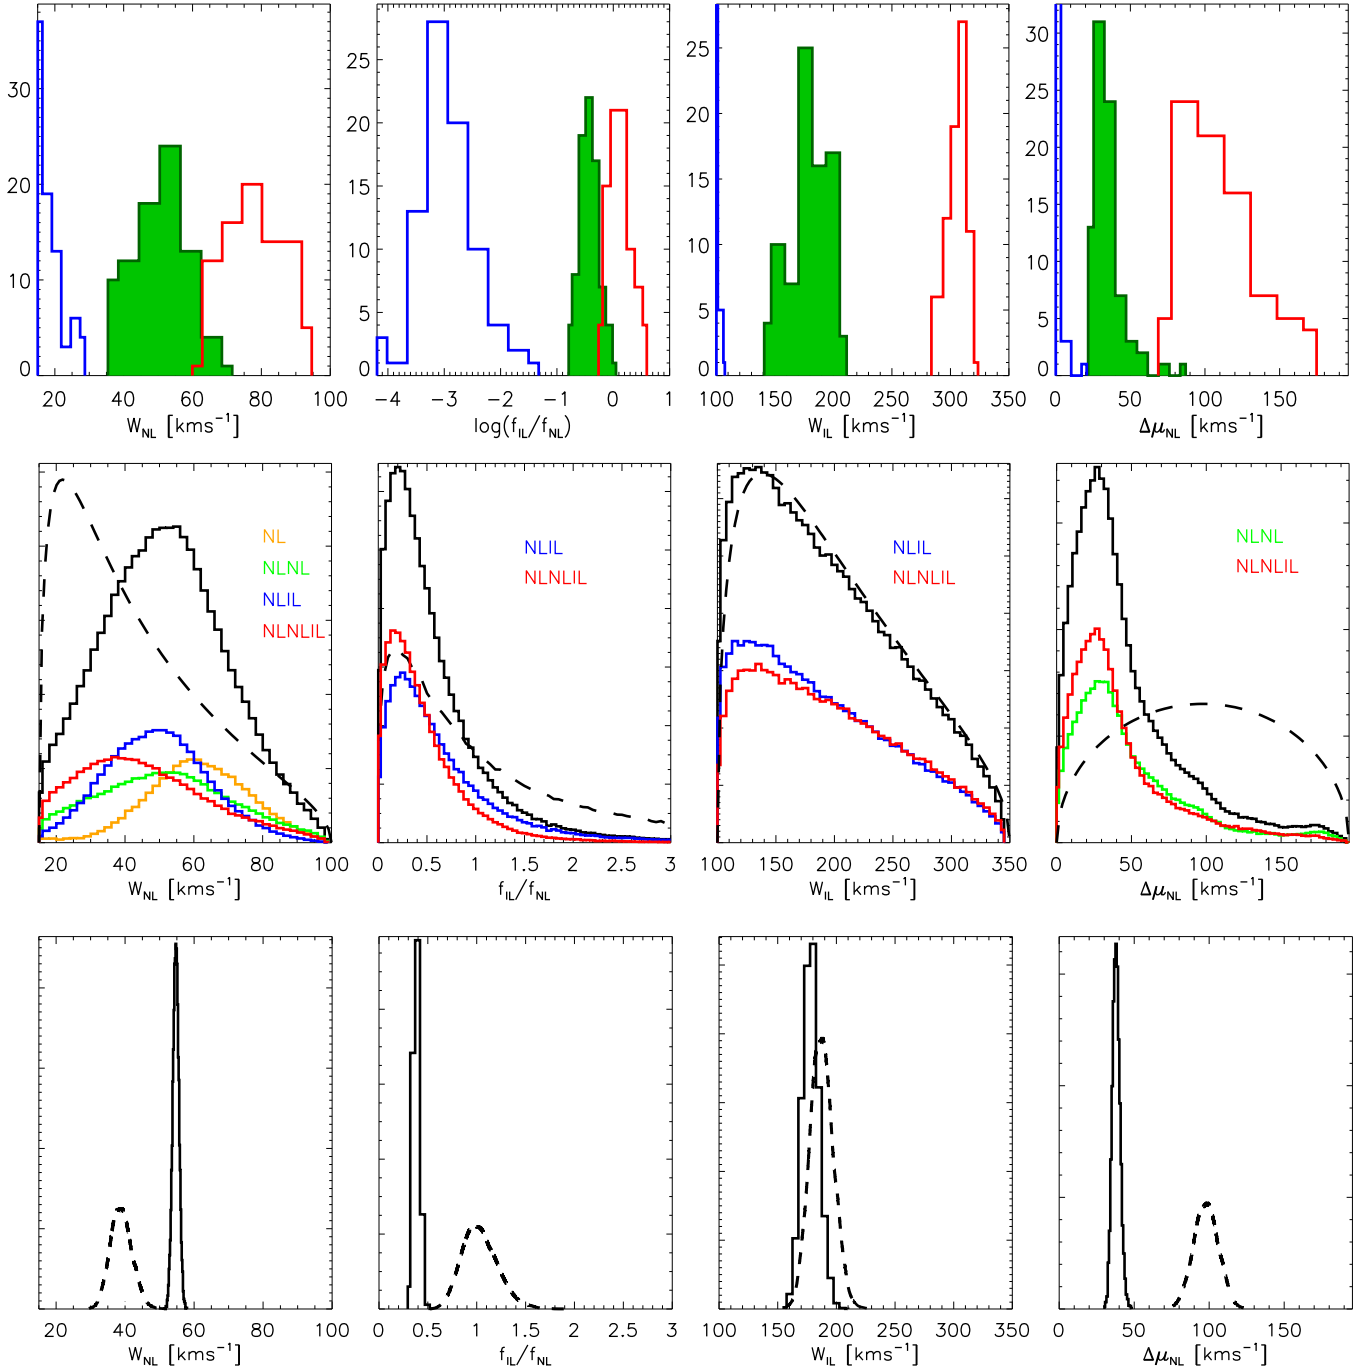**Figure 6.** Summary of results for the narrow-line width (W_{NL}, left panels), intermediate-to-narrow-line flux fraction (f_{IL}/f_{NL}, center-left), IL width (W_{IL}, center-right), and NL centroid separation ($\Delta\mu_{NL}$, rightmost panels). All figures are based on the marginalized posteriors of all 82 bins and 82×4 models, weighted within individual bins by relative model evidence. Top row: distribution of median (green) and highest density 95% confidence interval boundaries (blue, red), quantifying the variation across the filament. For the f_{IL}/f_{NL}, we adopted a log-scale in order to make the low-end confidence interval boundary histogram more visible. Middle: sum of all posteriors (solid black), illustrating the average posterior of an individual bin and the typical relative contributions of single-NL (orange), NLNL (green), NLIL (blue), and NLNLIL (red) models. Bottom: posterior of the median across all bins (solid), in effect using all data to constrain the respective parameter values. These are significantly narrower than the individual bin posterior. The prior is shown by the dashed curve in the middle and bottom panels, whereas the prior of the cross-bin median is given in the bottom panels.

increasing ϵ, the model of Morlino et al. (2013) predicts a simultaneous increase in W_{NL} and decrease in W_{IL} and f_{IL}/f_{NL}. We find a hint for such an anticorrelation between W_{NL} and W_{IL} (panel 7 in our Figure 8).

Average and variation of velocities. Our result for the median LSR-corrected LOS velocity, $V_{LSR} = (-34 \pm 2)$ km s^{-1},

is in agreement with the earlier study by Lee et al. (2007), who reported an NL Hα LSR velocity -30.3 km s^{-1} of "knot g." Furthermore, Ghavamian et al. (2017) reported a value of -45.6 km s^{-1}, which falls within our range $[-64, -25]$ km s^{-1} of observed LSR velocities in the NE rim. In H I 21 cm observations toward Tycho's NE rim, Reynoso et al. (1999)
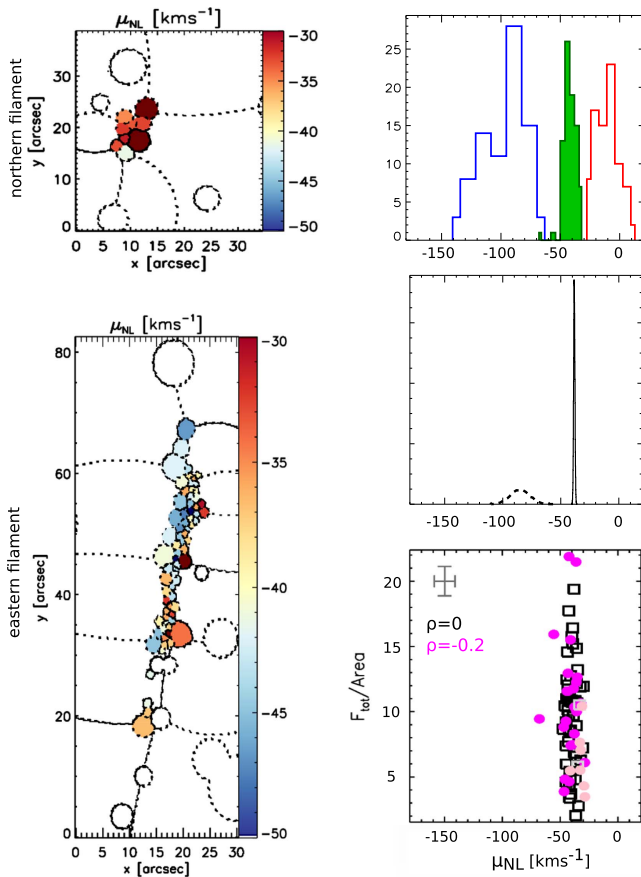
Figure 7. NL centroid (line-of-sight velocity) variation in the NE rim (left panels), distribution of median μ_{NL} (green) and highest density 95% confidence interval boundaries (blue, red) in the top-right panel, and posterior of the median across all bins (solid) using all data to constrain μ_{NL} (middle-right panel). The prior is overplotted with the dashed line. The northern (top-left) filament moves systematically with respect to the eastern filament, and the latter exhibits significant internal differential motion. The bottom-right panel shows that the NL centroid does not correlate with the surface brightness. The symbols and notation are the same as for Figure 8.

found $V_{LSR} = -51.5\,\mathrm{km\,s^{-1}}$ and associated the location of Tycho's SNR and the H I cloud with the Perseus arm. Similarly, ^{12}CO emission was found at $V_{LSR} = -62.5\ \mathrm{km\,s^{-1}}$ (Lee et al. 2004; see also Zhou et al. 2016) and might be associated with SN 1572's pre-shock gas. However, as pointed out by Tian & Leahy (2011), there is no clear evidence that either the H I cloud or CO cloud is physically associated with Tycho's NE rim. We leave further discussion of the V_{LSR} result and its possible interpretation to future work, where we will also investigate the shock and CR properties in more detail (see Section 6).

6. Summary and Conclusions

We present Hα spectroscopic observations of Tycho's NE Balmer filaments. This study provides spectroscopic data that, for the first time, are spatially resolved (spectro-imagery), with large coverage that comprises and resolves the entire NE filament. Our analysis is based on Bayesian inference, which enables a quantitative, probabilistic, and well-defined model comparison, and a comprehensive, complete characterization of the parameter probabilities.

We find that the broadening of the NL beyond $20\,\mathrm{km\,s^{-1}}$ that was noted in previous studies was not an artifact of the spatial integration, and that it extends across the whole filament, not only the previously covered "knot g." The NL width in the NE rim is typically found around $55\,\mathrm{km\,s^{-1}}$. Such a large width cannot be due to the superposition of multiple lines. In fact, our data analysis allows us to take projection effects into account when interpreting the data. We are able to distinguish between single-NL and double-NL models where we find significant evidence for a split in NL in 18% of the Voronoi bins. The widths of the two NLs are around $49\,\mathrm{km\,s^{-1}}$ and their centroid separations are $38\,\mathrm{km\,s^{-1}}$ on average.

An NL width of $55\,\mathrm{km\,s^{-1}}$ implies a temperature of the upstream gas of $\approx 68,000\,K$. If this were the temperature of the unperturbed ISM where the SNR is expanding, no neutral hydrogen would exist in the first place, contradicting the presence of Hα emission. Hence, our finding is the signature of the existence of a mechanism able to heat the upstream plasma in a region ahead of the shock much smaller than the collisional ionization length scale. As shown by the previous study of Morlino et al. (2013), a CR precursor is the best candidate to explain the widening of the NL, opening the possibility of studying particle acceleration in shocks using Hα emission. The fact that the NL width ranges from 35 to $72\,\mathrm{km\,s^{-1}}$ across the NE rim suggests that the amount of neutrals in the ambient medium varies, which imposes different degrees of damping of magnetic waves excited by CR streaming.

Likewise, we confirm the suspected presence of an IL and show it to be widespread (24% of the bins). Typical IL widths and intermediate-to-narrow flux ratios are $180\,\mathrm{km\,s^{-1}}$ and 0.41, respectively.

Our model parameters also comprise the LOS velocity centroids. After correction to the local standard of rest, their median is $V_{LSR} = (-34 \pm 2)\ \mathrm{km\,s^{-1}}$, in agreement with the Lee et al. (2007) investigation of "knot g."

Overall, our results reveal an interplay between two precursors in Tycho's NE rim: broadened NL widths point toward the evidence for the presence of a CR precursor, while detected IL reveals the presence of a BN precursor.

From the knowledge of the NL width only it is not possible to determine the CR acceleration efficiency, because such width depends on many parameters (shock speed, maximum energy of accelerated particles, electron−ion equilibration, and turbulent heating). Nevertheless, we can conclude that, assuming a shock speed between 2500 and $3000\,\mathrm{km\,s^{-1}}$, our result is compatible with having a maximum CR energy $>10\,\mathrm{TeV}$, a turbulent heating $>10\%$, and an acceleration efficiency $>$ few %. The degeneracy between these parameters could be broken using other information coming from the broad-line width and intensity (giving more precise information on shock speed and electron−ion equilibration) and X-ray/ gamma-ray observations (determining the maximum energy of accelerated particles). The difficulties in performing such calculations rely on the fact that the broad line is not known with the same accuracy as the NL and IL and that the gamma-ray emission does not have enough spatial and spectral resolution to fix unambiguously the maximum energy (Park 2015; Morlino & Blasi 2016). Improvements in this regards will surely come from the Cherenkov Telescope Array. Furthermore, parallel to the study of the GHαFaS narrow Hα-line profiles, we have conducted an investigation on the same part of the Tycho's remnant using OSIRIS (Optical System for

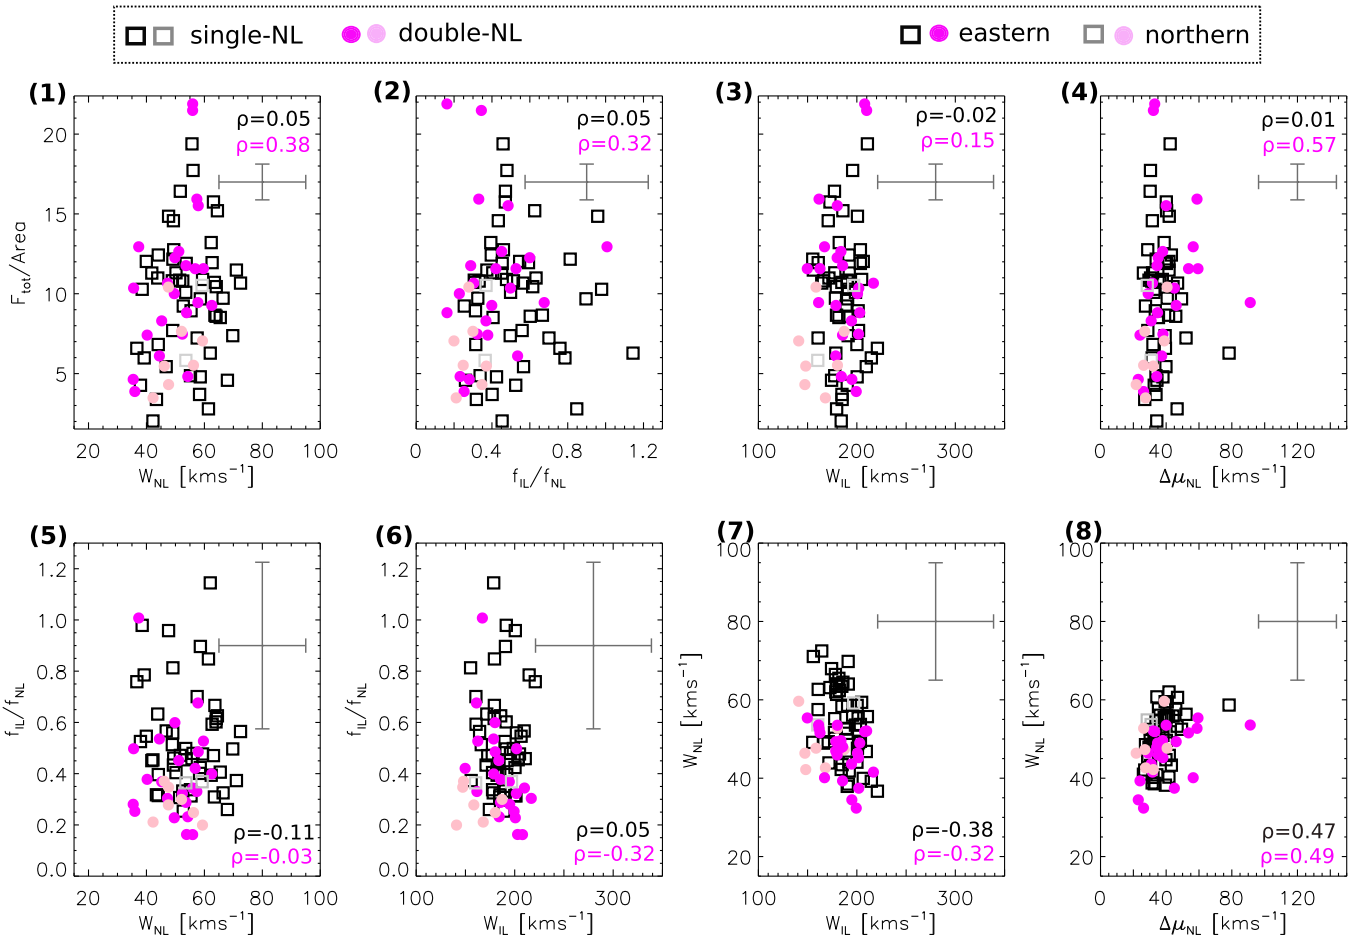**Figure 8.** Parameter correlation: top row shows the median of evidence-weighted W_{NL}, f_{IL}/f_{NL}, W_{IL}, and $\Delta\mu_{NL}$ posteriors and their dependence on the surface brightness, i.e., total model S flux over bin area (panels 1–4). In the bottom row, panels 5–8, the correlation among model parameters is tested. Squares (circles) refer to single (double) NL models in the eastern filament (black/magenta points) and northern filament (gray/pink points) based on the highest CV. Wherever we have the NLNL(IL) model as the one with the highest CV, W_{NL} is set to the average of the two NL widths. In panel 8 for bins favoring single-NL models, we use the posterior results for their NLNL(IL) models. Labeled are the Spearman's rank correlation coefficients ρ that correspond to the color-relevant points, where the black label refers to the black and gray points together, and the magenta label to the magenta and pink points. We indicate typical error bars on the surface brightness and parameters (average shortest 68% confidence intervals) in the right corner of every panel.
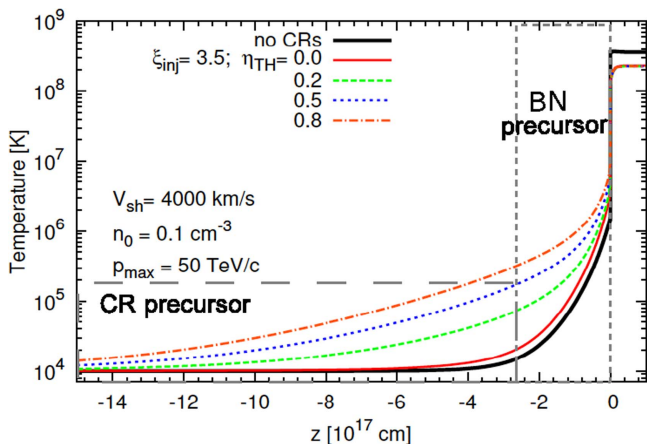

Figure 9. Heating and length scales of the CR and BN precursors for specific shock and CR properties (Morlino et al. 2013). Zero marks the location of the shock front, while the negative distance from the shock front represents the distance in the pre-shock region. As labeled, different lines present the extent and level of heating in the precursor for different amounts of turbulent heating η_{TH}. The black solid line shows the case without CRs and clearly shows the extent of the BN precursor alone. The immediate pre-shock region is affected by both BN and CR precursors where the temperature reaches several 10^6 K. The CR precursor extends much farther from the shock front ($>10^{17}$ cm) where the gas is heated to several 10^5 K.

Imaging and low-intermediate Resolution Spectroscopy) on the Gran Telescopio Canarias to observe the broad Hα-line profiles. In a forthcoming paper, we will present results of broad Hα components of the same spatial locations (bins) along the filaments as presented in this paper, which, combined with narrow Hα components and applied shock models, will give a better handle on the overall conditions in the shock and will enable us to quantify CR properties.

Appendix A
Data, Flatfield and Background

In order to compare our models of Tycho's SNR spectra to the data, we need to account for the variable sensitivity ("flatfield") and the background flux. We eschew the standard method of applying background subtraction and flatfield

Table 4
Median and Highest Density 95% Confidence Interval of the Parameter Posteriors in the Favored Model of Voronoi Bins

Bin	Model	W_{NL1} [km s^{-1}]	W_{NL2} [km s^{-1}]	W_{IL} [km s^{-1}]	$f_{\mathrm{IL}}/f_{\langle\mathrm{NL}\rangle}$	$\Delta\mu_{\mathrm{NL1}}$ [km s^{-1}]
1	NL	60.53 [43.84, 80.68]	⋯	⋯	⋯	⋯
2	NL	65.12 [48.10, 84.66]	⋯	⋯	⋯	⋯
3	NL	59.23 [44.61, 76.77]	⋯	⋯	⋯	⋯
4	NL	72.07 [49.95, 94.80]	⋯	⋯	⋯	⋯
5	NL	60.90 [44.14, 81.94]	⋯	⋯	⋯	⋯
6	NLIL	40.63 [24.78, 56.31]	⋯	191.02 [100.15, 308.96]	0.51 [0.01, 1.55]	⋯
7	NLNLIL	35.23 [15.02, 63.35]	31.68 [15.11, 79.41]	195.54 [102.05, 310.43]	0.26 [0, 0.64]	21.65 [0.04, 108.74]
8	NL	60.78 [45.34, 78.82]	⋯	⋯	⋯	⋯
9	NL	46.85 [32.28, 61.53]	⋯	⋯	⋯	⋯
10	NL	40.87 [28.68, 54.37]	⋯	⋯	⋯	⋯
11	NL	54.04 [35.29, 73.60]	⋯	⋯	⋯	⋯
12	NL	78.70 [62.54, 95.65]	⋯	⋯	⋯	⋯
13	NL	70.71 [49.50, 92.40]	⋯	⋯	⋯	⋯
14	NLNLIL	40.34 [15.76, 81.33]	35.18 [15.50, 72.99]	191.06 [102.59, 311.14]	0.33 [0, 0.80]	22.74 [0.16, 101.29]
15	NL	71.88 [43.78, 99.09]	⋯	⋯	⋯	⋯
16	NL	75.86 [56.59, 96.57]	⋯	⋯	⋯	⋯
17	NLIL	39.13 [22.30, 56.71]	⋯	212.36 [100.18, 307.88]	0.88 [0.01, 6.15]	⋯
18	NL	82.43 [64.85, 98.87]	⋯	⋯	⋯	⋯
19	NLNL	45.04 [15.39, 90.16]	65.68 [39.58, 90.64]	⋯	⋯	91.34 [41.02, 115.24]
20	NL	74.45 [54.71, 96.58]	⋯	⋯	⋯	⋯
21	NL	73.00 [47.82, 98.64]	⋯	⋯	⋯	⋯
22	NL	68.76 [50.01, 88.42]	⋯	⋯	⋯	⋯
23	NLNLIL	48.04 [15.86, 88.08]	49.68 [19.25, 84.92]	168.44 [100.05, 295.69]	0.42 [0, 1.78]	51.99 [11.42, 153.20]
24	NL	59.97 [43.21, 77.48]	⋯	⋯	⋯	⋯
25	NL	66.63 [36.35, 96.93]	⋯	⋯	⋯	⋯
26	NL	70.97 [47.59, 95.35]	⋯	⋯	⋯	⋯
27	NL	58.66 [39.30, 81.32]	⋯	⋯	⋯	⋯
28	NL	56.87 [41.50, 74.12]	⋯	⋯	⋯	⋯
29	NL	56.07 [36.44, 79.31]	⋯	⋯	⋯	⋯
30	NLIL	59.03 [21.68, 93.56]	⋯	181.53 [103.90, 288.14]	1.44 [0, 7.08]	⋯
31	NL	56.38 [37.83, 74.36]	⋯	⋯	⋯	⋯
32	NLIL	41.53	⋯	150.76	1.02	⋯

Table 4
(Continued)

Bin	Model	W_{NL1} [km s^{-1}]	W_{NL2} [km s^{-1}]	W_{IL} [km s^{-1}]	$f_{IL}/f_{\langle NL \rangle}$	$\Delta\mu_{NL}$ [km s^{-1}]
		[15.31, 76.64]		[100.17, 278.80]	[0, 3.29]	
33	NL	38.17	...	...	...	...
		[22.16, 56.19]				
34	NL	47.04	...	...	...	...
		[30.20, 65.05]				
35	NL	72.60	...	...	...	...
		[56.45, 89.57]				
36	NL	59.26	...	...	...	...
		[40.97, 80.03]				
37	NL	42.92	...	...	...	...
		[23.91, 62.21]				
38	NLIL	41.82	...	170.06	0.74	...
		[21.28, 66.79]		[100.12, 294.73]	[0, 2.26]	
39	NLNL	51.41	56.03	...	...	60.15
		[16.46, 90.57]	[29.12, 80.18]			[9.50, 84.17]
40	NLNL	38.06	63.21	...	...	39.90
		[15.61, 61.48]	[26.59, 96.32]			[1.32, 144.16]
41	NLIL	48.34	...	201.48	0.62	...
		[24.99, 70.31]		[103.46, 310.50]	[0.03, 1.67]	
42	NLIL	48.09	...	162.60	0.61	...
		[27.07, 72.44]		[100.39, 291.90]	[0.01, 1.82]	
43	NLNLIL	44.57	33.42	171.73	0.75	55.59
		[15.43, 87.20]	[15.03, 77.23]	[100.31, 295.94]	[0.01, 2.23]	[1.57, 97.03]
44	NL	65.49	...	...	...	...
		[48.42, 85.13]				
45	NL	55.89	...	...	...	...
		[38.00, 76.15]				
46	NLIL	35.29	...	221.61	0.89	...
		[17.55, 53.39]		[100.39, 317.73]	[0, 10.90]	
47	NL	71.18	...	...	...	...
		[51.66, 91.45]				
48	NL	49.41	...	...	...	...
		[32.16, 66.58]				
49	NL	43.21	...	...	...	...
		[27.87, 63.95]				
50	NL	61.30	...	...	...	...
		[42.56, 83.35]				
51	NLNLIL	50.84	37.90	201.87	0.31	37.67
		[16.21, 81.76]	[16.10, 66.80]	[104.16, 313.03]	[0, 0.89]	[1.03, 90.84]
52	NL	61.42	...	...	...	...
		[39.35, 85.21]				
53	NL	54.00	...	...	...	...
		[38.21, 71.73]				
54	NL	47.73	...	...	...	...
		[23.04, 68.53]				
55	NLNLIL	33.67	50.83	187.38	0.50	34.31
		[15.24, 72.78]	[17.70, 88.41]	[101.07, 312.58]	[0, 1.49]	[1.00, 82.17]
56	NL	58.08	...	...	...	...
		[38.88, 80.69]				
57	NLIL	45.30	...	199.89	1.12	...
		[17.24, 72.11]		[107.28, 312.43]	[0.02, 3.11]	
58	NL	74.99	...	...	...	...
		[54.48, 97.77]				
59	NLNL	57.49	58.12	...	...	63.14
		[19.54, 93.95]	[36.03, 82.85]			[11.53, 93.23]
60	NL	66.49	...	...	...	...
		[46.53, 89.49]				
61	NL	73.42	...	...	...	...
		[52.04, 96.65]				
62	NL	55.28	...	...	...	...
		[37.46, 75.62]				
63	NLIL	49.93	...	202.95	0.49	...
		[32.62, 69.52]		[106.47, 316.80]	[0, 1.34]	

Table 4
(Continued)

Bin	Model	W_{NL1} [km s^{-1}]	W_{NL2} [km s^{-1}]	W_{IL} [km s^{-1}]	$f_{IL}/f_{\langle NL \rangle}$	$\Delta\mu_{NL}$ [km s^{-1}]
64	NLIL	38.86 [19.78, 60.12]	⋯	191.52 [107.16, 298.62]	1.12 [0.12, 2.53]	⋯
65	NL	68.07 [47.10, 92.23]	⋯	⋯	⋯	⋯
66	NL	57.62 [42.10, 77.46]	⋯	⋯	⋯	⋯
67	NL	65.35 [46.28, 87.53]	⋯	⋯	⋯	⋯
68	NL	72.89 [57.12, 89.85]	⋯	⋯	⋯	⋯
69	NL	60.22 [41.25, 82.71]	⋯	⋯	⋯	⋯
70	NL	46.09 [27.58, 67.10]	⋯	⋯	⋯	⋯
71	NL	50.96 [29.85, 76.82]	⋯	⋯	⋯	⋯
72	NL	78.28 [61.10, 96.81]	⋯	⋯	⋯	⋯
73	NL	64.24 [42.79, 89.21]	⋯	⋯	⋯	⋯

Note. Shown here are the spatial bins in the eastern filament of Tycho's SNR. If none of the multi-line models is at least ≈3 times (or 0.5 dex) more likely, the single-line (NL) model is taken to be the favored one.

Table 5
Same as Table 4 Just for Voronoi Bins in Tycho's Northern Filament

Bin	Model	W_{NL1} [km s^{-1}]	W_{NL2} [km s^{-1}]	W_{IL} [km s^{-1}]	$f_{IL}/f_{\langle NL \rangle}$	$\Delta\mu_{NL}$ [km s^{-1}]
1	NLNL	33.85 [15.04, 61.64]	53.16 [21.03, 87.55]	⋯	⋯	28.04 [0.29, 156.41]
2	NLIL	58.50 [41.77, 75.74]	⋯	195.28 [102.92, 316.25]	0.39 [0, 1.13]	⋯
3	NLIL	54.32 [37.83, 69.52]	⋯	156.07 [100.12, 274.97]	0.41 [0.02, 1.08]	⋯
4	NL	59.63 [37.09, 78.26]	⋯	⋯	⋯	⋯
5	NLNL	51.95 [25.10, 73.85]	69.86 [34.12, 97.51]	⋯	⋯	39.81 [2.97, 173.42]
6	NLNL	42.30 [15.65, 83.53]	56.58 [20.46, 91.93]	⋯	⋯	20.37 [0.05, 67.38]
7	NLNL	37.89 [15.09, 66.94]	60.23 [26.97, 93.47]	⋯	⋯	42.00 [6.73, 184.78]
8	NLNLIL	44.66 [15.68, 70.78]	34.55 [15.22, 80.12]	149.45 [100.36, 291.76]	0.36 [0, 0.89]	29.63 [0.23, 86.09]
9	NL	56.38 [37.58, 74.56]	⋯	⋯	⋯	⋯

correction to compensate for both effects in the individual exposures. Instead, we reconstruct the flatfield and background in the final product of the data reduction pipeline, which is a cube (position, wavelength) of co-added, "stacked" individual observations. Notably, the flatfield will also have a wavelength dependence in addition to being position dependent. We eventually include the flatfield and background in our models. The advantage of our approach is the preservation of photon statistics and therefore accurate uncertainties of the model parameters and evidence.

Before describing the construction of the individual as well as the co-added flatfield and background, we formalize the process through which the data are generated and how they are propagated by the data reduction pipeline.

A.1. Measurement Process

Each datum (measurement) $D_{xy,i}$ with pixel indices (x,y) and exposure index i is the response of the telescope and instrument to the incoming, seeing-convolved flux $d_0 = d_0(\alpha, \delta, \lambda)$,

which varies with sky coordinates (α, δ) and wavelength λ:

$$D_{xy,i} = \int_{-\infty}^{\infty} d_{0,i}(x, y, \lambda)\, F_{xy}\, R_{xy,i}(\lambda)\, d\lambda, \qquad (3)$$

where

$$d_{0,i}(x, y, \lambda) \equiv d_0(\alpha = \alpha_i(x, y), \delta = \delta_i(x, y), \lambda)$$

has the astrometric solution $\alpha = \alpha_i(x, y)$, $\delta = \delta_i(x, y)$, and the integral over the area of pixel (x,y) implicitly applied.

F_{xy} is the spatially varying sensitivity of the detector and optical system—the flatfield image. As expected, and as verified by us (see below), it is the same for all exposures.

$R_{xy,i}(\lambda)$ is the pixel- and exposure-specific wavelength filter imposed by the Fabry–Pérot interferometer (etalon). In our case, it is accurately described by a universal line-spread function (LSF), $r(\lambda)$:

$$R_{xy,i}(\lambda) = r(\Lambda_{xy,i} - \lambda), \qquad (4)$$

where the wavelength calibration $\Lambda_{xy,i}$ returns the central, maximum-throughput wavelength for each pixel and exposure (tuning of the etalon). By definition, the LSF is centered on (peaks at) the origin. In our case, it is a Gaussian with dispersion measured from calibration spectra.

Inserting Equation (4) into Equation (3), one sees that the datum $D_{xy,i}$ is the LSF-convolved spectral flux $d_i(x, y, \lambda)$, evaluated at $\Lambda_{xy,i}$:

$$\begin{aligned}
D_{xy,i} &= \int_{-\infty}^{\infty} d_{0,i}(x, y, \lambda)\, F_{xy}\, r(\Lambda_{xy,i} - \lambda)\, d\lambda \\
&= F_{xy} \cdot (d_{0,i} *_{\lambda} r)(x, y, \lambda = \Lambda_{xy,i}) \\
&\equiv F_{xy} \cdot d_i(x, y, \Lambda_{xy,i}). \qquad (5)
\end{aligned}$$

A.2. Data Processing

By way of converting the observed images into data subcubes, $D_{xy,i} \rightarrow D_{xyl,i}$, we make the wavelength information contained in them explicit and obtain a format that allows direct co-addition. Each of the 48 "slices" (third index) of a subcube corresponds to a wavelength, λ_l. Each observed $\Lambda_{xy,i}$ is bracketed by two slices, and the corresponding flux $D_{xy,i}$ is assigned to them via linear interpolation:

$$\begin{aligned}
D_{xyl,i} &= D_{xy,i} \cdot T_{xyl,i} \\
T_{xyl,i} &= t_l(\Lambda_{xy,i}) \\
&\equiv \begin{cases} 1 - |\Lambda_{xy,i} - \lambda_l|/\Delta\lambda & : |\Lambda_{xy,i} - \lambda_l| \leqslant \Delta\lambda \\ 0 & : |\Lambda_{xy,i} - \lambda_l| > \Delta\lambda \end{cases}, \qquad (6)
\end{aligned}$$

where $\Delta\lambda$ is the "size" of each slice, i.e., the distance between the slices' central wavelengths, λ_l. Another way to describe this assignment is that each slice imposes a triangle filter $t_l(\lambda) = \max(0, 1 - |\lambda - \lambda_l|/\Delta\lambda)$. We will use it again to construct the co-added flatfield and background cubes.

As $D_{xy,i}$ itself is a filtered version of the flux, the subcube $D_{xyl,i}$ at a given pixel (x, y) is nearly "empty," except for two slices l. Adequate coverage of the spectrum therefore necessitates multiple tunings of the etalon, or imaging the source in varying locations on the detector, since even for unchanged tuning, Λ depends on x and y. The resulting multiple subcubes are then co-added, albeit after projecting and spatially resampling them onto a common astrometrically
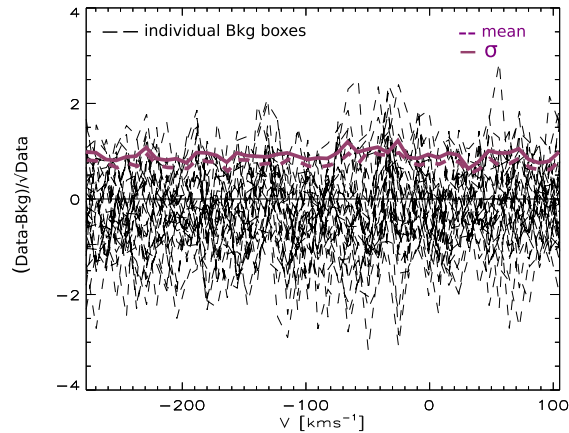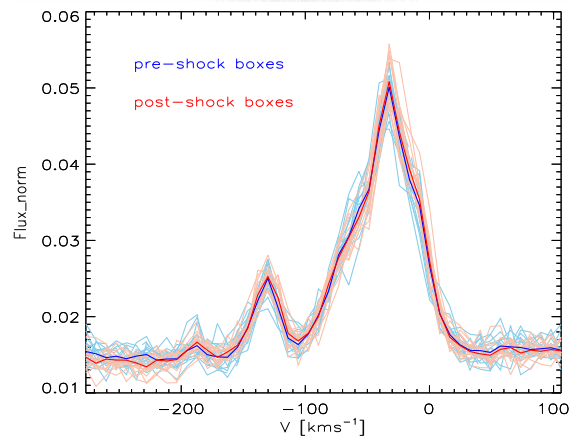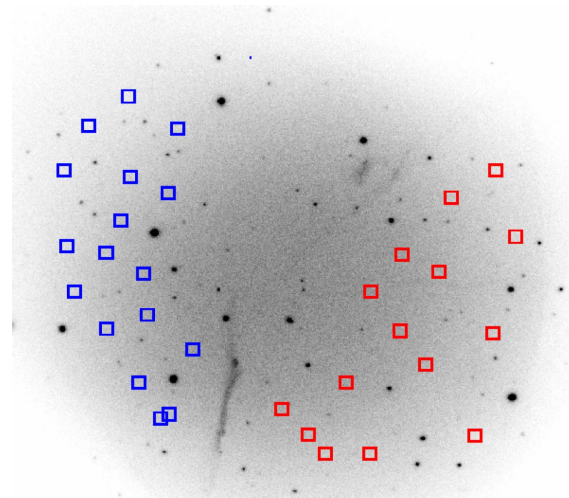

Figure 10. The top panel shows the wavelength-integrated data cube $D_{x'y'}$. The effect of the flatfield is clearly visible as a radial surface-brightness pattern of the background flux in addition to inhomogeneous coverage (empty corners). Overlaid are $(25\,\text{pix})^2 \approx (5'')^2$ boxes in the pre-shock/post-shock (blue/red) region, which we used to estimate the incident background spectrum $b(\lambda)$ and, in turn, the flatfield. The middle panel presents the normalized spectra of the boxes (light blue/red lines), while the thick blue and red lines are the mean profiles of the pre- and post-shock boxes. The shapes of the background spectra are similar across the entire FOV. The bottom panel indicates differences between the actual background ("Data") and the background model ("Bkg"), normalized by the Poisson noise. The dashed black lines represent the differences in the 32 background boxes separately, the dashed purple line is the mean absolute difference, and the solid purple line is the standard deviation between the boxes.

calibrated frame (x', y'):

$$D_{xyl,i} \xrightarrow{\alpha,\delta} D_{x'y'l,i} \tag{7}$$

$$D_{x'y'l} = \sum_i D_{x'y'l,i}. \tag{8}$$

$D_{x'y'l}$ is the final, co-added data cube which, apart from spatial binning, directly constrains the SNR shock models. We account for the spatial sensitivity variations and any residual non-constant spectral sampling rate by including them in the models, in the form of the co-added flatfield cube, $F_{x'y'l}$, which we derive below along with the background cube, $B_{x'y'l}$.

A.3. Propagation of Flatfield and Background

We represent the incoming flux $d_{0,i}$ as the sum of the "source" flux ($s_{0,i}$) that comprises celestial objects, in particular the SNR filament that we are interested in, and the background flux ($b_{0,i}$). Because of the linearity of the integral in Equation (5), the same applies to the LSF-convolved flux d_i:

$$d_{0,i}(x, y, \lambda) = s_{0,i}(x, y, \lambda) + b_{0,i}(x, y, \lambda)$$
$$s_i(x, y, \lambda) = (s_{0,i} *_\lambda r)(x, y, \lambda)$$
$$b_i(x, y, \lambda) = (b_{0,i} *_\lambda r)(x, y, \lambda)$$
$$\xRightarrow{(3)} d_i(x, y, \lambda) = s_i(x, y, \lambda) + b_i(x, y, \lambda)$$
$$D_{xy,i} = F_{xy} \cdot [s_i(x, y, \Lambda_{xy,i}) + b_i(x, y, \Lambda_{xy,i})]$$
$$\equiv F_{xy} \cdot S_{xy,i} + B_{xy,i} \ .$$

Note that in contrast to the source component of the data,

$$S_{xy,i} = s_i(x, y, \Lambda_{xy,i}), \tag{9}$$

we have absorbed the flatfield in the definition of the data's background component,

$$B_{xy,i} = F_{xy} \cdot b_i(x, y, \Lambda_{xy,i}). \tag{10}$$

Due to the linearity of the wavelength assignment (6),

$$D_{xyl,i} = T_{xyl,i} F_{xy} S_{xy,i} + T_{xyl,i} B_{xy,i}$$
$$\equiv F_{xyl,i} S_{xy,i} + B_{xyl,i}.$$

Here we have defined the background subcube $B_{xyl,i} = T_{xyl,i} B_{xy,i}$, and the flatfield subcube $F_{xyl,i} = T_{xyl,i} F_{xy}$, which are the background components of the data and the flatfield image (the response to a flat spectrum), respectively, transformed to a subcube using the same slice-assignment (6) that was used for the data itself. Also, as for the data, background and flatfield subcubes are projected onto the same frame and co-added:

$$F_{xyl,i} \xrightarrow{\alpha,\delta} F_{x'y'l,i}$$
$$S_{xy,i} \xrightarrow{\alpha,\delta} S_{x'y',i}$$
$$B_{xyl,i} \xrightarrow{\alpha,\delta} B_{x'y'l,i}$$
$$D_{x'y'l} = \sum_i (F_{x'y'l,i} S_{x'y',i} + B_{x'y'l,i})$$

$S_{x'y',i} = s_i(x', y', \Lambda_{x'y',i})$ is $s_i(x, y, \Lambda_{xy,i})$ with $x = x_i(x')$, $y = y_i(y')$ given by the astrometric solution. The same transformation applies to b. Since the co-added frame's (x', y') corresponds to a unique sky position, and we assume the filament emission (but not necessarily the background!) to be

constant between exposures, $s_i(x', y', \lambda) = s(x', y', \lambda)$. Then,

$$D_{x'y'l} = \sum_i F_{x'y'l,i} \, s(x', y', \Lambda_{x'y',i})$$
$$+ F_{x'y'l,i} \, b_i(x', y', \Lambda_{x'y',i})$$
$$\approx \left(\sum_i F_{x'y'l,i}\right) s(x', y', \lambda_l)$$
$$+ \sum_i F_{x'y'l,i} \, b_i(x', y', \lambda_l)$$
$$\equiv F_{x'y'l} \cdot s(x', y', \lambda_l) + B_{x'y'l}. \tag{11}$$

In the second line, we took $F_{x'y'l,i}$ to be non-zero only for l nearest to the sampled wavelength, $\lambda_l \approx \Lambda_{x'y',i}$. Hence, $F_{x'y'l,i} \cdot s(x', y', \Lambda_{x'y',i}) \approx F_{x'y'l,i} \cdot s(x', y', \lambda_l)$ for all $l = 1, \ldots, 48$, and analogously for b. However, the background does not factor out of the co-addition, as it may change between exposures. In the last step, we have defined the flatfield and background cubes $F_{x'y'l}$ and $B_{x'y'l}$ as the sum of the flatfield and background subcubes.

A.4. Background and Flatfield Model

In order to isolate the source flux component of the data, we still need to model the flatfield and background.

We begin by investigating the background on the co-added data (8), (11). Using SEXTRACTOR, we mask sources (stars, galaxies, and the filament itself) on the wavelength-integrated frame $D_{x'y'} = \sum_{l=1}^{48} D_{x'y'l}$. This frame is the deepest, highest S/N data product we have available, hence the mask is as complete as possible. We obtain the set of unmasked pixels, (x'_u, y'_u), which we can use to measure the background without "contamination" by celestial sources:

$$D_{x'_u y'_u l} = \sum_i B_{x_u' y_u' l,i} = B_{x_u' y_u' l}.$$

Next, in order to ameliorate potential undetected low surface-brightness source flux incursion to measure the variability of B and to average over spectral sensitivity variations in $D_{x'y'l}$, we select 32 boxes of $(25 \text{ pix})^2 \approx (5'')^2$, in particularly object-poor locations. For each box $\mathcal{B}_j = \{(x'_u, y'_u)_j\}$, we measure the spatially integrated spectrum:

$$B_l^{(j)} = \sum_{(x'_u, y'_u) \in \mathcal{B}_j} B_{x'_u y'_u l}$$
$$\overset{\text{Equation (11)}}{=} \sum_{(x'_u, y'_u) \in \mathcal{B}_j} \sum_i F_{x'_u y'_u l,i} \, b_i(x'_u, y'_u, \lambda_l).$$

It turns out that, apart from flux normalization, *all $B_l^{(j)}$ are nearly the same*, with differences at the percent level (see Figure 10). This allows us to conclude that

(i) the spectral part of the co-added flatfield is spatially invariant, as desired and expected after co-adding ~ 2500 exposures with different pointing, orientation and wavelength tuning plus averaging over a $\sim 5''$ box

(ii) the spectral shape of the background component is spatially invariant:

$$b_i(x', y', \lambda) = a_i(x', y') \cdot b(\lambda), \tag{12}$$
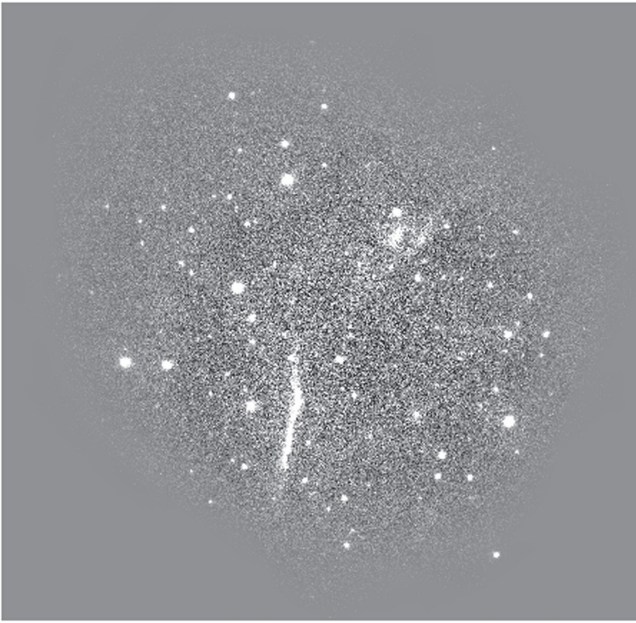
where a_i is the amplitude of the background.

Figure 11. Wavelength-integrated background residuals $(D - B)_{x'y'}$. The gray scale is linear from -10 to $+10$ counts/pixel. The smooth gray regions have low signal or are even zero due to low effective exposure time and hence lower flux (the flatfield was not divided out). The sources (white) were masked in the flatfield and background construction procedure and additionally avoided by the background-probing boxes. This applies in particular also to the apparent faint but extended brightness around the eastern (left) and northern (top) filaments.

Proof comes by contradiction: if the spectral shape of the 5″ box-averaged flatfield or background varied with location, $B_l^{(j)}$ would vary. Notably, we also see no systematic change of the background between the pre- and post-shock regions. Variations of the background flux on smaller ($<5″$) spatial scales are physically unlikely and in any case must be small as the observed variability can entirely be accounted for by Poisson (measurement) noise (the middle and bottom panels in Figure 10). In the bottom panel of Figure 10, we present differences between the actual background and the modeled background. There is no systematic effect, and the spatial variation is within the measurement uncertainty, which in turn is much smaller than the signal. The variation between background boxes, as measured by the 32 element sample standard deviation normalized by the photon noise, is 0.95 on average across the spectrum (solid purple line), with maximum value of 1.22. The average absolute deviation (dashed purple line) is 0.77 on average, close to the theoretical value of 0.8 expected for the absolute value of a standard normally distributed variable. Therefore, any differences between the actual background ("Data") and model ("Bkg") are fully explained by measurement uncertainties.

We can therefore measure $b(\lambda)$ on the co-added data cube. By combining all 32 background boxes, we additionally minimize noise and residual systematics:

$$b(\lambda) = \sum_j B_l^{(j)} \equiv \tilde{b}(\lambda).$$

The tilde indicates that $\tilde{b}$ is a model of the background spectrum. The normalization of $b(\lambda)$ is absorbed in a_i. We now assume that a_i is spatially invariant, i.e., that the background amplitude changes only between exposures but

not across the FOV:

$$a_i(x', y') = a_i.$$

Therefore, also $a_i(x, y) = a_i$ is spatially constant, and we have

$$D_{x_u y_u, i} = B_{x_u y_u, i} = F_{x_u y_u} \cdot a_i\, \tilde{b}(\Lambda_{x_u y_u, i}).$$

Here, (x_u, y_u) are the non-masked pixels of the co-added frame, reprojected onto the individual exposures. We do not know the background amplitudes a_i yet, but can already use the knowledge of $b(\lambda)$ to model F_{xy}:

$$\left(\sum_i \frac{D_{\tilde{x}_u \tilde{y}_u, i}}{\tilde{b}(\Lambda_{\tilde{x}_u \tilde{y}_u, i})} \right) \xrightarrow{\text{fit+norm.}} \tilde{F}_{xy}. \tag{13}$$

This way, we "divide out" the non-constant background spectrum, which is imprinted on the measurement via $\Lambda_{xy,i}$. The sum on the left side is only carried out for pixels $(\tilde{x}_u, \tilde{y}_u)$ that are not masked in any of the exposures. The result of the sum is a non-normalized flatfield image, which is then modeled by a fourth-order polynomial. The fit eliminates pixel noise and interpolates over masked pixels. It is followed by normalization, such that $\sum_{x,y} \tilde{F}_{xy} = N_{\text{pix}}$.

Now, we justify that the background amplitude is spatially constant, and even that the flatfield is indeed constant as normally expected. We cannot prove this for individual exposures, as the fluxes are impractically small. However, we can restrict the sum in Equation (13) to different subsets of exposures. This way, we derived $\tilde{F}_{xy}$ for exposures of only one specific "channel" (tuning) and for subsets of observations at different times. The results are invariant within a few percent: random variations of 3.2% (mean pixel standard deviation) between flatfields of different observing runs, and 1.5% among channel-specific flatfields. Again, the conclusion is by contradiction: given that exposures vary in pointing, field rotation, and tuning, a non-constant flatfield or a spatially variable background amplitude would lead to systematic variations of $\tilde{F}$, but we do not observe such variations.

With $\tilde{F}_{xy}$ on hand, we measure a_i as the flux scaling required to match the data,

$$a_i = \frac{\sum_{x_u, y_u} D_{x_u y_u, i}}{\sum_{x_u, y_u} \tilde{F}_{xy} \cdot \tilde{b}(\Lambda_{x_u y_u, i})} \,,$$

and reconstruct the background in each exposure:

$$\tilde{B}_{xy, i} = \tilde{F}_{xy} \cdot a_i\, \tilde{b}(\Lambda_{xy, i}). \tag{14}$$

We then transform the flatfield and background models $\tilde{F}_{xy}$ and $\tilde{B}_{xy,i}$ to subcubes in the common astrometric frame and co-add them, in the same way as the data in Equations (6)–(8). As for the "real" flatfield and background, Equation (11) ensures that the resulting co-added cubes correspond to the actual flatfield and background components in the data cube.

We use $\tilde{B}_{x'y'l}$ and $\tilde{F}_{x'y'l}$ to check once more whether our assumption of a uniform background and invariant flatfield are fulfilled: the wavelength-integrated $(B/F)_{x'y'} = \sum_l B_{x'y'l}/F_{x'y'l}$ is indeed flat, and the background residuals $(D - B)_{x'y'}$ are zero apart from Poisson noise and objects (see Figure 11).
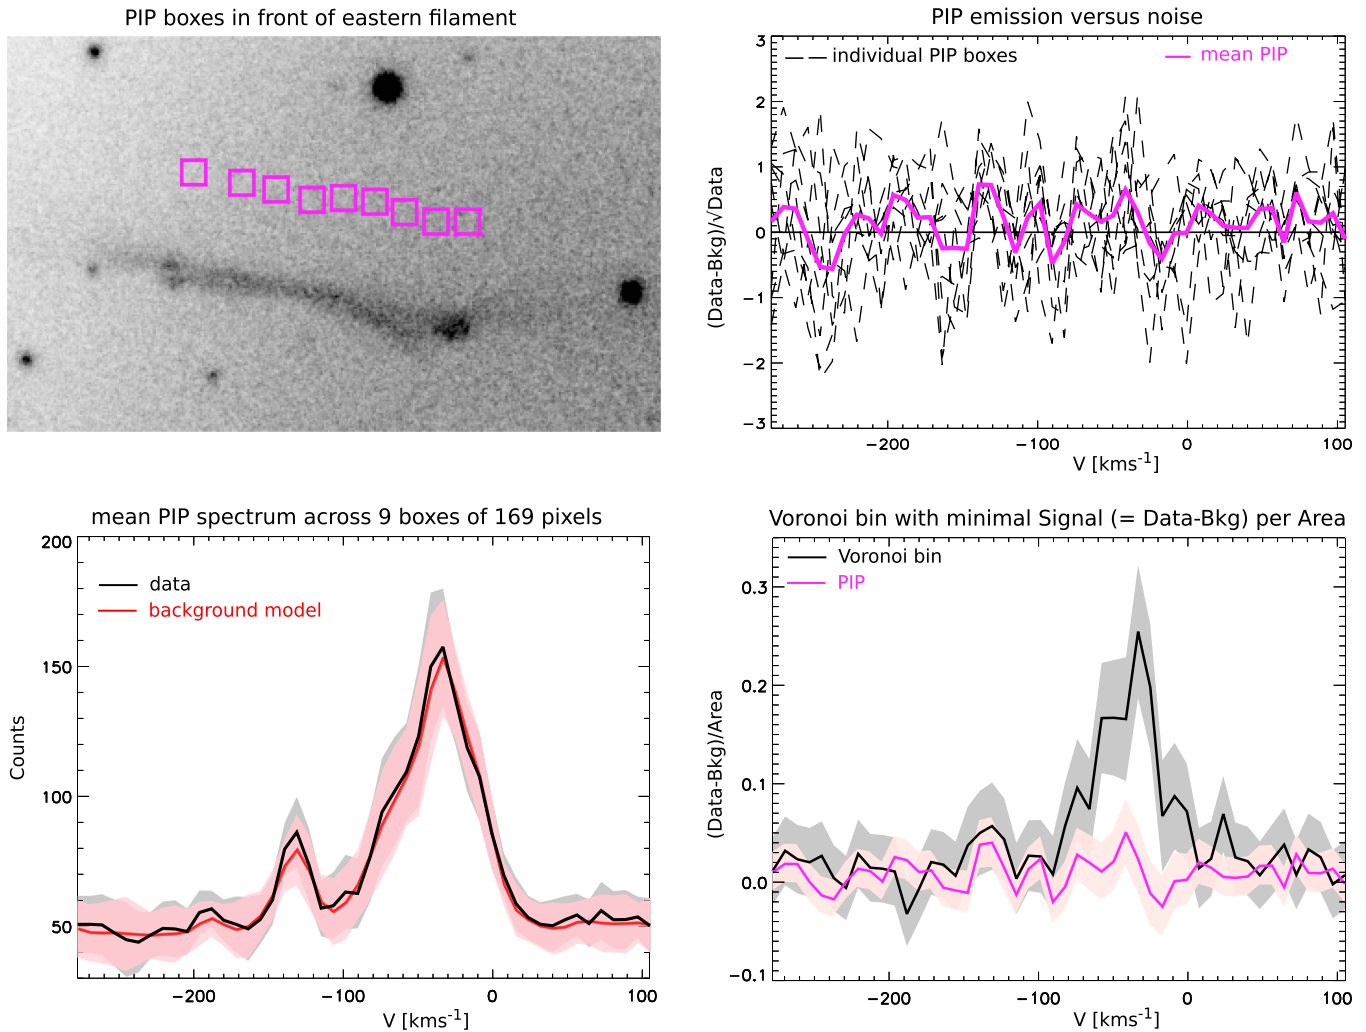
Figure 12. Magenta boxes in the top-left panel indicate the region where we extracted spectra in search of a potential PIP signal. The difference between the data and background model in individual pre-shock boxes (black-dashed) and mean PIP level (solid-magenta), divided by the measurement errors, shows that the residual emission is consistent with Poisson noise. The mean flux (solid black) and the background model (solid red) in the regions of putative PIP emission are directly compared in the bottom-left panel, with their 1σ uncertainties indicated with the corresponding color-shaded regions. The putative PIP flux and the background are entirely consistent with one another. In the bottom-right panel, we compare the filament emission (black line and gray-shaded region) in the bin with the smallest SB to the PIP region signal (magenta line and pink-shaded region), again illustrating that any possible PIP signal is negligible.

Appendix B
Photoionization Precursor (PIP)

When gas starts to cool and recombine downstream, the photons produced escape to the pre-shock region and form a PIP (Raymond 1979). Although non-radiative shocks lack recombination zones, PIP can still be created, where the main sources of the photons produced downstream are He I $\lambda584$ Å and He II $\lambda304$ Å. Ghavamian et al. (2000) reported not only on diffuse Hα, but also [N II] and [S II] emission extending over $1'$ in front of Tycho's NE rim, and suggested that it arises in a PIP. They predicted that the pre-shock gas was heated in the PIP to $\sim12,000$ K. Subsequently, Lee et al. (2007) measured the PIP spectrum in front of "knot g": narrow Hα with $W_{NL} \approx 34$ km s^{-1}, and [N II] $\lambda6583$ Å with the width of 23 km s^{-1}.

The similarity of the spectra in the pre- and post-shock background boxes (middle panel in Figure 10), where pre-shock boxes partially cover the region of the suspected PIP, unambiguously shows that the PIP emission in our data is

negligible. Furthermore, since the signal of a PIP increases toward the shock front, we searched for its signature in nine pre-shock, 13^2 pixel regions (magenta boxes in Figure 12) that were taken to be closer to the filament than the background boxes ($\approx7''$ away from the filament), but still far enough so that we do not pick up on projected filament emission and emission in the CR precursor. The top-right panel shows the comparison of the putative PIP signal and background model, normalized by measurement noise. The mean normalized PIP level is indicated by the solid -magenta line and is 0.13 on average. Therefore, the possible PIP signal is consistent with zero. We also show the comparison between the mean PIP and background flux, including their 1σ uncertainties (bottom-left panel). Finally, in the bottom-right panel we demonstrate that the PIP signal is negligible in comparison to the filament flux. We chose here the location (bin) with the smallest signal per area; the signal is larger in all other bins, and the putative PIP even smaller in relation to the signal. We therefore conclude that there is no or negligible PIP contribution to our filament flux models.
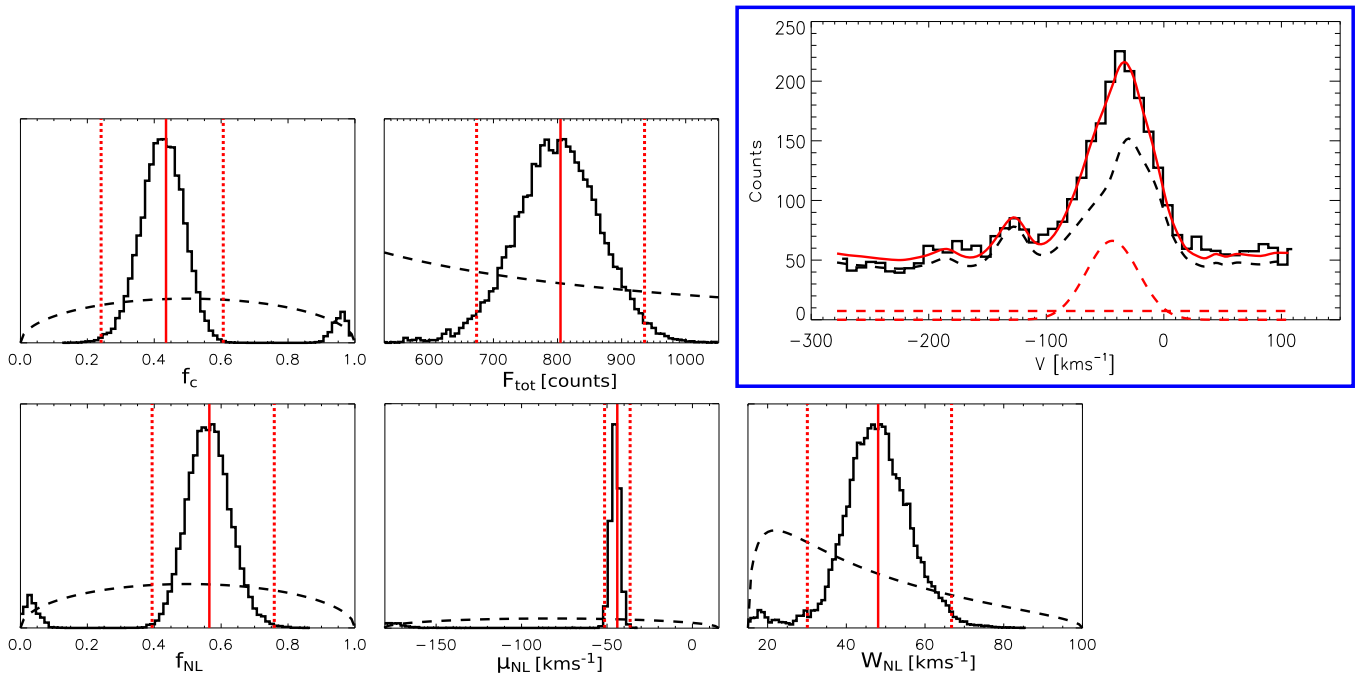
Figure 13. Parameter estimation of an NL model via Bayesian inference for the bin in the NE filament of Tycho's SNR for which we presented the posterior of the NLIL model parameters in Figure 2 (see its caption for explanation).

Appendix C
Spatial Binning of the Data

The Voronoi binning (Cappellari & Copin 2003) was performed setting two criteria: the targeted S/N and the minimum bin size. The minimum bin size has to be set in order to account for the seeing of $\simeq 1''$, which for the spatial scale of $0''.2$ pixel^{-1} gives the minimal bin size of $\simeq 19$ pixels. We set the targeted S/N = 10. The code of Cappellari & Copin (2003) in its standard version tends to create elongated-shaped bins in the direction perpendicular to the shock filament. The *HST* image of Tycho's SNR seen in Figures 1 and 2 in Lee et al. (2010) shows a complex filamentary structure including several very bright knots as a result of different shock emission projected along the LOS. Small-scale differential gradients are very noticeable in the direction of the shock normal, which is the reason why we used the Weighted Voronoi Tessellation adaptation of Cappellari & Copin (2003; Diehl & Statler 2006) that created a couple of rounder bins in the same direction instead. This way, we created 85 and 15 spatial Voronoi bins in the eastern and northern parts of the rim (Figure 3), respectively. Although we see only 2% level variations between boxes of 25^2 pixel size, these residual background variations become important relative to the S/N for bin sizes of 400 pixels and larger. This yields the upper size limit of 400 pixels for the bins used in our analysis. Seventy-three bins (out of 85) in the eastern and nine bins (out of 15) in the northern filament fulfill the above criteria and are further considered while the bins with more than 400 pixels are excluded from the analysis.

Appendix D
Model Parameters and Prior PDFs

As already mentioned in Section 3.2, we define the total flux, continuum and line flux fractions, line centroids, and widths as model parameters. In the case of the NLNL and NLNLIL models, we define the NL centroid mean $\langle \mu_{NL} \rangle$ and the separation between the two NLs $\Delta \mu_{NL}$ as parameters. The IL centroid is introduced with its offset from the NL centroid (mean) $\Delta \mu_{IL}$.

The parameters that we actually sample from are slightly different from the parameters that we use in Section 3; the difference enables the direct application of the prior PDFs (Dirichlet or Beta distributions) to the parameters. These prior distributions require parameter sets defined in the range (0, 1), except for the naturally based logarithm of the total flux $\ln(F_{tot})$ for which we use a flat unbound prior. The continuum and line flux fractions are by definition in the (0, 1) range, where we set the continuum flux fraction to be dependent on the flux fractions in the lines $f_c = 1 - \sum f_i$. The fact that the flux fractions sum up to 1 and that they are in the range (0, 1) makes the Dirichlet distribution $\prod_i f_i^{\alpha_i - 1}$, $i = [NL, NL1, NL2, IL, c]$ the perfect choice for their prior PDF. Since we do not favor any of the flux components, we use a symmetric Dirichlet distribution with the same index α which we set to $\alpha = 1.5$ to disfavor zero fluxes.

The line parameters μ_{NL}, $\langle \mu_{NL} \rangle$, $\Delta \mu_{NL}$, and $\Delta \mu_{IL}$ are all defined with the following functional form: $x' = (x - x_{min})/(x_{max} - x_{min})$. μ_{NL} or $\langle \mu_{NL} \rangle$ are defined in the range $[V_{cen} - V_{FSR}/4, V_{cen} + V_{FSR}/4]$, where V_{cen} is the center of the free velocity range V_{FSR} (FSR in velocity units) and $\Delta \mu_{NL}$ is in the range $[0, V_{FSR}/2]$ so that in the most extreme case the two NL centroids are at the edges of the spectral coverage. $\Delta \mu_{IL}$ is within $\pm V_{FSR}/4$, having the absolute upper boundary set to the upper (lower) boundary of the NL (IL) width ≈ 100 km s^{-1}. Instead of line widths W_{NL} and W_{IL} defined in the range [15, 100] km s^{-1} and [100, 350] km s^{-1}, respectively (Morlino et al. 2012, 2013), we used their log-widths (natural logarithm) denoted as $w_{NL(IL)} = \ln W_{NL(IL)}$ in the same functional form as for the centroids and separations. For all
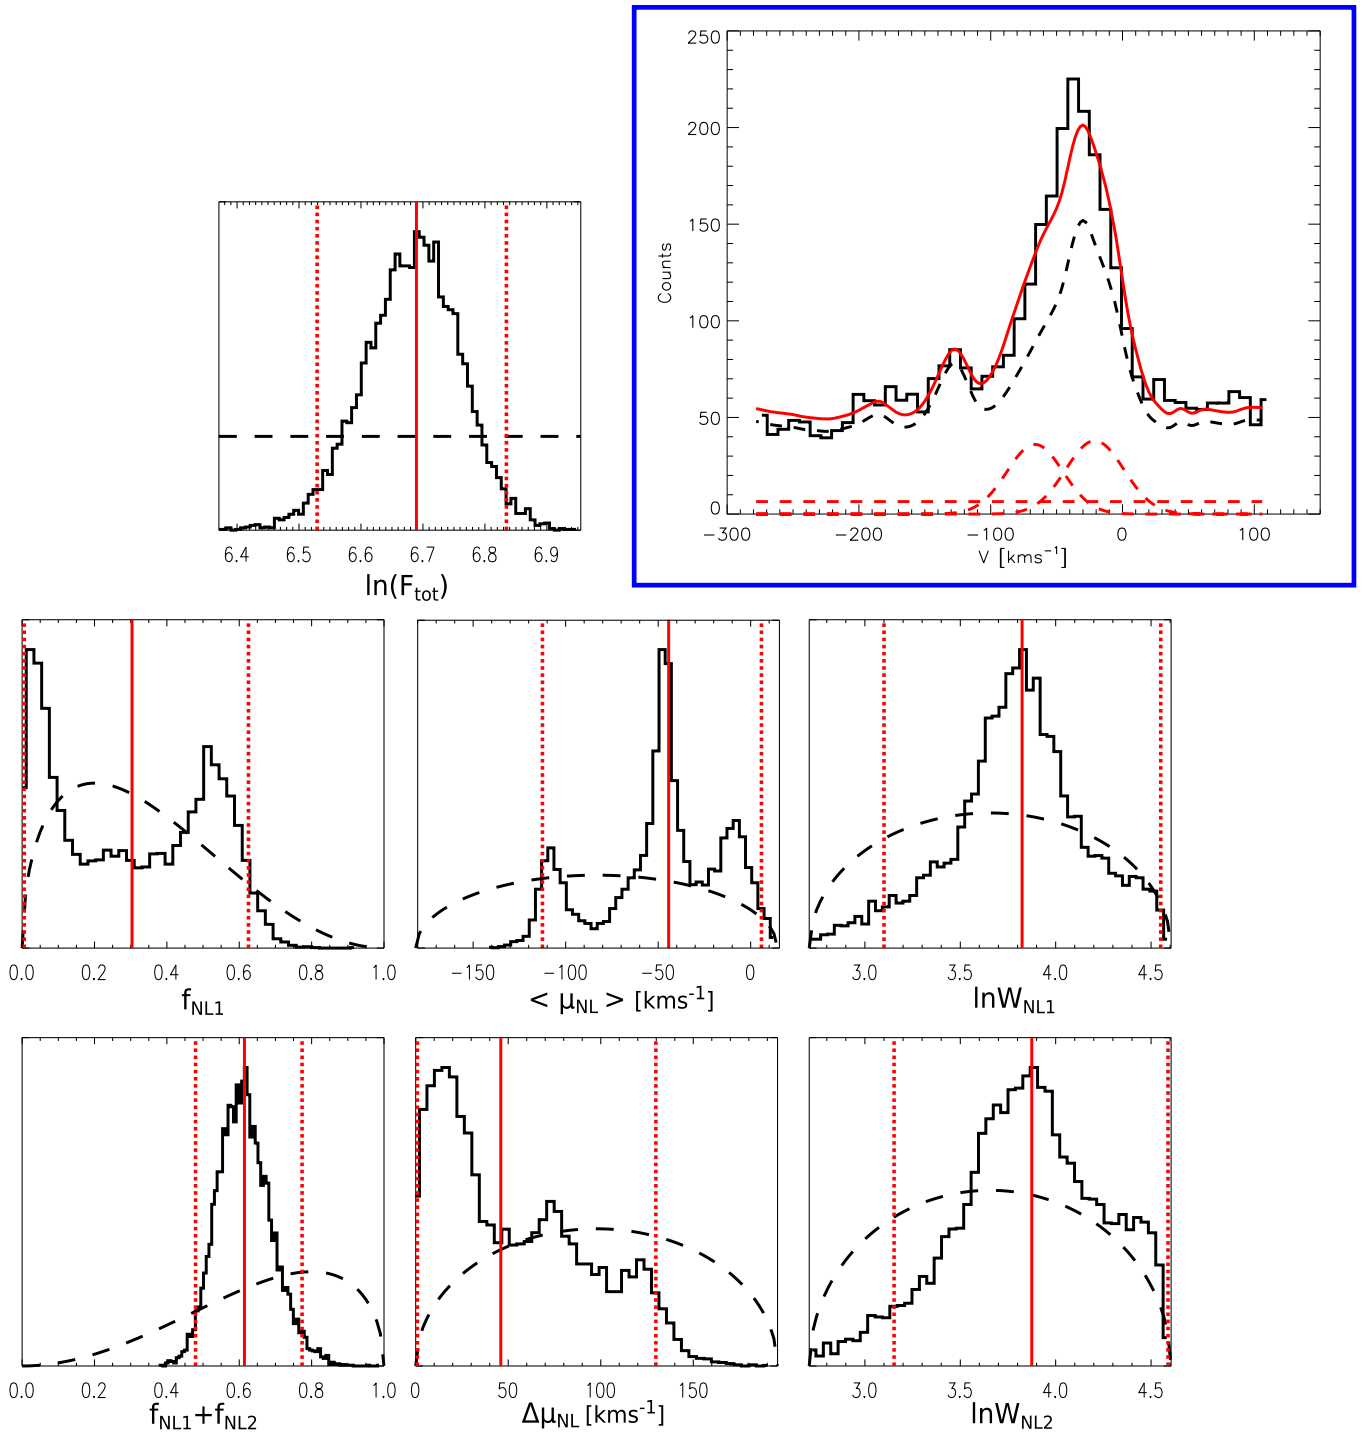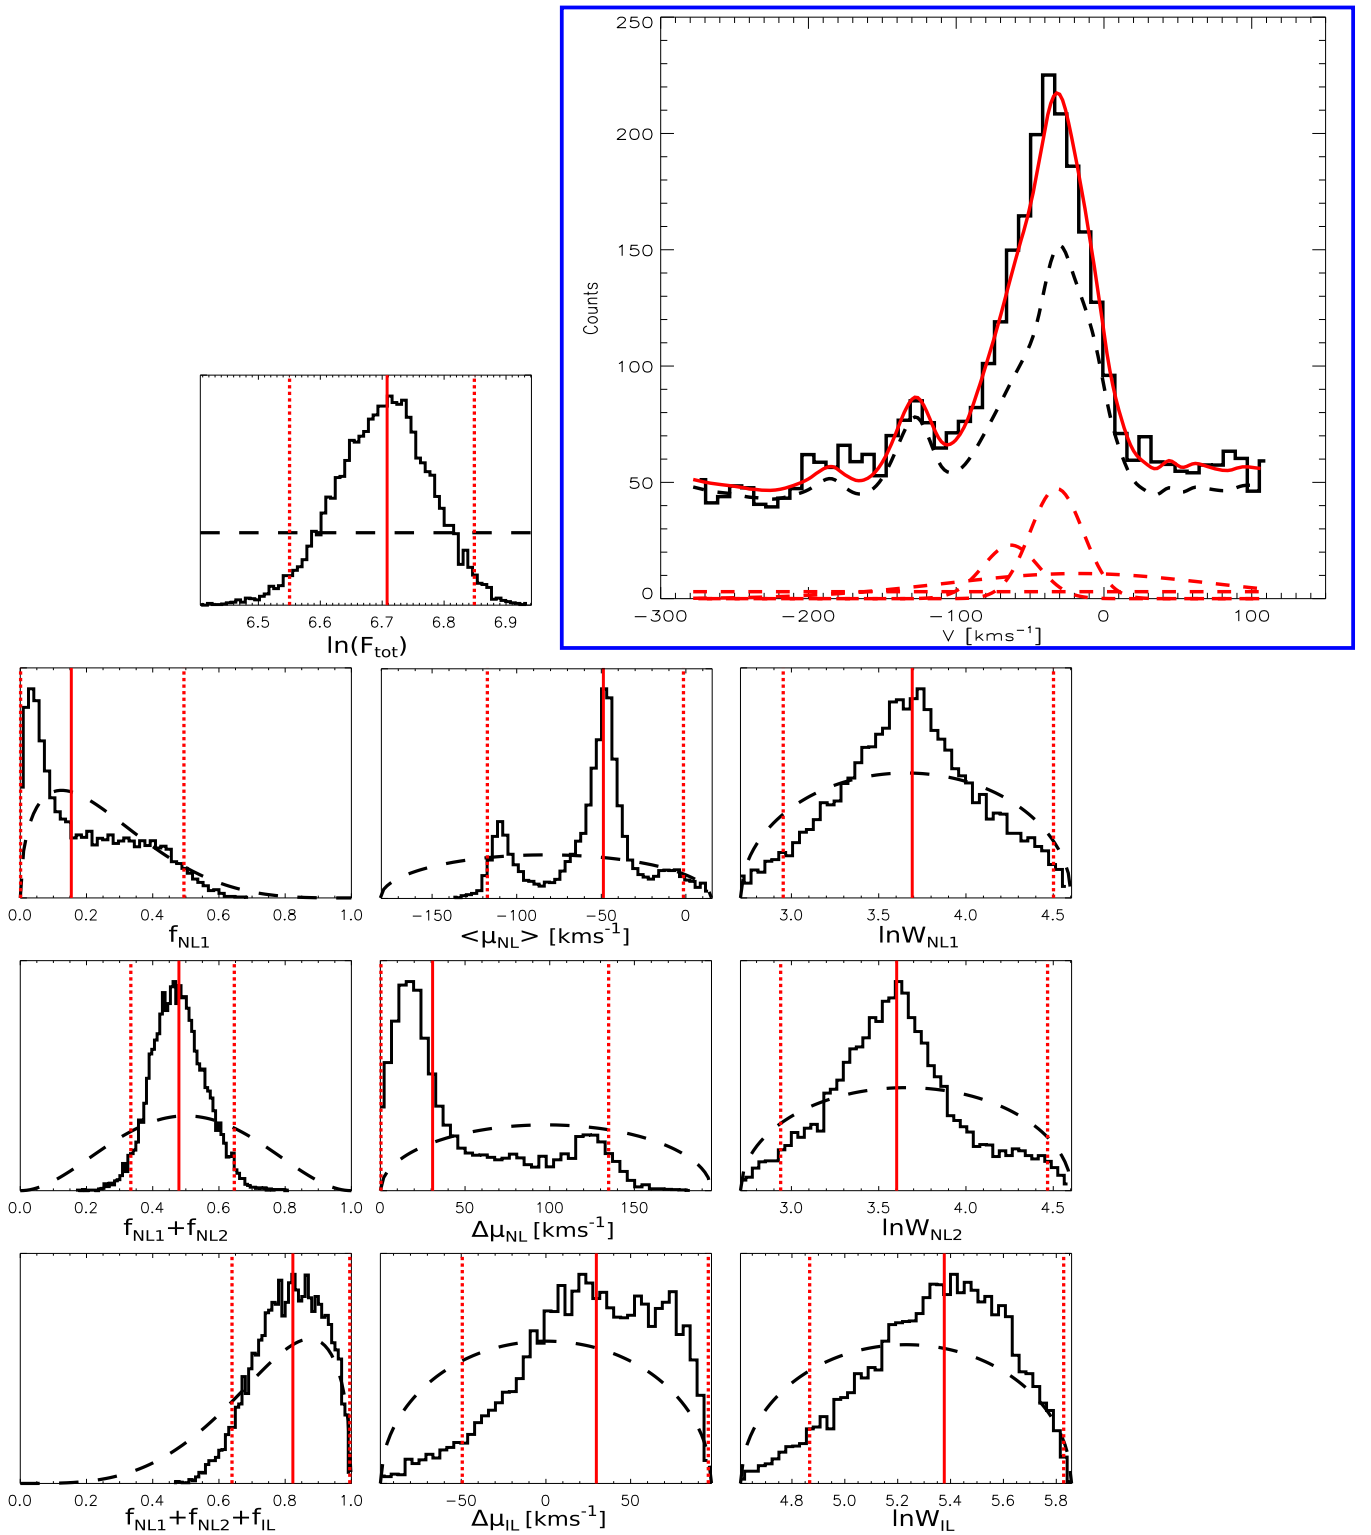
Figure 14. Parameter estimation of an NLNL model via Bayesian inference for the bin in the NE filament of Tycho's SNR for which we presented the posterior of the NLIL model parameters in Figure 2 (see its caption for an explanation). Instead of $F_{\rm tot}$, $W_{\rm NL1}$, $W_{\rm NL2}$, $f_{\rm NL1}$, and $f_{\rm NL2}$, we present the log (natural logarithm) of the total intrinsic flux and intrinsic line widths (ln $(F_{\rm tot})$, ln$W_{\rm NL1}$, ln$W_{\rm NL2}$), and cumulative flux fractions ($f_{\rm NL1}$, $f_{\rm NL1} + f_{\rm NL2}$).

of these parameters, we define a symmetric Beta distribution prior $x^{\alpha-1}(1-x)^{\beta-1}$ with the arguments $\alpha = \beta = 1.5$—slightly favoring the central values of the defined parameter ranges. The model parameters and their priors are summarized in the Table 1.

In addition to Figure 2, where we plotted posteriors for the favored NLIL model for one of the bins, in Figures 13–15 we show the posteriors for the NL, NLNL, and NLNLIL models,

where the latter two figures show the parameters that we actually sample from.

Appendix E
MCMC Sampling

Posterior samples were drawn from an ensemble MCMC sampler (Goodman & Weare 2010), an implementation of

25

Figure 15. Parameter estimation of an NLNLIL model via Bayesian inference for the bin in the NE filament of Tycho's SNR for which we presented the posterior of the NLIL model parameters in Figure 2 (see its caption for an explanation). Instead of F_{tot}, W_{NL1}, W_{NL2}, W_{IL}, f_{NL1}, f_{NL2}, and f_{IL}, we present the log (natural logarithm) of the total intrinsic flux and intrinsic line widths ($\ln(F_{tot})$, $\ln W_{NL1}$, $\ln W_{NL2}$, $\ln W_{IL}$), and cumulative flux fractions (f_{NL1}, $f_{NL1}+f_{NL2}$, $f_{NL1}+f_{NL2}+f_{IL}$).

which has been popularized as "emcee" (Foreman-Mackey et al. 2013). Among other advantages, this method provides for (near-) optimal tuning at every stage of the sampling, which would otherwise be a substantial challenge and obstacle in the way of efficient sampling considering 82 different data sets,

four different models for each of them, up to 10 model parameters, and our intent to test the sampler and results based on hundreds of additional simulated data sets. We draw the initial parameters for 128 parallel chains (walkers) uniformly between the prior boundaries. The unnormalized log-posterior

of a model is computed as a sum of unnormalized log-prior and log-likelihood for the proposed walker position. Since our data (fluxes in the spectral bins) result from Poisson processes, the likelihood is the product of each datum's probability under a Poisson distribution with the expectation value equal to the model prediction. After taking the logarithm,

$$\ln L = \sum_{i=1}^{N=48} d_i \ln(m_i) - m_i - \ln(\Gamma(d_i + 1)), \qquad (15)$$

where m_i are the model predictions at each spectral bin i, and d_i are the corresponding data. The last term in Equation (15) (the factorial term) was left out of the posterior sampling (but not the evidence calculation) because it is model independent. We refer to Foreman-Mackey et al. (2013) and Goodman & Weare (2010) for details of the sampling algorithm. Before checking the chains' convergence, we disregard the first 25%, but at least 512, of the samples of each walker ("burn-in"), and further thin the chains until the autocorrelation time of the thinned sample is smaller that 5. In order to achieve low noise and to set a first minimum threshold for the convergence of the chains, we require at least $2^{13} = 8096$ total samples (all walkers combined) to be kept after thinning. Once this minimum number of samples is reached, we additionally impose the following convergence (stopping) criterion: we split the sample into subsamples and compute the desired estimators (maximum posterior sample, median, and 95% confidence interval boundaries) for each subsample. The variances of the subsample estimators are then required to be smaller than 5% of the mean parameter value. In order to reduce the probability of coincidentally favorable (small) variances from possible "modes" in the chains, we repeat the process twice, first taking each of the 128 chains as one subsample, and second, taking each of the 128 element walker states as subsamples, and use the arithmetic mean of the resulting eight relative estimator standard deviations toward our 5% criterion.

To ensure that the applied procedure gives the correct results, we performed tests of the posterior sampling routine, using simulated data with known model parameters. We tested models with S/N in the range [5, 50] and varying background levels (0%, 50%, and 90% of the total flux). First, we checked if the posterior distribution reproduces the prior in the $S/N = 0$ limit, but also if the posterior approached a delta function in the infinity limit, i.e., for very large S/N.

Second, we check if the model parameters of the input model are reproduced statistically. We ran the algorithm for 200 different realizations of the same model, i.e., each time drawing the data from the same model prediction with its specified uncertainty included (each a set of 48 Poisson distributions). The mean of the distribution of the median values is always consistent with the input model parameter and the typical scatter of this mean is 10%–20% for the range of S/N in our real data.

Finally, we vary model parameters by randomly choosing them 200 times from within the prior boundaries. Again, for each of the resulting simulated data, we sample the posterior as described above and evaluate it in the form of the median of the marginalized posteriors. We find that the measured median

values scatter symmetrically around the 1:1 relation with the input model parameters, with the scatter being roughly equal to the individual posteriors' standard deviation, as desired. The distribution of the measured values becomes biased as the input parameters approach the parameter range boundaries, as expected for our priors.

Appendix F
Evidence Calculation via the LOO-CV Likelihood

We use the CV likelihood, specifically its "LOO" variant, to compute model evidences and to compare models (Bailer-Jones 2012). We prefer it over the standard numeric ("Bayesian") integral, because it draws samples from the posterior instead of the prior. It is hence more efficient, less dependent on the choice of prior, and in some cases numerically more stable. The idea of the CV likelihood is to evaluate the likelihood of part of the data, given the model and the rest of the data. In our case, we have 48 data points and measure how well any 47 data points (the complement, D_{-k}) under the model M predict the 48th data point (the partition, D_k), as quantified by the complement's posterior $P(\theta|D_{-k}, M)$ and its prediction for D_k. The process is repeated for all possible (48) partitions. Each time we leave out one datum (D_k), its partition likelihood L_k is given by

$$
\begin{aligned}
L_k &= P(D_k|D_{-k}, M) \\
&= \int_\theta P(D_k|\theta, M)P(\theta|D_{-k}, M)d\theta. \qquad (16)
\end{aligned}
$$

The first term in the integrand is the likelihood of D_k, while the second term is the posterior PDF after considering the information contained in D_{-k}. We can numerically (Monte Carlo) integrate by drawing a number of samples N from $P(\theta|D_{-k}, M)$:

$$L_k \approx \frac{1}{N} \sum_{n=1}^{N} P(D_k|\theta_n, M). \qquad (17)$$

Assuming that the data points are independent, the LOO-CV likelihood is the product of all partition likelihoods, $L_{\text{LOO-CV}} = \prod_{k=1}^{48} L_k$, or

$$\ln L_{\text{LOO-CV}} = \sum_{k=1}^{48} \ln L_k. \qquad (18)$$

It can be shown that $L_{\text{LOO-CV}}$ for model M is equal to the Bayesian evidence ($E(M)$), and we henceforth use it to compute the Bayes factors, $E(M_1)/E(M_2)$, to compare models.

We tested the ability of the Bayes factors to indicate the correct model by employing the simulated data. We find that the typical numerical precision of $\ln L_{\text{LOO-CV}}$ is better than 0.05 dex. At the same time, for data generated from the parameters and with S/Ns that are typical for the actual data, the "right" model's $\ln L_{\text{LOO-CV}}$ is 0.2 dex ($\gtrsim 50\%$) better than any of the alternative models. Nevertheless, on heuristic grounds (what probability is considered statistically "significant enough"?) and in order to bracket the practically limited scope of such tests, we adopt a more conservative $+0.5$ dex threshold
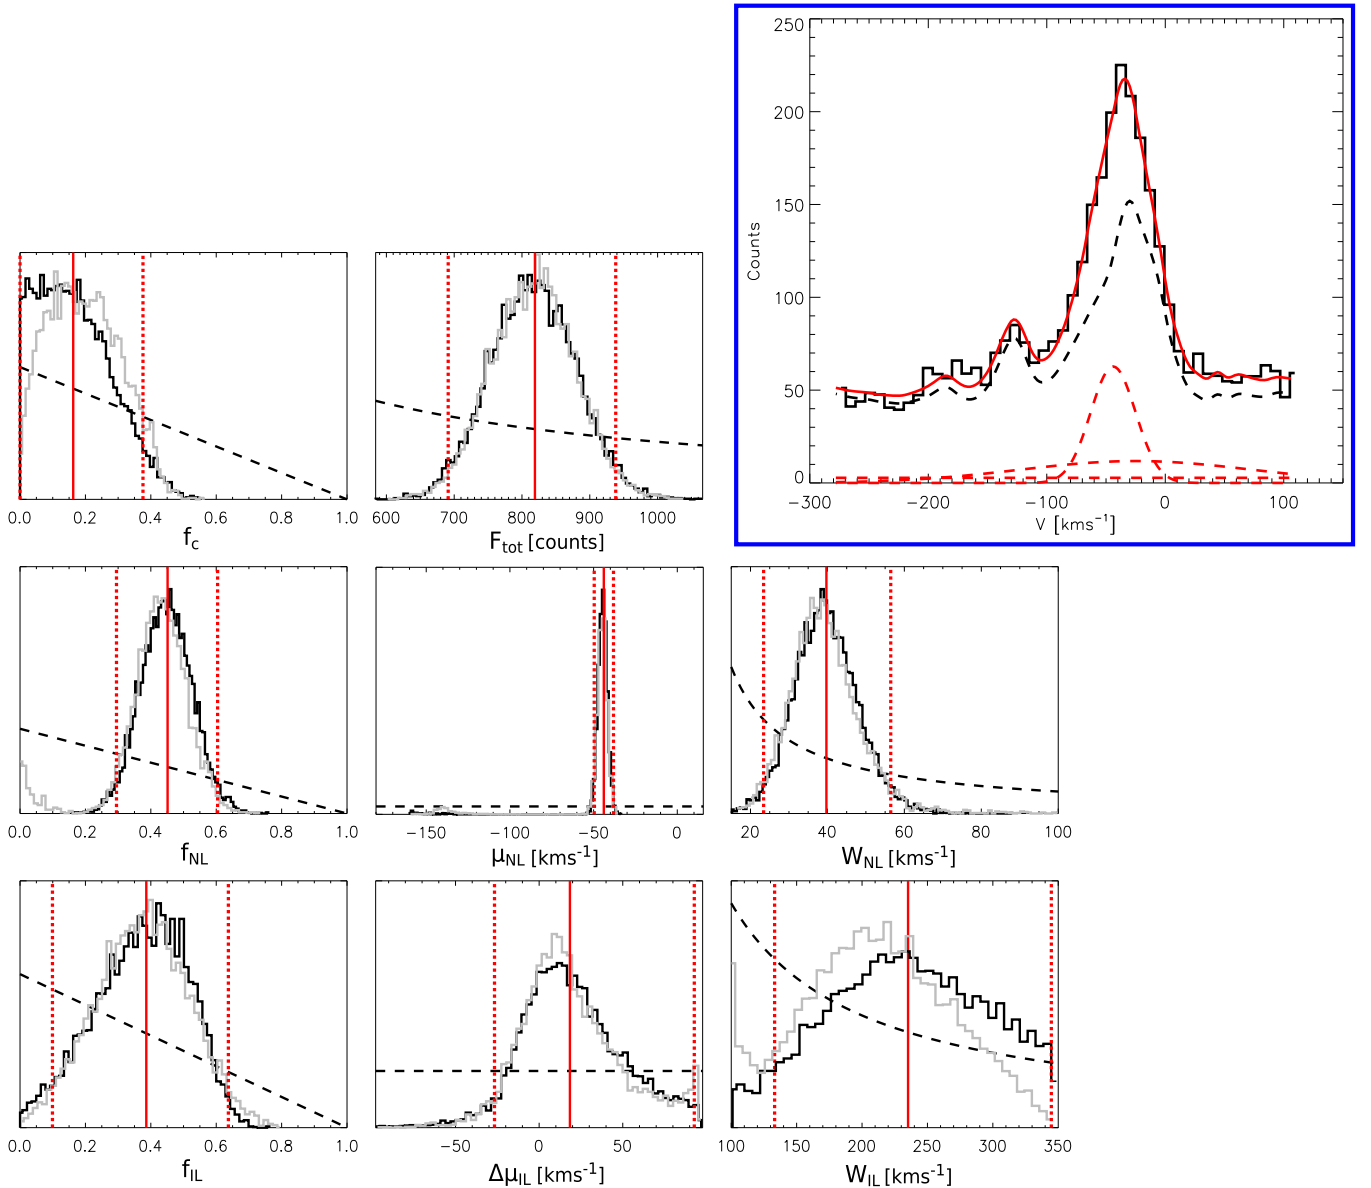
Figure 16. Parameter estimation of an NLIL model via Bayesian inference for the bin in the NE filament of Tycho's SNR for which we presented the posterior in Figure 2 (see its caption for an explanation). Posteriors in black are calculated for flat Dirichlet and Beta priors for all parameters. Posteriors from Figure 2 are overplotted in gray. Among all parameters, the W_{IL} posterior is the most sensitive to the prior choice. Even so, the median and 95% confidence intervals agree to within $\approx$10%, and its basic shape remains robust.

(3 : 1 probability) before we consider a model to be preferable over another.

We have also tested the dependence of the evidence ratios on the choice of prior. We used the same functional form for the priors (Table 1), but different α, β distribution parameters. Specifically, we tested α_D values for the Dirichlet distribution, and $\alpha_B = \beta_B$ for the Beta distribution with values of $(\alpha_D, \alpha_B) = \{(1.5, 1.5), (1, 1.5), (1.5, 1), (1, 1)\}$, where $(1.5, 1.5)$ was used for the results presented in the main part of the paper. We found that for all bins, the mean standard deviation of the appropriate evidence ratios is $\lesssim 0.07$ dex.

In addition, in Figures 16 and 17 we show the posteriors, medians, and 95% confidence intervals when flat priors, i.e., $(\alpha_D, \alpha_B) = (1, 1)$ instead of the fiducial $(1.5, 1.5)$, are applied.

We present results for the same two Voronoi bins shown in Figures 2 and 5. The shape of the posteriors and their medians are very similar to those obtained with the adopted priors, demonstrating that our results do not depend strongly on the choice of prior. W_{IL} posteriors for the adopted and flat prior are different, but similar to the degree that the median and confidence interval boundaries change only by 30 km s^{-1} ($\approx$10%). Even with the flat prior, 210 km s^{-1} is clearly more probable than the other W_{IL} parameter values; in particular, it is preferred over values close to the W_{IL} limits. This shows that the preference for the central W_{IL} values is not just borne out of the prior shape or range, but genuinely reflects constraints provided by the data, even if they are not as strong as for other parameters.
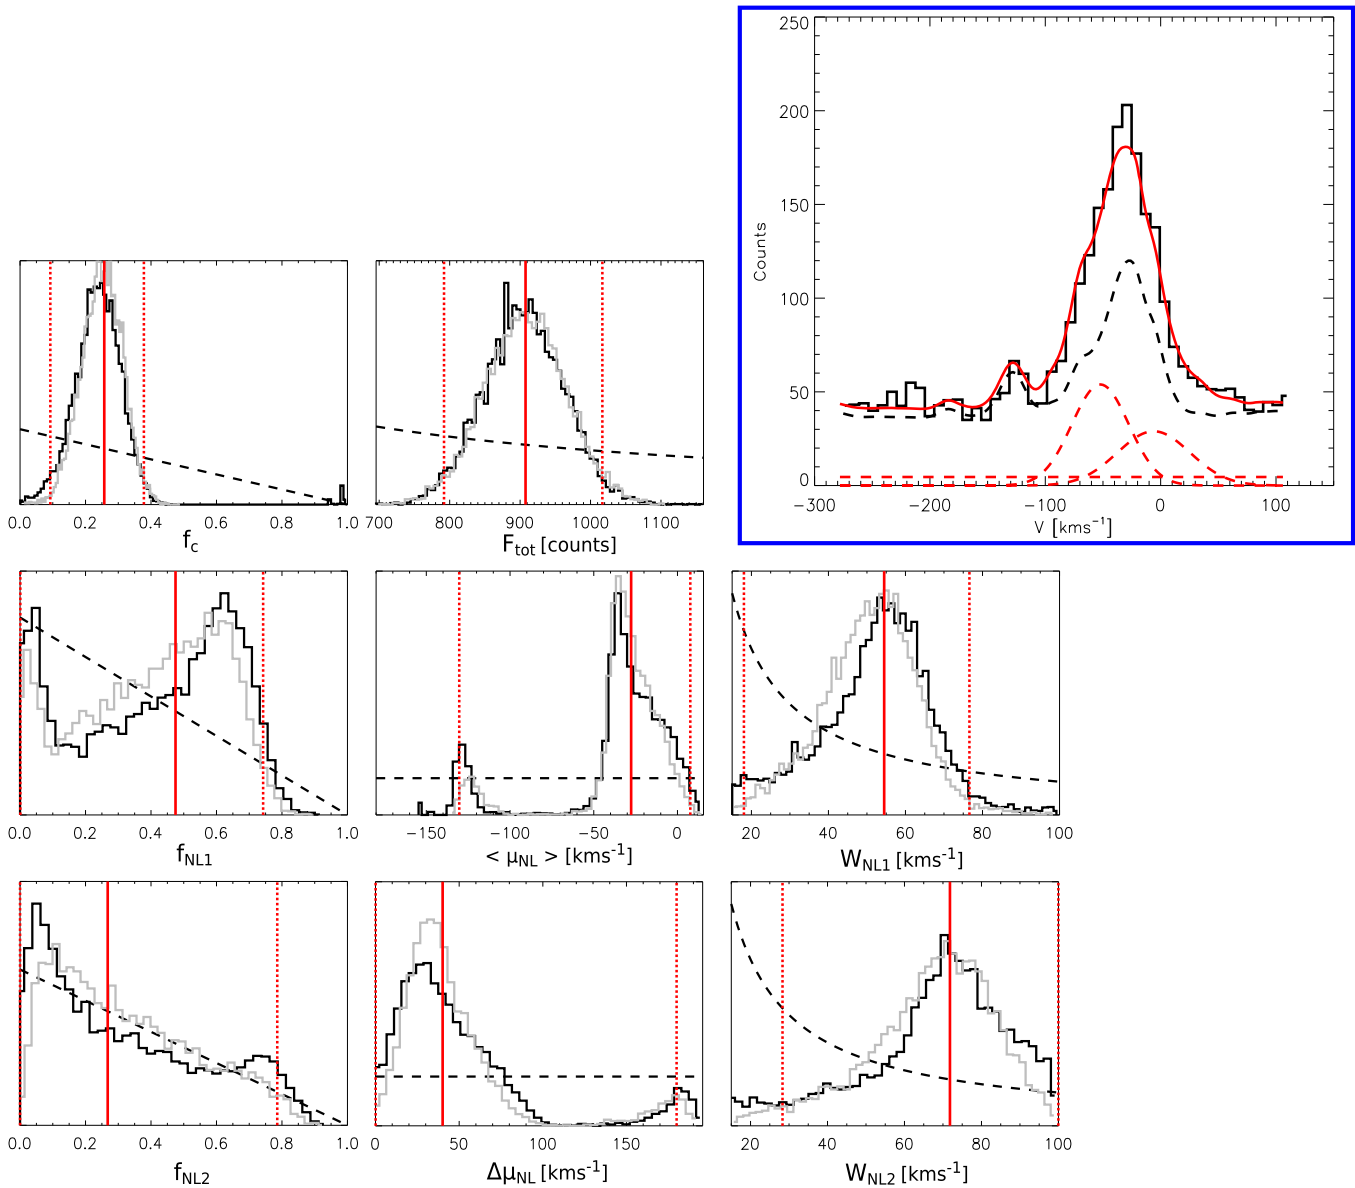
Figure 17. Parameter estimation of an NLNL model via Bayesian inference for the bin in the NE filament of Tycho's SNR for which we presented the posterior in Figure 5 (see its caption for an explanation). Posteriors in black are calculated for flat Dirichlet and Beta priors for all parameters. Posteriors from Figure 5 are overplotted in gray.

ORCID iDs

Sladjana Knežević https://orcid.org/0000-0003-1416-8069
John C. Raymond https://orcid.org/0000-0002-7868-1622
Giovanni Morlino https://orcid.org/0000-0002-5014-4817
Parviz Ghavamian https://orcid.org/0000-0002-9886-0839
Kevin Heng https://orcid.org/0000-0003-1907-5910

References

Acciari, V. A., Aliu, E., Arlen, T., et al. 2011, ApJL, 730, L20
Ackermann, M., Ajello, M., Allafort, A., et al. 2013, Sci, 339, 807
Bailer-Jones, C. A. L. 2012, A&A, 546, A89
Bamba, A., Yamazaki, R., Yoshida, T., Terasawa, T., & Koyama, K. 2005, ApJ, 621, 793
Blasco-Herrera, J., Fathi, K., Beckman, J., et al. 2010, MNRAS, 407, 2519
Blasi, P., Morlino, G., Bandiera, R., Amato, E., & Caprioli, D. 2012, ApJ, 755, 121
Cappellari, M., & Copin, Y. 2003, MNRAS, 342, 345
Caprioli, D. 2015, https://arxiv.org/pdf/1510.07042.pdf
Cassam-Chenaï, G., Hughes, J. P., Reynoso, E. M., Badenes, C., & Moffett, D. 2008, ApJ, 680, 1180
Chevalier, R. A., Kirshner, R. P., & Raymond, J. C. 1980, ApJ, 235, 186
Diehl, S., & Statler, T. S. 2006, MNRAS, 368, 497
Eriksen, K. A., Hughes, J. P., Badenes, C., et al. 2011, ApJL, 728, L28
Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306
Ghavamian, P., Raymond, J. C., Hartigan, P., & Blair, W. P. 2000, ApJ, 535, 266
Ghavamian, P., Raymond, J. C., Hartigan, P., & Blair, W. P. 2017, ApJ, 843, 77
Ghavamian, P., Raymond, J. C., Smith, R. C., & Hartigan, P. 2001, ApJ, 547, 995
Ghavamian, P., Schwartz, S. J., Mitchell, J., Masters, A., & Laming, J. M. 2013, SSRv, 178, 633
Giordano, F., Naumann-Godo, M., Ballet, J., et al. 2012, ApJL, 744, L2
Goodman, J., & Weare, J. 2010, Comm. App. Math. Comp. Sci, 5, 65
Hayato, A., Yamaguchi, H., Tamagawa, T., et al. 2010, ApJ, 725, 894
Helder, E. A., Kosenko, D., & Vink, J. 2010, ApJL, 719, L140
Helder, E. A., Vink, J., Bassa, C. G., et al. 2009, Sci, 325, 719
Helder, E. A., Vink, J., Bykov, A. M., Ohira, Y., Raymond, J. C., & Terrier, R. 2012, SSRv, 173, 369

Heng, K. 2010, PASA, 27, 23
Hernandez, O., Fathi, K., Carignan, C., et al. 2008, PASP, 120, 665
Hester, J. J., Raymond, J. C., & Blair, W. B. 1994, ApJ, 420, 721
Katsuda, S., Petre, R., Hughes, J. P., et al. 2010, ApJ, 709, 1387
Lee, J.-J., Koo, B.-C., Raymond, J. C., et al. 2007, ApJL, 659, L133
Lee, J.-J., Koo, B.-C., & Tatematsu, K. 2004, ApJL, 605, L113
Lee, J. J., Raymond, J. C., Park, S., & Blair, W. P. 2010, ApJL, 715, L146
Morlino, G., Bandiera, R., Blasi, P., & Amato, E. 2012, ApJ, 760, 137
Morlino, G., & Blasi, P. 2016, A&A, 589, A7
Morlino, G., Blasi, P., Bandiera, R., Amato, E., & Caprioli, D. 2013, ApJ, 768, 148
Morlino, G., & Caprioli, D. 2012, A&A, 538, A81
Nikolić, S., Van de Ven, G., Heng, K., et al. 2013, Sci, 340, 45
Park, N. & for the VERITAS Collaboration 2015, arXiv:1508.07068

Raymond, J. C. 1979, ApJS, 39, 1
Reynolds, S. P., & Ellison, D. C. 1992, ApJL, 399, L75
Reynolds, S. P., & Keohane, J. W. 1999, ApJ, 525, 368
Reynoso, E. M., Velázquez, P. F., Dubner, G. M., & Goss, W. M. 1999, ApJ, 117, 1827
Slane, P., Lee, S.-H., Ellison, D. C., et al. 2014, ApJ, 783, 33
Smith, R. C., Raymond, J. C., & Laming, J. M. 1994, ApJ, 420, 286
Stroman, W., & Pohl, M. 2009, ApJ, 696, 1864
Tian, W. W., & Leahy, D. A. 2011, ApJL, 729, L15
Wagner, A. Y., Lee, J. J., Raymond, J. C., Hartquist, T. W., & Falle, S. A. E. G. 2009, ApJ, 690, 1412
Warren, J. S., Hughes, J. P., Badenes, C., et al. 2005, ApJ, 634, 376
Williams, B. J., Borkowski, K. J., Ghavamian, P., et al. 2013, ApJ, 770, 129
Zhou, P., Chen, Y., Zhang, Z.-Y., et al. 2016, ApJ, 826, 34