

Spica and the annual cycle of PKS B1322–110 scintillations

Hayley Bignall,¹★ Cormac Reynolds,¹ Jamie Stevens,² Keith Bannister,³
 Simon Johnston,³ Artem V. Tuntsov,⁴† Mark A. Walker,⁴ Sergei Gulyaev,⁵
 Tim Natusch,⁵ Stuart Weston,⁵ Noor Masdiana Md Said,⁶ Matthew Kratzer.⁷

¹CSIRO Astronomy and Space Science, 26 Dick Perry Avenue, Kensington, WA 6151, Australia

²CSIRO Paul Wild Observatory, 1828 Yarrie Lake Road, Narrabri, NSW 2390, Australia

³CSIRO Astronomy and Space Science, PO Box 76, Epping NSW 1710, Australia

⁴Manly Astrophysics, 15/41-42 East Esplanade, Manly, NSW 2095, Australia

⁵Auckland University of Technology, 120 Mayoral Drive, Auckland 1010, New Zealand

⁶University of Tasmania, Hobart, TAS 7001, Australia

⁷The University of Queensland, QLD 4072, Australia

Accepted 2019 May 31. Received 2019 May 06; in original form 2018 December 07

ABSTRACT

PKS B1322–110 is a radio quasar that is located only 8.5 in angular separation from the bright B star Spica. It exhibits intra-day variability in its flux density at GHz frequencies attributed to scintillations from plasma inhomogeneities. We have tracked the rate of scintillation of this source for over a year with the Australia Telescope Compact Array, recording a strong annual cycle that includes a near-standstill in August and another in December. The cycle is consistent with scattering by highly anisotropic plasma microstructure, and we fit our data to that model in order to determine the kinematic parameters of the plasma. Because of the low ecliptic latitude of PKS B1322–110, the orientation of the plasma microstructure is poorly constrained. Nonetheless at each possible orientation our data single out a narrow range of the corresponding velocity component, leading to a one-dimensional constraint in a two-dimensional parameter space. The constrained region is consistent with a published model in which the scattering material is associated with Spica and consists of filaments that are radially oriented around the star. This result has a 1% probability of arising by chance.

Key words: ISM: general – ISM: structure – radio continuum: galaxies – radio continuum: transients – circumstellar matter – stars: individual: Spica

1 INTRODUCTION

PKS B1322–110 is a flat spectrum radio source (Griffith *et al.* 1994) which was recently discovered to undergo extreme intra-day variations (IDV, also used to refer to the class of sources) at GHz frequencies (Bannister *et al.* 2017). Rapid flux density variations in these sources are scintillations resulting from plasma inhomogeneities in our Galaxy (Jauncey *et al.* 2016; Rickett 1990), and can be used to constrain the apparent size or brightness temperature of their radio emitting regions (Dennett-Thorpe & de Bruyn 2000; Kedziora-Chudczer *et al.* 1997).

Scintillation can be thought of as a spatial flux pattern – i.e. the source projected through the transparent plasma

“screen” – that drifts relative to the Earth. For sources at cosmological distances, the velocity of the pattern is essentially that of the screen (Cordes & Rickett 1998), and therefore the change in the velocity of the Earth as it orbits the Sun can strongly affect the variation timescales. This annual modulation has so far been reported in a handful of sources: J1819+3845 (Dennett-Thorpe & de Bruyn 2001, 2003); QSO B0917+624 (Jauncey & Macquart 2001; Rickett *et al.* 2001; Fuhrmann *et al.* 2002); PKS 1257–326 (Bignall *et al.* 2003); PKS B1519–273 (Jauncey *et al.* 2003); PKS B1622–253 (Carter *et al.* 2009); S5 0716+714 (Liu *et al.* 2012); 0925+504 (Liu & Liu 2015; Liu *et al.* 2017); S4 0954+65 (Marchili *et al.* 2012); 1156+295 (4C+29.45, Liu *et al.* 2013); and J1128+5925 (Gabányi *et al.* 2007b,a). Together with the two-station experiments (Dennett-Thorpe & de Bruyn 2002; Bignall *et al.* 2006), which directly demonstrated the existence of a spatial flux pattern, the annual

★ E-mail: Hayley.Bignall@csiro.au

† E-mail: Artem.Tuntsov@manlyastrophysics.org

cycles provided the key evidence that proved IDV to be scintillation. No clear annual cycle could be established for the prototypical intra-hour variable PKS 0405–385 due to the intermittency of its variations (Kedziora-Chudczer 2006).

Annual cycles of the two best-studied intra-hour variables, J1819+3845 and PKS 1257–326 have been shown to be consistent with highly anisotropic, essentially one-dimensional scattering (Walker et al. 2009), and the orientation of the respective anisotropy axes along with the projected velocity of the screen were determined. However, although the properties of the screens have been precisely characterised, the physical context of the scattering material remains unclear. In the case of J1819+3845, whose screen is expected to be relatively close, less than 10 pc from Earth, it was previously suggested that the plasma might be associated with Vega, a nearby A star that is close to the source in the sky (Dennett-Thorpe & de Bruyn 2002). Curiously, the anisotropy axis of J1819+3845 does point towards Vega.

The possibility of a connection between IDVs and hot stars was reinforced by the realisation that the new IDV PKS B1322–110 is just 8.5 arcminutes away from Spica, the Sun’s closest B star neighbour, prompting Walker et al. (2017) to examine the stars foreground to J1819+3845 and PKS B1257–326. A conclusion of that study was that the scintillations of both sources are due to plasma associated with nearby A stars — the star being Alhakim (ι Cen), in the case of PKS B1257–326. The picture that was suggested by Walker et al. (2017) is of plasma filaments that are radially oriented around the host star, and co-moving with it.

The possible connection between Spica and the IDV PKS B1322–110 was left out of that analysis. The reason for the omission is that the close alignment between Spica and PKS B1322–110 motivated the idea of an association, and therefore cannot be used as a test. On the other hand, at the time of writing of Walker et al. (2017), less than three months after discovering IDV in PKS B1322–110, its annual cycle was not established and the kinematics of the screen were unknown.

In this paper we report the results of tracking the rate of flux density variations in PKS B1322–110 for just over a year, from February 2017 to February 2018, in which an evolution in the scintillation timescale is clearly seen. We interpret this evolution in purely kinematic terms – i.e. we assume that it is an annual cycle arising from the Earth’s orbit – but with only one year of data we are unable to demonstrate the repetition that is expected for an annual cycle. In Section 2 we present the observations and data reduction. Section 3 describes our inference of variability rates; our method allows us to characterise the scintillation rate during slow phases of the cycle, where traditional methods of analysis struggle. Section 4 fits the kinematic parameters of the annual cycle to the data and compares the results to the predictions of the model that connects the scattering with Spica. We discuss the degeneracy in our constraints Section 5 before concluding in Section 6.

2 OBSERVATIONS AND DATA REDUCTION

We observed PKS B1322–110 with the Australia Telescope Compact Array (ATCA) at 18 epochs, taking between 30 and 110 spectra extending from approximately 4.3 GHz to

Table 1. Parameters of the 18 observational epochs on which long light curves were obtained. The right column shows the number of data points remaining in the $(5.5 \pm 0.25)/(10 \pm 0.25)$ GHz sub-bands, as used in Section 3, after RFI excision.

Epoch	Date	D.o.Y.	MJD (mean)	#points
1	2017/02/02	32	57785.70	104/103
2	2017/02/07	37	57790.69	75
3	2017/02/19	49	57802.77	30
4	2017/02/21	51	57804.72	73
5	2017/02/23	53	57806.68	90
6	2017/03/23	81	57834.65	105/104
7	2017/03/24	82	57835.65	102/103
8	2017/04/11	100	57853.67	75
9	2017/05/10	129	57882.52	105
10	2017/05/11	130	57883.51	109
11	2017/08/14	226	58979.25	39
12	2017/08/30	242	58995.15	70
13	2017/10/01	274	58027.10	109/108
14	2017/10/02	275	58028.11	0/0
15	2017/11/02	306	58059.02	105/103
16	2017/12/15	348	58101.90	107
17	2017/12/16	349	58102.90	90
18	2018/02/24	54(+365)	58172.72	46

11 GHz using two quasi-simultaneous tunings, of 150 s integration time each. The summary of these data is given in Table 1.

To form the light curves used in the variability rate analysis below, we first filtered outliers from each recorded spectrum in the sub-bands of interest, by discarding data points that deviated from the mean of the group of their 10 closest neighbours by more than 3 times the r.m.s. values of the group, repeating this procedure twice on the updated spectra. We then visually inspected the full dynamic spectra and dropped those remaining data points that were clearly affected by RFI or other instrumental issues. In particular, we have excluded the entire data set from 2017/10/02 due to a persistent low-amplitude ‘moire’ pattern in the dynamic spectrum; the origin of this pattern is unclear. Figure 1 presents the light curves of PKS B1322–110 observed at all 18 epochs, averaged over two 0.5 GHz-wide bands centred at 5.5 GHz and 10 GHz.

3 VARIABILITY RATE INFERENCE

Traditionally, IDV have been analysed by direct computation of the temporal auto-correlation function (ACF) (e.g., Bignall et al. 2003), or the structure function (e.g., Gabányi et al. 2007a), from the measured flux densities; and then extracting a characteristic timescale. However, either of these methods is difficult to use during slow phases in the cycle when the entire observing run might not be long enough to record a single oscillation in the light curve, leading to a biased estimate of the ACF. The epochs of slow variation are nevertheless particularly useful in constraining the kinematics of the screen, as we will see in §4.3.2, and we are therefore motivated to try a new approach to estimating the variability timescale.

We model the light curves as a realisation of a stationary Gaussian process with an epoch-dependent time axis scaling;

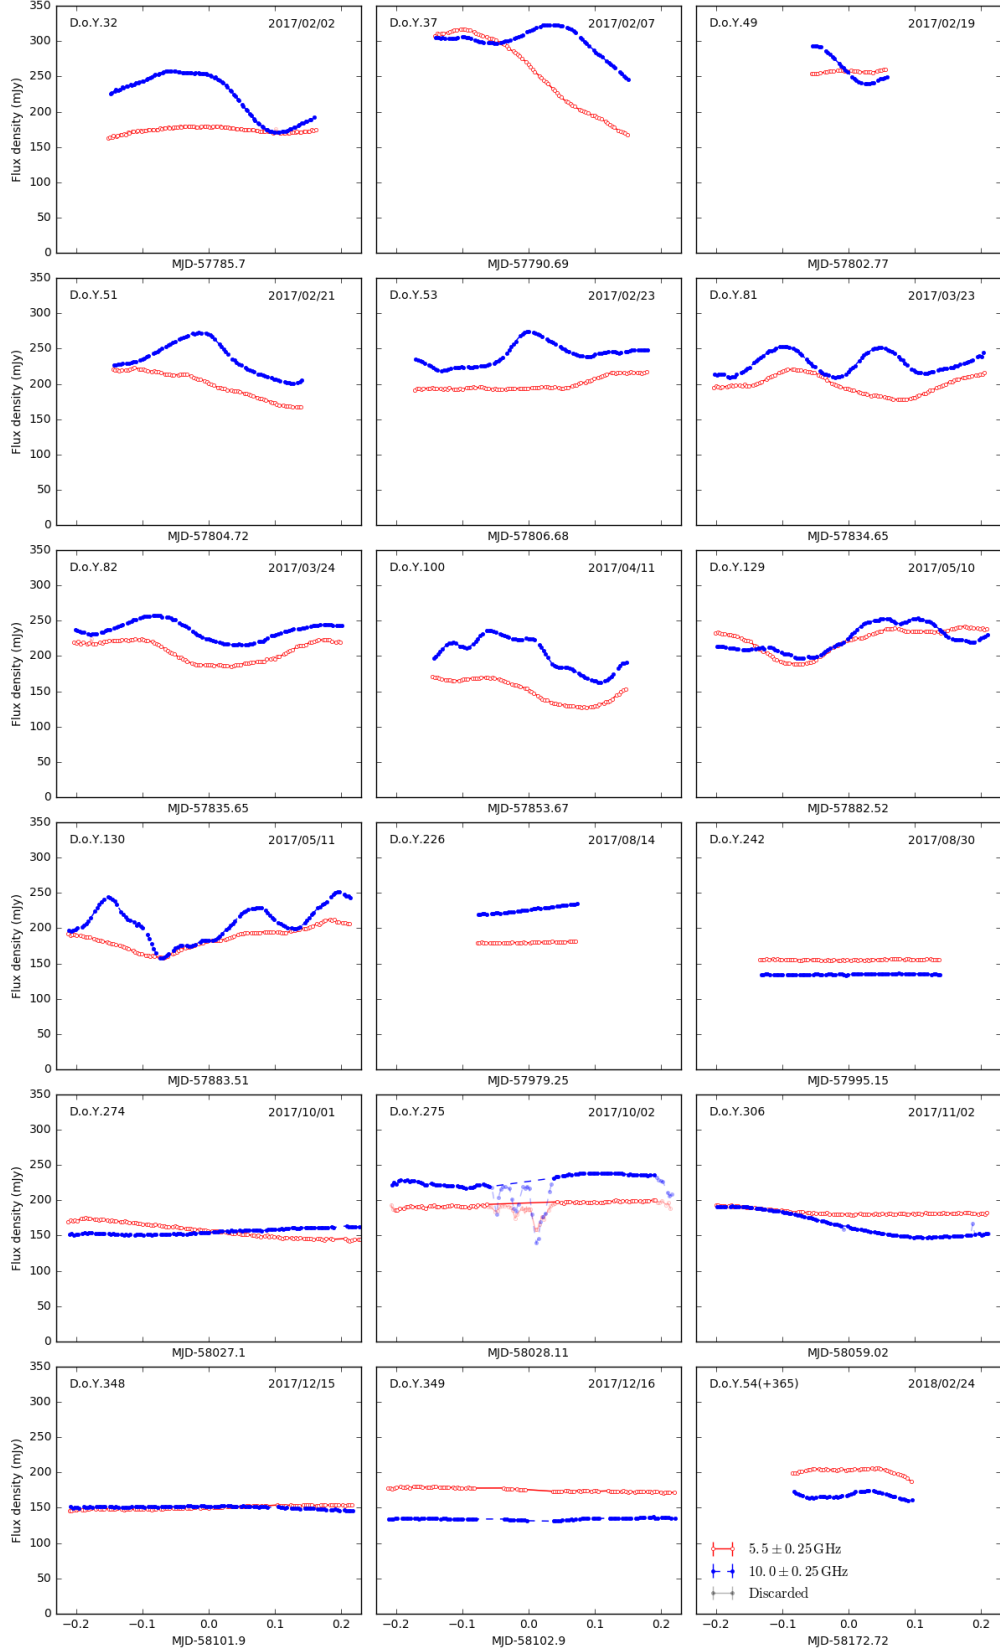


Figure 1. Observed light curves of PKS B1322–110, averaged over 0.5GHz-wide bands near the bottom and top of our bandwidth. The error bars, mostly too small to be seen, estimate the uncertainty of the mean. The points shown in semi-transparent were discarded from the analysis.

this is an assumption that is motivated by simplicity and the availability of suitable tools for the subsequent analysis. Let f_{ij} be the j -th data point of the i -th light curve measured at time t_{ij} with some uncertainty σ_{ij} uncorrelated between the measurements. Neglecting the correlations between any of our I epochs and assuming that the variations on a given epoch are due to the time-dependent magnification of the intrinsic flux density \bar{f}_i , the likelihood of observing these data for a Gaussian process is given by the product

$$p(\{f_{ij}\}, \{t_{ij}\}) = \prod_{i=1}^I \det^{-1/2}(2\pi\mathbf{M}_i) \times \exp\left[-\frac{1}{2} \sum_{jj'} (f_{ij} - \bar{f}_i) (\mathbf{M}_i)_{jj'}^{-1} (f_{ij'} - \bar{f}_i)\right], \quad (1)$$

where \mathbf{M}_i is the covariance matrix given by the sum of the uncorrelated measurement noise and the matrix of the autocorrelation function (ACF) values at the observed time lags

$$M_{ijj'} = \delta_{jj'} \sigma_{ij}^2 + \bar{f}_i^2 K_i(t_{ij} - t_{ij'}). \quad (2)$$

The auto-correlation functions $K_i(\Delta t)$ are unknown but we will assume that they derive from the same underlying $K(\Delta t)$ by an epoch-dependent rescaling of its argument:

$$K_i(\Delta t) = K(r_i \Delta t), \quad i = \overline{1, I}. \quad (3)$$

That assumption reflects the expectation that the variations are statistically uniform in space, with the rate varying only because of the changing orbital velocity of the Earth. The likelihood (1) thus encodes information on the rates, r_i , via the Bayes theorem.

At the outset of our study it was unclear whether or not a Gaussian process should provide a good description of the IDV phenomenon, and our choice of model was driven mainly by the need for a method that is both tractable and unbiased when the variability timescales are long. We will see later that our approach has proved only partially successful.

In practice, the most computationally expensive part of evaluating (1) is calculating the determinant of \mathbf{M} and the value of the quadratic form in the argument of the exponent, given K_i . Unless the kernel is chosen very carefully, with numerical efficiency in mind, Bayesian inference is not computationally feasible even for moderately-sized datasets. Recently a fast algorithm, *celerite*, was developed (Foreman-Mackey et al. 2017, see also Rybicki & Press 1995) that can be used to compute the likelihoods (1) very efficiently for a class of kernels $K_i(\Delta t)$ represented by a sum of exponentials with complex coefficients, of which we only consider even functions:

$$K(\Delta t) = \sum_{n=1}^N a_n e^{-c_n \Delta t} \cos d_n \Delta t, \quad a_n, c_n, d_n > 0. \quad (4)$$

This form appears well matched to the damped oscillations seen in the correlation functions of published IDV light-curves (e.g. Bignall et al. 2003). We make use of the algorithm realised as a Python package, *celerite*, released with the Foreman-Mackey et al. (2017) paper along with the *emcee* package (Foreman-Mackey et al. 2013) for Markov Chain Monte Carlo (MCMC) implementation. We experimented

with the number of terms in (4) by running the optimisation code *NLOpt* (Johnson 2010¹, Powell 2009²) for a fixed number of objective function evaluations and various N and compared them using the adjusted Akaike information criterion (e.g. Maier 2013), which produced more consistent results compared to similarly used Bayesian information criteria. In most cases the optimum value turned out to be $N = 1$ and it was never above four; moreover, the results for r_i did not seem to be much affected if just a single term was used. We thus settled on $N = 1$ in the MCMC calculations.

Rather than a single light curve, we record thousands of spectral channels, which all contain information on the rate of flux density variation. However, neighbouring channels are not independent with a decorrelation scale of a few GHz on most epochs, and the likelihood of the entire data set would need to account for correlations along the frequency axis by adding a pair of indices in addition to j, j' in (1) – and thus greatly increasing the dimensionality of the parameter space. This correlation structure is of little interest for the present work but adds substantial computational expense. To use as much data as possible and at the same time keep calculations practically feasible we extract two light curves by frequency-averaging the data near the edges of the observed bandwidth, one centred at 5.5 GHz, the other at 10 GHz, and replace (1) with a product of two such expressions, one for each sub-band (which assumes that the two light curves are not correlated). We use 0.5 GHz-wide intervals, the width where the empirical uncertainty of the mean over the interval (which includes both noise and real variations with frequency) approaches the expected thermal noise in the interval. This value is ~ 0.3 mJy for both sub-bands, and we use the empirical uncertainty of the mean as a measure of σ_{ij} in (2). We use the arithmetic mean of the light curve as an estimate of the intrinsic flux density, $\bar{f}_i = \langle f_{ij} \rangle_j$, allowing for its variation from epoch to epoch as appropriate for a compact flat spectrum source. We attempted to explore the \bar{f}_i parameter space with the MCMC method but failed to achieve convergence even after running the chains for several days – presumably due to the significant dimensionality added. The factors r_i are the same for both light curves. As the likelihood is invariant to scaling of all rates by the same factor while simultaneously scaling c_n, d_n coefficients by its inverse, a reference rate is specified by keeping $r_1 \equiv 1$ in the code. There are therefore $I - 1 + 2 \cdot 3$ parameters (rates for $I - 1$ epochs relative to the first epoch plus 3 model ACF parameters for each of 2 sub-bands) which are all positive and assumed *a priori* distributed uniformly in log between e^{-15} and e^{15} , measured in d^{-1} for c_n, d_n , and dimensionless otherwise.

Figure 2 presents the inferred absolute rates, R_i , defined as the inverse of the ACF half-width at half-maximum (HWHM):

$$R_i = \frac{1}{\tau_i} = \frac{r_i}{\tau}, \quad \text{where } \tau : K(\tau) = \frac{K(0)}{2}. \quad (5)$$

We choose to plot the scintillation rates rather than timescales because, for a one-dimensional model, the former are proportional to a component of the effective velocity, v_{eff}^\perp .

¹ <http://ab-initio.mit.edu/nlopt>

² http://www.damtp.cam.ac.uk/user/na/NA_papers/NA2009_06.pdf




Figure 2. Absolute variation rates (R_i , from equation 5), versus day of the year, at 5.5 GHz (the 10 GHz sub-band behaves very similarly— by construction, they only differ in the normalisation, as relative rates are the same for both sub-bands). The open symbol shows our single epoch of observation in 2018. Error bars extend over 68 per cent variation of the MCMC samples; points mark the medians of the distributions. The dotted line shows the prediction of a highly anisotropic model that associates the scattering plasma with Spica (Walker et al. 2017) – specifically, assumes that the plasma co-moves with the star while its anisotropy axis points at it and the source has a brightness temperature of $T_b = 10^{13}$ K.

As such the rates are expected to be a sinusoidal function of time, and the information content of the data is readily perceived. In §4.3 we give a qualitative analysis of the kinematic constraints that can be obtained from our data; for now we note that the phase of the sinusoid reflects the orientation of the plasma anisotropy, and the offset reflects the corresponding component of the plasma velocity.

4 FITTING THE ANNUAL CYCLE

4.1 Performance of the Walker et al. 2017 model

The dotted line in Figure 2 shows the prediction of the Walker et al. (2017) model, in which the scattering plasma is co-moving with Spica and highly anisotropic with its major axis pointing at the star; it has no free parameters. The absolute rate of the variations,

$$R = \frac{v_{\text{eff}}^{\perp}}{a_{\perp}}, \quad (6)$$

is determined by the effective transverse velocity (Cordes & Rickett 1998),

$$\mathbf{v}_{\text{eff}} = \mathbf{v}_{\text{screen}} - \mathbf{v}_{\oplus}, \quad (7)$$

and a_{\perp} , the HWHM of the ACF of the spatial structure of the scintillation pattern.

Qualitatively, the model prediction appears broadly consistent with the shape of the inferred variation rate, which depends on the kinematics only, but appears slightly off in normalisation – i.e., the model value of a_{\perp} is too high. Walker et al. (2017) associate the latter with the angular width of the source projected through the screen, at distance D_s , into the observer’s plane, assuming a Gaussian

source of peak brightness temperature $T_b = 10^{13}$ K. The ACF HWHM expected at the wavelength λ is then

$$a_{\perp} = \lambda D_s \sqrt{\frac{\bar{f}_i \log 2}{\pi k_B T_b}} \quad (8)$$

and Figure 2 uses an average \bar{f}_i of 191 mJy, resulting in $a_{\perp} \approx 2.26 \times 10^5$ km for 5.5 GHz and (Spica) distance of 77 pc. In reality a_{\perp} is also influenced by other aspects of the problem — the Fresnel scale and the strength of the scattering, for example (Goodman & Narayan 2006). However, apart from normalisation, the model appears to be consistent with the positions of standstills as well as both the positions and relative amplitudes of the two peaks of the rate annual curve.

In quantitative terms, the model presented by the dotted line in Figure 2 is a poor fit to the data; its χ^2 computed from the numbers in the figure is 101.4. As the fitting procedure treats the 17 rates inferred in Section 3 as effective measurements (with associated uncertainties) and the model it is compared to has not been fitted to this data, the expected value of χ^2 is 17, much lower than measured. Treating a_{\perp} (estimated to be uncertain by ~ 0.5 dex in Walker et al. 2017), as an additional free parameter reduces the discrepancy measure to 83.5, still much above the expectation of $16 = 17 - 1$. Interestingly though, by varying all three parameters of a one-dimensional model – a_{\perp} , orientation of the plasma anisotropy axis and plasma velocity relative to this axis, *c.f.* (9) and immediately thereafter – the best fit one can achieve has a χ^2 of 75.3, still far above the expected value of $14 = 17 - 1 - 2$. Although various interpretations are possible, this comparison suggests that the error bars in Figure 2 are underestimated because our statistical model is not entirely adequate.

That would not be surprising, given that the Gaussian process assumption and the chosen ACF parameterisation were motivated largely by a need for feasible computations³. Although any reasonable choice of the likelihood functional form would push the model to some sort of match with the data, MCMC methods use the likelihood ratios to decide how often a certain region of the parameter space is to be explored, and the convergence dynamics may be adversely affected by a poor statistical description. The foregoing concerns about the statistical model are reinforced by examination of our MCMC determinations of the scintillation rates for the individual epochs. In Figure 2 we can see examples where consecutive days – D.o.Y. 81 – 82 and D.o.Y. 129 – 130 – yield highly significant differences between the inferred rates, despite qualitatively similar light curves (Figure 1). By contrast, no significant difference is expected between consecutive days if the evolution of the rates is purely of kinematic origin, as the Earth’s orbital velocity changes only slightly from day to day. We also note that epochs close to the expected standstills, where very low rates are inferred, yield small fractional rate uncertainties, whereas we expect fractional uncertainty of order unity, because of sample variance (less than one oscillation sampled within our observing window). Rescaling the error bars by the square-root of the 75.3 : 14 ratio of χ^2 values (actual:expected) so as to bring the best-fit model reduced χ^2 to unity makes the Walker

³ Unfortunately we are at present unable to suggest any better statistical models.

et al. (2017) model consistent with the data at better than 1σ level. Figure 4 illustrates the effect of rescaling on consistency of rate estimates on consecutive days.

4.2 Kinematic MCMC fitting

The least squares analysis of the kind described in the previous subsection does not take into account correlations between the rate estimates, which arise due to the dependence of the likelihood (1) on the global parameters of the ACF model as well as the individual r_i . It is therefore most appropriate to fit the kinematic model directly to the light curves by substituting $r_i = R(t_i)/R(t_1)$ with absolute rates R_i given by (6) and running the MCMC code on the parameter space of the kinematic model.

For a highly anisotropic plasma screen, the expected variation rate is proportional to the (absolute value of the) component of transverse effective velocity across the anisotropy axis,

$$a_{\perp} R_i = |\mathbf{v}_{\text{eff}}(t_i)\hat{\mathbf{e}}^{\perp}| = |v_{\text{screen}}^{\perp} - v_{\oplus}(t_i) \sin[\text{PA}_{\oplus}(t_i) - \text{PA}]|. \quad (9)$$

The parameter space of our kinematic model is thus two-dimensional, spanned by the position angle of the anisotropy axis, PA, and the component of the screen velocity orthogonal to that axis, $v_{\text{screen}}^{\perp}$. The velocity component parallel to the anisotropy axis does not affect rates for the one-dimensional model and cannot be constrained. It is convenient to keep $v_{\text{screen}}^{\perp}$ non-negative with $\text{PA} \in [0, 2\pi)$ such that

$$\hat{\mathbf{e}}^{\perp} \equiv \hat{\mathbf{e}}_{\text{Dec}} \cos\left(\text{PA} + \frac{\pi}{2}\right) + \hat{\mathbf{e}}_{\text{RA}} \sin\left(\text{PA} + \frac{\pi}{2}\right) \quad (10)$$

and $\mathbf{v}_{\text{screen}}$ make an acute angle:

$$v_{\text{screen}}^{\perp} \equiv (\hat{\mathbf{e}}^{\perp} \cdot \mathbf{v}_{\text{screen}}) \geq 0. \quad (11)$$

We calculate the Earth barycentric velocity transverse to the line of sight using the `get_body_barycentric()` function of the `astropy.coordinates` package (The Astropy Collaboration et al. 2018). To handle the non-trivial topology of the PA domain in the MCMC exploration, technically we reparametrise the problem

$$v_{\text{screen}}^{\perp}, \text{PA} \rightarrow x, y : \begin{pmatrix} x \\ y \end{pmatrix} = v_{\text{screen}}^{\perp} \begin{pmatrix} \cos \text{PA} \\ \sin \text{PA} \end{pmatrix} \quad (12)$$

The priors on x, y are flat within $\pm 300 \text{ km s}^{-1}$. The MCMC implementation is otherwise the same as for the rate estimates.

Figure 3 presents the results of this modelling, reparametrised back to the $(v_{\text{screen}}^{\perp}, \text{PA})$ space. The maximum density of the samples, represented as shades, is formally achieved near $(v_{\text{screen}}^{\perp}, \text{PA}) \approx (11.2 \text{ km s}^{-1}, 151^{\circ})$ but an extended region of the parameter space is consistent with the data. In particular, the position of the Walker et al. (2017) model, $(v_{\text{screen}}^{\perp}, \text{PA}) \approx (16.8 \text{ km s}^{-1}, 177^{\circ})$, is consistent with the data at just inside the 2σ level. Also marked at $(v_{\text{screen}}^{\perp}, \text{PA}) \approx (9.5 \text{ km s}^{-1}, 144^{\circ})$ is the model that minimises the χ^2 difference between the kinematic model and the rates as shown in Figure 2. Solid lines show the confidence intervals of this χ^2 statistic whereas light dotted lines are the same levels after rescaling the error bars so as to bring the reduced χ^2 of the best fit to unity.

A word of caution concerning the kinematic fitting is that our Markov Chains could not achieve convergence, as




Figure 3. Posterior distribution of the parameters of the anisotropic scattering screen. Colours represent the confidence intervals of the MCMC posterior distribution, as described in §4.2, around its densest point, marked with a cross. Contour lines show the results of χ^2 fitting to the rate estimates shown in Figure 2 (ignoring correlations between those estimates), as described in §4.1. Contours are drawn at the levels marked in the colour bar, which correspond to the standard 1-, 2- and 3- σ levels of a 1D Gaussian distribution. Solid contours show χ^2 computed using the rates with unscaled error bars (as shown in Figure 2); dotted contours show χ^2 computed using the rates with error bars scaled up by a factor of 2.32 (as shown in Figure 4), which correspond to a minimum χ^2 of 14. The location of the minimum is marked with a plus. The star corresponds to the model of Walker et al. (2017) that attributes scattering to radial plasma filaments co-moving with Spica; the model is near the 2- σ contour of the MCMC and unscaled χ^2 analysis, and not significantly different from the best χ^2 fit if error bars are inflated.

judged by the conventional heuristics (Hogg & Foreman-Mackey 2017). However, we do not observe any indications of unusual behaviour of the `emcee` walkers, nor do we see significant multi-modality of the posterior distribution. Why the autocorrelation time estimates continue to rise almost linearly with the chain length is not clear; one possibility is an inadequacy of the statistical model, as discussed above. As a result, we are not certain which set of lines or shades in Figure 3 best represents our uncertainty about the kinematic parameters of the scattering plasma, but it is clear that this region extends over a large range of position angles.

Figure 4 partly explains the significant extent of this region by showing the performance of the three models pinpointed in Figure 3 in fitting the variation rates inferred from the light curves. Qualitatively, the three model curves perform similarly well and appear hardly distinguishable given the quality of constraints presented.

4.3 Qualitative analysis

We will now explain the origin of the degeneracy in the kinematic parameters seen in Figure 3, see what inferences can be drawn directly from the rate curve and describe how the annual cycle can be analysed qualitatively.




Figure 4. Annual cycles of the variation rate for the three models marked in Figure 3 compared to the values inferred in Section 3 (with inflated error bars). Despite considerable difference in the input parameters, particularly in the orientation of the anisotropy, the differences in the predicted curves are small and, for these error bars, not significant.

4.3.1 Hodograph

In this section we will extensively use the hodograph of the velocity vector, which is the locus of the terminal points of a variable vector as its initial point is held fixed. In particular, the top panel of Figure 5 plots the hodograph of the component \mathbf{v}_\oplus of the Earth velocity that is transverse to the line of sight to PKS B1322–110. Roman numerals along the curve mark the start of the corresponding months. The hodograph is, to a high accuracy, an ellipse, and the low ecliptic latitude of PKS B1322–110 (it is just 2° below the ecliptic plane) gives the hodograph its highly elongated shape.

According to (7), the effective velocity \mathbf{v}_{eff} on a particular date is given by the vector from the corresponding point on the hodograph curve to a fixed point in the plot, the transverse velocity of the screen, $\mathbf{v}_{\text{screen}}$. As an example, Figure 5 shows the transverse velocity of Spica with an open star symbol. For highly anisotropic scattering, only one component of the effective velocity affects the rate of scintillation, the component v_{eff}^\perp , orthogonal to the major axis of the illumination pattern, at a position angle of PA. Geometrically, this component is equal to the distance between the respective point on the hodograph curve and the straight line that passes through $\mathbf{v}_{\text{screen}}$ and has the position angle of the pattern anisotropy axis. In contrast, the effective velocity component along this line has no bearing on the scintillation timescales; conversely, this component cannot be deduced from the annual cycle analysis. Hence, the screen transverse velocity implied by the annual cycle can lie anywhere along this line and therefore the line itself represents the kinematic model:

$$\mathbf{v}_{\text{screen}} \cdot \hat{\mathbf{e}}^\perp = v_{\text{screen}}^\perp \quad (13)$$

Its two parameters are the position angle PA and distance from the origin, v_{screen}^\perp , with an $n \times \pi$ ambiguity in the direction of the major axis resolved by (11) such that $v_{\text{screen}}^\perp \geq 0$. Line representation is useful for qualitative analysis of the most salient features of the rate curve – the position of its standstills (where present), and the positions and relative

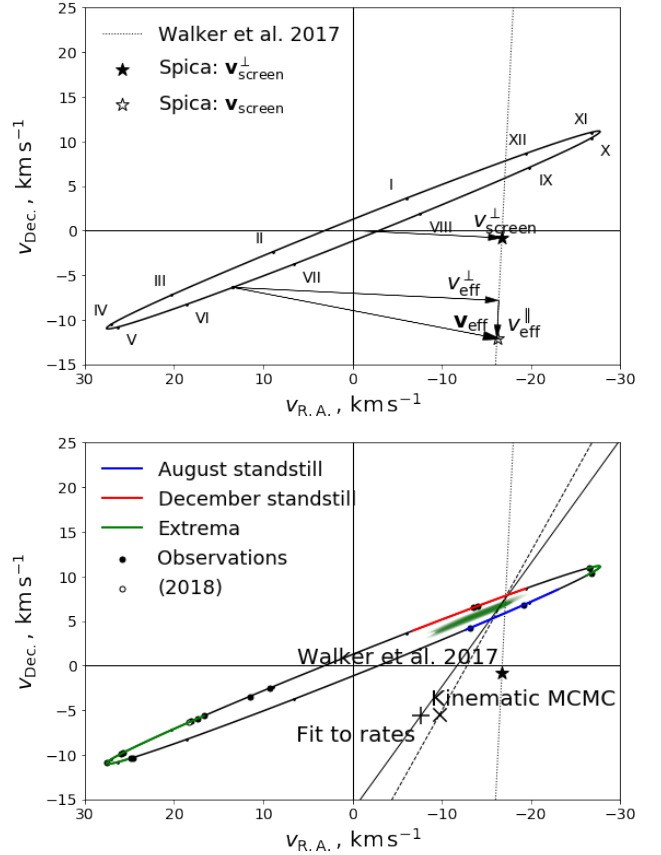


Figure 5. (Top) Hodograph of the Earth orbital velocity transverse to the direction to PKS B1322–110. Roman numerals mark the beginning of respective months. The effective velocity \mathbf{v}_{eff} is given by the vector from the point on the hodograph to the transverse velocity of the screen (shown with an open star for Spica). In highly anisotropic scattering, the scintillation rate depends on the component v_{eff}^\perp orthogonal to the anisotropy axis (shown with a dotted line for PKS B1322–110–Spica orientation). (Bottom) The hodograph displaying the location of standstills and extrema of the rate from Figure 2 along with the observation dates. The successful kinematic model should pass through the uncertainty intervals of both standstills on the hodograph as well as a shaded region inside the curve, representing the constraint (17) from the extrema. Also shown are the three kinematic models marked in Figure 3; all do a reasonable job reproducing standstills but struggle somewhat to fit the observed extremal rates.

magnitudes of the rate local maxima. They have a very simple interpretation on the hodograph plot: the standstills are the intersections of the hodograph with the model line (so that the rate, proportional to v_{eff}^\perp is zero), and at the extrema the hodograph tangents are parallel to the line (so that the line-hodograph distance is stationary). The bottom panel of Figure 5 shows lines representing the three kinematic models identified in Figure 3.

An alternative – and more compact – representation of the kinematic model is by the end point of the vector

$$\mathbf{v}_{\text{screen}}^\perp \equiv v_{\text{screen}}^\perp \hat{\mathbf{e}}^\perp \quad (14)$$

(which is $\neq \mathbf{v}_{\text{screen}}$, and neither to v_{screen}^\perp as it also depends on the orientation of the anisotropy axis); a filled star in Fig-

ure 5 marks $\mathbf{v}_{\text{screen}}^{\perp}$ of the Walker et al. (2017) model. While less natural for qualitative analysis, this representation is convenient when considering multiple models simultaneously – e.g., when comparing them or representing uncertainty regions, which are difficult to visualise unambiguously for sets of straight lines.

4.3.2 Standstills

If the annual curve contains standstills, where the rate drops to zero, accurate knowledge of their positions is sufficient to determine PA and $v_{\text{screen}}^{\perp}$ with no input from variable epochs. Since the rate is given by the distance from the point on the hodograph to the straight line representing the kinematic model ($\text{PA}, v_{\text{screen}}^{\perp}$), this distance should be zero at standstill – *i.e.*, the line should pass through the standstill. This is also true of the second standstill, and the two therefore completely define the kinematic model. Formally, by requiring (9) to vanish, one obtains a sinusoid in the $(v_{\text{screen}}^{\perp}, \text{PA})$ plane,

$$v_{\text{screen}}^{\perp} = v_{\oplus}^{\text{stand}} \sin(\text{PA}_{\oplus}^{\text{stand}} - \text{PA}), \quad (15)$$

where $v_{\oplus}^{\text{stand}}$, $\text{PA}_{\oplus}^{\text{stand}}$ are the magnitude and position angle of the Earth transverse velocity at the time of the standstill, both known. In the $\mathbf{v}_{\text{screen}}^{\perp}$ representation, the condition is

$$\mathbf{v}_{\oplus}^{\text{stand}} \cdot \hat{\mathbf{e}}^{\perp} = v_{\text{screen}}^{\perp} \Leftrightarrow (\mathbf{v}_{\oplus}^{\text{stand}} \cdot \hat{\mathbf{e}}^{\perp}) \hat{\mathbf{e}}^{\perp} = \mathbf{v}_{\text{screen}}^{\perp} \quad (16)$$

– *i.e.*, that $\mathbf{v}_{\text{screen}}^{\perp}$ is an orthogonal projection of $\mathbf{v}_{\oplus}^{\text{stand}}$; the locus of such $\mathbf{v}_{\text{screen}}^{\perp}$ is a circle of which $\mathbf{v}_{\oplus}^{\text{stand}}$ is a diameter. The full solution is obtained by locating the intersection of the two standstill sinusoids $(v_{\text{screen}}^{\perp}, \text{PA})$ or circles $(\mathbf{v}_{\text{screen}}^{\perp})$.

In practice, the position of a standstill is known with some uncertainty due to gaps in variability monitoring. In fact, it is not possible to know for sure if an observed lull in fluctuations is due to the Earth being stationary with respect to the structure in the flux distribution, or because the flux distribution happens to have a locally flat area there. This replaces a pair of points on the hodograph track with a pair of uncertainty regions that contain the standstills, and any line passing through both regions is a potential solution. Likewise, the constraint lines in the parameter planes are replaced by sinusoidal $(v_{\text{screen}}^{\perp}, \text{PA})$ or circular $(\mathbf{v}_{\text{screen}}^{\perp})$ bands of non-zero (and varying) width.

Generally, the further apart the two uncertainty regions in the hodograph, the better constrained the position angle and the orthogonal component of the screen velocity are. However, the case in Figure 5 seems to be closer to the other extreme, with the kinematic parameters quite uncertain. Looking at the light curves in Figure 1 and inferred rates in Figure 2 one might argue that the standstills are observed around D.o.Y. 240 ± 15 (in particular, the light curve on D.o.Y. 242 seems featureless in both subbands) and D.o.Y. 350 ± 15 , as highlighted in the bottom panel of Figure 5. Because of the low ecliptic latitude of PKS B1322–110 the hodograph is tightly squeezed in the latitudinal direction and the two uncertainty regions are quite close to each other. As a result, the position angle of the line passing through these constraints can be anywhere from $\sim 120^{\circ}$ to 290° (N \rightarrow E) and its distance from the origin similarly varies from $\approx 3 \text{ km s}^{-1}$ to $\approx 22 \text{ km s}^{-1}$.

We note that the configuration of the standstill seasons observed in PKS B1322–110 is close to as bad as it gets in this sense. A different arrangement could, in principle, have been much better constraining; but in practice one is less likely to obtain a favourable configuration for a source that has low ecliptic latitude. One can typically expect much better constraints from standstills in the annual cycles of IDVs that are far from the ecliptic plane.

4.3.3 Extrema

Another easily interpretable trait of rate cycles, whether displaying standstills or not, is the timing and relative magnitude of the local extrema. For a one-dimensional model, the tangents to the hodograph at the extrema positions are parallel to the anisotropy axis; to a very high accuracy, they should be opposite to each other in the hodograph (and very close to half a year apart in the rate curve). The relative magnitudes of the rate extrema in turn require the kinematic constraint line to pass through a point \mathbf{v}_r on the line connecting extrema positions $\mathbf{v}_{\oplus,1,2}^{\text{ext}}$, whose distances to the two points are in the same ratio as the respective extremal rates, $r_{1,2}$. There are two such points, one in between the extrema positions and the other behind the extremum $\mathbf{v}_{\oplus,1}^{\text{ext}}$ of smaller magnitude r_1 :

$$\mathbf{v}_r = \mathbf{v}_{\oplus,1}^{\text{ext}}(1 - q) + \mathbf{v}_{\oplus,2}^{\text{ext}}q, \quad q = \pm \frac{1}{\frac{r_2}{r_1} \pm 1}, \quad r_1 < r_2; \quad (17)$$

only the in-between case (+) is consistent with cycles displaying standstills for a one-dimensional model.

The requirement that the constraint line passes through this point

$$\mathbf{v}_r \cdot \hat{\mathbf{e}}^{\perp} = v_{\text{screen}}^{\perp} \quad (18)$$

represents a circle of diameter \mathbf{v}_r in the $\mathbf{v}_{\text{screen}}^{\perp}$ space, or a sinusoid,

$$v_{\text{screen}}^{\perp} = v_r \sin(\text{PA}_r - \text{PA}), \quad (19)$$

in the $(\text{PA}, v_{\text{screen}}^{\perp})$ space. In the Solar system barycentre frame we expect $\mathbf{v}_{\oplus,1}^{\text{ext}} \approx -\mathbf{v}_{\oplus,2}^{\text{ext}}$ very accurately, hence

$$\mathbf{v}_r \approx \mathbf{v}_{\oplus,1}^{\text{ext}}(1 - 2q) \Leftrightarrow v_r \approx v_{\oplus,1}^{\text{ext}} \frac{r_2 \mp r_1}{r_2 \pm r_1}, \quad \text{PA}_r \approx \text{PA}_1^{\text{ext}}; \quad (20)$$

the upper sign is to be taken for cycles with standstills.

Similarly to the standstill constraints, uncertainties in the measured positions and relative magnitudes of the extrema blur the location of the point through which the kinematic constraint line should pass. Figure 2 suggests that the extrema are located around D.o.Y. 290 ± 20 and D.o.Y. 100^{+30}_{-50} , with the latter particularly uncertain, and the ratio of the extremal rates is about 3^{+2}_{-1} . The constraint that corresponds to these estimates is shown in Figure 5 as a shaded region inside the hodograph of the Earth's transverse velocity, and provides an upper limit on the magnitude of $\mathbf{v}_{\text{screen}}^{\perp}$ as well as a sign of its projection on the ecliptic plane. On the other hand, the extrema themselves are located near the 'pointy' ends of the hodograph where the position angle of the tangent line makes effectively a full swing rendering the inference of PA unreliable. Similarly to the case of standstills, this is partly due to the low ecliptic latitude of PKS B1322–110.

The standstill and extrema constraints are not automatically consistent but they have to be so for a one-dimensional scattering model, which might help identify the limits of its applicability to the data. For more general anisotropy the situation can be quite different – for example, standstills are not generally expected in such models, as that would require the two components of the effective velocity to vanish simultaneously. Likewise, the extrema are not generally expected to lie on the opposite sides of the hodograph, with a six month separation. The analysis of the more general case goes beyond the scope of this paper.

5 DISCUSSION

As is clear from the qualitative analysis of the hodograph, accurate knowledge of two standstill positions is sufficient to uniquely determine the two kinematic parameters of the totally-anisotropic scattering model. However, near standstills the scintillation rate is difficult to constrain via traditional ACF estimation methods, and that was part of the motivation for our Bayesian analysis. For any cycle displaying standstills it is therefore highly desirable to obtain good coverage near the standstill epochs, to pin down the timing as tightly as possible. That is not practicable in the first monitoring season because the standstill positions are initially unknown, but becomes easier in subsequent years; we are currently obtaining such data for PKS B1322–110.

In this paper we have mostly concerned ourselves with the analysis of the shape of the annual cycle, while paying comparatively little attention to its absolute scale. The latter, however, is of interest in connection with the distance to the scattering screen. The prescription used by Walker et al. (2017) associates the spatial scale of the scintillation pattern with the source angular size multiplied by the screen distance. In the case of PKS B1322–110, and a screen placed at the distance of Spica, that prescription matched the value fit for the same ($v_{\text{screen}}^{\perp}, \text{PA}$) to better than 10 per cent although, given the simplicity of the prescription – Walker et al. (2017) suggested likely errors of 0.25 dex – this match must be considered a coincidence. One thing to bear in mind is that although specific assumptions about the character of the scattering plasma – conventionally, a Gaussian random field with a power-law spectrum of density inhomogeneities – allows one to calculate a precise value of the HWHM of the ACF of the scintillation pattern as a function of screen distance, pinning down the appropriate values of those parameters is not easy, and the unknown source structure (frequency-dependent size and shape) always remains a part of the mix. Indeed, the screen distance estimates obtained by such detailed modelling for the three best-studied IDVs (PKS 0405-385, J1819+3845 and PKS 1257-326) have uncertainties of a factor of a few, and might vary by an order of magnitude between different models (Bignall et al. 2006; Dennett-Thorpe & de Bruyn 2000; Macquart & de Bruyn 2006, 2007; Rickett et al. 2002). Furthermore, theoretical models often assume the scattering to be isotropic (e.g., Goodman & Narayan 2006) and are thus not applicable to the present data, for which the shape of the annual cycle indicates strong anisotropy; a similar resource for anisotropic plasma would be valuable.

5.1 Does Spica host the scattering plasma?

One of our main interests in following the scintillations of PKS B1322–110 over the last year was the possibility of testing the suggestion of Walker et al. (2017) that the structures responsible for IDV are associated with foreground, hot stars. We have already shown (§4) that our data are consistent with that model; but that could be a fortuitous agreement, arising in the context of a completely different model, and it is useful to evaluate the probability of such a coincidence. To do that we need to construct a statistical model for the distribution of the relevant parameters – position angle and perpendicular velocity – of the population of blobs of scattering plasma. Our adopted model is this: isotropic distribution of plasma microstructure orientation; and, an isotropic Gaussian distribution of the transverse velocity components. These properties are generic to a large class of models, and as such are reasonable assumptions. To fully specify the prior we need, in addition, a value of the dispersion of the velocity distribution. Rather than making an ad hoc assumption, we note that very small and very large values of the velocity dispersion would both yield vanishingly small probabilities of obtaining our results by chance. We therefore proceed by determining the value of the velocity dispersion that maximises the chances of a coincidence, so that the probability we obtain is an upper limit.

To compute that limit we integrate the probability distribution of our prior over the region of parameter space that matches our data at least as well as the Walker et al. (2017) prediction (Figure 4). Doing so we find a probability of 0.0051 (at a velocity dispersion of 13.0 km s^{-1}), if the MCMC kinematic distribution is used to measure the quality of fit. If, on the other hand, we use the χ^2 results then the probability is 0.0073 (at a velocity dispersion of 11.9 km s^{-1}). Therefore we could indeed have obtained the observed agreement by chance, but the likelihood of doing so is less than 1%. That figure is small enough to conclude that our data support the Walker et al. (2017) model, but not small enough to put it beyond doubt.

6 CONCLUSIONS

Monitoring of PKS B1322–110 has revealed a strong annual cycle in the rate of its scintillations. The cycle, which appears to include two standstills, is consistent with a highly anisotropic model of the scattering plasma. Quantifying the timescale of the scintillations is challenging because it is often comparable to, and sometimes much larger than, the extent of the observing window. Using the *celerite* MCMC package we evaluated the scintillation rate for each epoch, by determining a scaling factor for a temporal autocorrelation function whose shape is assumed constant over the year. The rates determined in this way exhibit highly significant differences between consecutive days, whereas the Earth’s orbital velocity hardly changes at all, and we conclude that our statistical model underestimates the uncertainties. To bring the reduced χ^2 of the best-fit model to unity we needed to increase the error bars on the rates by a factor of 2.3.

Although our data for PKS B1322–110 prefer a model in which the plasma microstructure points $\sim 30^\circ$ away from Spica, the preference is weak, and a radial orientation – as

suggested by Walker et al. (2017) – is included in the 68% confidence interval for the χ^2 surface (after rescaling the error bars on the rates). The very broad angular distribution of the acceptable kinematic models is understood as a direct consequence of the the low ecliptic latitude (-2°) of PKS B1322–110, which makes the shape of the annual cycle insensitive to microstructure orientation.

Despite the large uncertainty in the position angle of the plasma filaments, the relatively narrow range of preferred velocities at each orientation means that we have a one-dimensional constraint region in a two-dimensional space. Making use of a generic, isotropic prior for the statistics of the kinematics of the scattering material, we find a 1% probability of discovering, by chance, a screen whose properties are at least as close to the data as those of the Walker et al. (2017) prediction. Our data thus provide some support for that model.

ACKNOWLEDGEMENTS

We thank Bill Coles and J.-P. Macquart for insightful suggestions. AVT thanks Svetlana Nikolaevna Shlenova for introducing him to the concept of hodograph. The Australia Telescope Compact Array is part of the Australia Telescope National Facility which is funded by the Commonwealth of Australia for operation as a National Facility managed by CSIRO. Observations reported in this paper were made under ATCA project codes C2914, C2965 and C3214.

REFERENCES

- Bannister K., Bignall H., Johnston S., Reynolds C., Stevens J., Tuntsov A., Walker M., 2017, *The Astronomer’s Telegram*, **10024**
- Bignall H. E., et al., 2003, *ApJ*, **585**, 653
- Bignall H. E., Macquart J.-P., Jauncey D. L., Lovell J. E. J., Tzioumis A. K., Kedziora-Chudczer L., 2006, *ApJ*, **652**, 1050
- Carter S. J. B., Ellingsen S. P., Macquart J.-P., Lovell J. E. J., 2009, *MNRAS*, **396**, 1222
- Cordes J. M., Rickett B. J., 1998, *ApJ*, **507**, 846
- Dennett-Thorpe J., de Bruyn A. G., 2000, *ApJ*, **529**, L65
- Dennett-Thorpe J., de Bruyn A. G., 2001, in Laing R. A., Blundell K. M., eds, *Astronomical Society of the Pacific Conference Series Vol. 250, Particles and Fields in Radio Galaxies Conference*. p. 133
- Dennett-Thorpe J., de Bruyn A. G., 2002, *Nature*, **415**, 57
- Dennett-Thorpe J., de Bruyn A. G., 2003, *A&A*, **404**, 113
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *Publications of the Astronomical Society of the Pacific*, **125**, 306
- Foreman-Mackey D., Agol E., Ambikasaran S., Angus R., 2017, *The Astronomical Journal*, **154**, 220
- Fuhrmann L., et al., 2002, *Publ. Astron. Soc. Australia*, **19**, 64
- Gabányi K. É., et al., 2007a, *Astronomische Nachrichten*, **328**, 863
- Gabányi K. É., et al., 2007b, *A&A*, **470**, 83
- Goodman J., Narayan R., 2006, *ApJ*, **636**, 510
- Griffith M. R., Wright A. E., Burke B. F., Ekers R. D., 1994, *ApJS*, **90**, 179
- Hogg D. W., Foreman-Mackey D., 2017, preprint, ([arXiv:1710.06068](https://arxiv.org/abs/1710.06068))
- Jauncey D. L., Macquart J.-P., 2001, *A&A*, **370**, L9
- Jauncey D. L., Johnston H. M., Bignall H. E., Lovell J. E. J., Kedziora-Chudczer L., Tzioumis A. K., Macquart J.-P., 2003, *Ap&SS*, **288**, 63
- Jauncey D., et al., 2016, *Galaxies*, **4**, 62
- Kedziora-Chudczer L., 2006, *MNRAS*, **369**, 449
- Kedziora-Chudczer L., Jauncey D. L., Wieringa M. H., Walker M. A., Nicolson G. D., Reynolds J. E., Tzioumis A. K., 1997, *ApJ*, **490**, L9
- Liu J., Liu X., 2015, *Ap&SS*, **357**, 165
- Liu X., Song H.-G., Marchili N., Liu B.-R., Liu J., Krichbaum T. P., Fuhrmann L., Zensus J. A., 2012, *A&A*, **543**, A78
- Liu B.-R., Liu X., Marchili N., Liu J., Mi L.-G., Krichbaum T. P., Fuhrmann L., Zensus J. A., 2013, *A&A*, **555**, A134
- Liu X., et al., 2017, *MNRAS*, **469**, 2457
- Macquart J.-P., de Bruyn A. G., 2006, *A&A*, **446**, 185
- Macquart J.-P., de Bruyn A. G., 2007, *MNRAS*, **380**, L20
- Maier R. S., 2013, arXiv e-prints, p. [arXiv:1305.5493](https://arxiv.org/abs/1305.5493)
- Marchili N., Krichbaum T. P., Liu X., Song H.-G., Gabányi K. É., Fuhrmann L., Witzel A., Zensus J. A., 2012, *A&A*, **542**, A121
- Rickett B. J., 1990, *ARA&A*, **28**, 561
- Rickett B. J., Witzel A., Kraus A., Krichbaum T. P., Qian S. J., 2001, *ApJ*, **550**, L11
- Rickett B. J., Kedziora-Chudczer L., Jauncey D. L., 2002, *ApJ*, **581**, 103
- Rybicki G. B., Press W. H., 1995, *Phys. Rev. Lett.*, **74**, 1060
- The Astropy Collaboration et al., 2018, preprint, ([arXiv:1801.02634](https://arxiv.org/abs/1801.02634))
- Walker M. A., de Bruyn A. G., Bignall H. E., 2009, *MNRAS*, **397**, 447
- Walker M. A., Tuntsov A. V., Bignall H., Reynolds C., Bannister K. W., Johnston S., Stevens J., Ravi V., 2017, *ApJ*, **843**, 15

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.