Monthly Notices
of the
ROYAL ASTRONOMICAL SOCIETY

MNRAS **546,** 1–17 (2026) https://doi.org/10.1093/mnras/staf2048

No evidence for local H_0 anisotropy from Tully–Fisher or supernova distances

Richard Stiskalek ,[1]★ Harry Desmond [2] and Guilhem Lavaux [3]

[1]*Astrophysics, University of Oxford, Denys Wilkinson Building, Keble Road, Oxford OX1 3RH, UK*
[2]*Institute of Cosmology and Gravitation, University of Portsmouth, Dennis Sciama Building, Portsmouth PO1 3FX, UK*
[3]*CNRS and Sorbonne Université, Institut d'Astrophysique de Paris (IAP), UMR 7095, 98 bis bd Arago, F-75014 Paris, France*

Accepted 2025 November 14. Received 2025 November 11; in original form 2025 September 19

ABSTRACT

Claims of local ($z \lesssim 0.05$) anisotropy in the Hubble constant have been made based on direct distance tracers such as Tully–Fisher galaxies and Type Ia supernovae. We revisit these using the CosmicFlows-4 Tully–Fisher W1 subsample, 2MTF and SFI++ Tully–Fisher catalogues, and the Pantheon+ supernova compilation (all restricted to $z < 0.05$), including a dipole in either the Tully–Fisher zero-point or the standardized supernova absolute magnitude. Our forward-modelling framework jointly calibrates the distance relation, marginalizes over distances, and accounts for peculiar velocities using a linear-theory reconstruction. We compare the anisotropic and isotropic model using the Bayesian evidence. In the CosmicFlows-4 sample, we infer a zero-point dipole of amplitude 0.087 ± 0.019 mag, or 4.1 ± 0.9 per cent when expressed as a dipole in the Hubble parameter. This is consistent with previous estimates but at higher significance: model comparison yields odds of 877 : 1 in favour of including the zero-point dipole. In Pantheon+ we infer zero-point dipole amplitude of 0.049 ± 0.013 mag, or 2.3 ± 0.6 per cent when expressed as a dipole in the Hubble parameter. However, by allowing for a radially varying velocity dipole, we show that the anisotropic zero-point model captures local flow features (or possibly systematics) in the data rather than an actual linearly growing effective bulk flow caused by anisotropy in the zero-point or expansion rate. Crucially, inferring a more general bulk flow curve we find results fully consistent with expectations from the standard cosmological model.

Key words: galaxies: distances and redshifts – distance scale – large-scale structure of the universe.

1 INTRODUCTION

The cosmological principle (CP) – that the Universe is isotropic and homogeneous on sufficiently large scales – plays a foundational role in modern cosmology. Combined with the assumption of General Relativity as the theory of gravity, it leads to the Friedmann–Robertson–Walker metric, which underpins the concordance model of cosmology, Lambda cold dark matter (ΛCDM). It is therefore crucial to test if the CP is satisfied and hence fit for the purpose of providing such a bedrock. Recent years have seen increased scrutiny of the CP, for example in the cosmic microwave background (CMB; J. Jones et al. 2023; E. Gaztañaga & K. Sravan Kumar 2024), Type Ia supernovae (SNe; N. Horstmann, Y. Pietschke & D. J. Schwarz 2022; C. Krishnan et al. 2022; W. Rahman et al. 2022; Z. Zhai & W. J. Percival 2022; J. A. Cowell, S. Dhawan & H. J. Macpherson 2023; B. Kalbouneh, C. Marinoni & J. Bel 2023; R. Mc Conville & E. Ó Colgáin 2023; F. Sorrenti, R. Durrer & M. Kunz 2023a; J. P. Hu et al. 2024; A. Sah et al. 2025; F. Sorrenti, R. Durrer & M. Kunz 2025), direct distance tracers (R. Watkins et al. 2023; P. Boubel et al. 2025), galaxy clusters (K. Migkas et al. 2021; A. Pandya et al. 2024), and distant radio galaxies, quasars, and gamma-ray bursts (O. Luongo et al. 2022; N. J. Secrest et al. 2022a; L. Dam, G. F. Lewis & B. J. Brewer 2023; N. Secrest et al. 2025) (for a review, see P. K. Aluri & others 2023). Moreover, increasing attention has been directed towards defining 'expansion' in metrics beyond Friedmann–Robertson–Walker to enable model-independent tests (B. Kalbouneh, C. Marinoni & R. Maartens 2024; R. Maartens et al. 2024; B. Kalbouneh et al. 2025; M. Sarma et al. 2025). Any evidence for anisotropy must be carefully validated before violation of so fundamental a principle as the CP can be concluded.

One powerful test of the CP involves examining the *local* expansion rate H_0 in different directions, which can be achieved by using direct distance tracers such as Tully–Fisher relation (TFR) galaxies and Type Ia SNe. This is doubly important as it may also be relevant for the Hubble tension, which is derived assuming isotropic expansion (A. G. Riess et al. 2022; E. Di Valentino et al. 2025). Without calibration through a lower rung of the distance ladder, an anisotropy in H_0 is completely degenerate with an anisotropy in the normalization of the TFR and the absolute magnitude of SNs. While it is therefore not possible to assess an H_0 anisotropy directly with this method, it seems implausible that a null detection of the degenerate combination would arise from compensated anisotropies in both H_0 and the normalization of the standardizing relation.

In this paper, we revisit claims of anisotropy in H_0 by fully forward-modelling the TFR and SN observables at $z < 0.05$. Our starting point is the recent study of P. Boubel et al. (2025) (hereafter B25), who conclude weak evidence for anisotropy in the

---

★ E-mail: richard.stiskalek@physics.ox.ac.uk

CosmicFlows-4 (CF4) W1 TFR subsample using a partial forward model. Following B25, we adopt the first-order deviation from isotropy in spherical harmonics, namely a dipolar modulation to the degenerate combination of the TFR or SN normalization and H_0. Our approach jointly calibrates the distance relation, marginalizes over distances and latent parameters describing the true values of the observables, and corrects for peculiar velocities using a linear-theory reconstruction based on J. Carrick et al. (2015) (hereafter C15). We generalize the analysis of B25 to three further independent data sets to provide a more comprehensive and robust assessment of local anisotropy, namely the 2MTF and SFI++ TFR catalogues and the Pantheon+ SN compilation (restricted to $z < 0.05$). We also, for the first time, explicitly investigate the impact of dust extinction on the results by investigating three qualitatively different maps with different priors on the extinction coefficients. This is potentially an important systematic because extinction is anisotropic across the sky.

However, we caution that although we search for a CP-violating anomaly, this is done under the assumption of the C15 reconstruction to account for peculiar velocities, which itself assumes the CP. In principle, the C15 reconstruction could already contain part of a cosmological dipole, which would then be subtracted in our analysis. Nevertheless, the bulk flow in C15 shows no significant deviation from ΛCDM (see also e.g. S. S. Boruah, M. J. Hudson & G. Lavaux 2020; R. Stiskalek et al. 2025a). Thus, the results presented here should be interpreted as a dipole in H_0 that would be superimposed on the dipole already present in the C15 peculiar velocity field, which is not CP-violating. Nevertheless, we also consider the total bulk flow curve from both C15 and a superimposed radially varying velocity dipole to test whether the resulting flow is in tension with ΛCDM expectations.

The remainder of this paper is structured as follows. In Section 2, we describe the data sets used in our analysis, including the CF4 TFR sample, the 2MTF sample, the SFI++ catalogue, and the Pantheon+ SN compilation. In Section 3, we present our methodology, including the forward-modelling framework, the peculiar velocity modelling scheme, the Bayesian evidence calculation, and the mock data procedure. Section 4 presents the results on both mock and observed data, while Section 5 presents further comparison with B25, discusses the more general ramifications of our results, and concludes. Appendix A describes the flow model in full, and Appendix B shows the complete posterior distribution for the CF4 W1 TFR inference.

All logarithms are base-10 unless otherwise stated. We use the notation $\mathcal{N}(x; \mu, \sigma)$ to denote a one-dimensional normal distribution with mean μ and standard deviation σ, evaluated at x; in higher dimensions μ is a vector and σ is replaced by a covariance matrix. We define $h \equiv H_0/\left(100\,\mathrm{km\,s^{-1}\,Mpc^{-1}}\right)$, where H_0 is the Hubble constant.

2 DATA

To test the hypothesis of a dipole in the zero-point calibration of redshift-independent distance indicators, or, equivalently, in the Hubble constant, we analyse multiple low-redshift catalogues. Our primary data set is the TFR subsample of the CosmicFlows-4 catalogue (CF4; R. B. Tully et al. 2023), restricted to photometry in the all-sky W1 band. In addition, we consider two other all-sky TFR samples: 2MTF (K. L. Masters, C. M. Springob & J. P. Huchra 2008; T. Hong et al. 2019) and the SFI++ catalogue (C. M. Springob et al. 2007). We also analyse the Pantheon+ compilation of Type Ia SN (D. Scolnic et al. 2022), along with its reanalysis by Z. G. Lane et al. (2024).

Since any large-scale dipole in the zero-point is degenerate with the local peculiar velocity field, we account for peculiar velocities using the linear reconstruction of C15. This choice, widely adopted in the literature (e.g. S. S. Boruah et al. 2020; K. Said et al. 2020; A. Carr et al. 2022; P. Boubel et al. 2024b, 2025), allows for direct comparison with recent studies, particularly B25, and facilitates consistent treatment of large-scale flows across all samples.

2.1 Tully–Fisher samples

Our first method for obtaining redshift-independent distance estimates is the TFR (R. B. Tully & J. R. Fisher 1977), which relates a spiral galaxy's rotational velocity to its absolute magnitude M. Given an observed apparent magnitude, this relation yields a distance modulus and thereby a peculiar velocity, since the observed redshift is a function of the distance and peculiar velocity. The linewidth parameter η is defined as

$$\eta \equiv \log \frac{W}{\mathrm{km\,s^{-1}}} - 2.5, \tag{1}$$

where W is the observed width of a galaxy's spectral line (typically H I), serving as a proxy for its rotational velocity. Throughout, we refer to η simply as the galaxy linewidth. We adopt the following quadratic form of the TFR:

$$M(\eta) = \begin{cases} a_{\mathrm{TFR}} + b_{\mathrm{TFR}}\eta + c_{\mathrm{TFR}}\eta^2 & \text{if } \eta > 0 \\ a_{\mathrm{TFR}} + b_{\mathrm{TFR}}\eta & \text{otherwise} \end{cases}, \tag{2}$$

where a_{TFR} and b_{TFR} are the zero-point and slope, respectively, and c_{TFR} models the curvature of the relation for high-linewidth (i.e. high-mass) galaxies.

2.1.1 CosmicFlows-4 TFR sample

We use the CF4 TFR catalogue, a subset of the broader CF4 compilation (R. B. Tully et al. 2023), consisting of 9792 galaxies with $z_{\mathrm{CMB}} \lesssim 0.05$ and no strict apparent magnitude limit (E. Kourkchi et al. 2020a, b). Our analysis uses photometry exclusively in the all-sky Wide-field Infrared Survey Explorer (WISE) W1 band (same as B25), with additional selection criteria applied. We require $\eta > -0.3$ (to eliminate dwarf and low-mass galaxies, which may follow a different TFR or have a higher scatter), Galactic latitude $|b| > 7.5^\circ$ to exclude the Galactic Zone of Avoidance, and a photometric quality flag of 5 (best). After these cuts, the final W1 sample contains 3246 galaxies. The highest photometric quality requirement excludes ~1300 galaxies from the W1 subsample, while the Galactic latitude cut removes only 44 galaxies from the final sample. The Zone of Avoidance is typically defined as $|b| < 5^\circ$ (e.g. L. Staveley-Smith et al. 1998), making our choice relatively conservative. In contrast, B25 do not apply either cut. Since the publicly released CF4 catalogue does not provide magnitude uncertainties, we adopt a conservative fiducial uncertainty of 0.05 mag, following E. Kourkchi et al. (2019). However, this uncertainty is subdominant to the intrinsic scatter in the TFR, which is approximately 0.35 mag (see fig. 2 or Fig. B1). Moreover, unlike B25, who impose a lower redshift cut of $cz_{\mathrm{CMB}} > 3000\,\mathrm{km\,s^{-1}}$, we do not apply a lower redshift limit in our main analysis, but discuss its impact in Section 5.

2.1.2 2MTF sample

Next, we consider the 2MASS Tully–Fisher Survey (2MTF), an all-sky sample of 2062 spiral galaxies with TFR measurements extending to redshifts of $z_{\mathrm{CMB}} \lesssim 0.03$ (K. L. Masters et al. 2008;

T. Hong et al. 2019). The survey is selected in the K band with an apparent magnitude limit of $K < 11.25$. We use the version of the catalogue compiled by S. S. Boruah et al. (2020), which removes duplicates from the SFI++ sample, includes only K-band photometry, and applies a quality cut on linewidths, retaining only galaxies with $-0.1 < \eta < 0.2$. These cuts leave 1247 galaxies.

2.1.3 *SFI++ TFR sample*

We also consider the SFI++ catalogue, an all-sky sample comprising galaxies and groups with TFR measurements extending to redshifts of $z_{\rm CMB} \lesssim 0.05$ (K. L. Masters et al. 2006; C. M. Springob et al. 2007). Unlike 2MTF, the survey does not impose a strict apparent magnitude limit, and photometry is provided in the I band. In this work, we use the galaxy-only version of the catalogue compiled by S. S. Boruah et al. (2020), who apply quality cuts to select galaxies within the 2M++ footprint, impose a lower linewidth threshold to exclude low-mass galaxies, and use an iterative procedure to reject TFR outliers. While S. S. Boruah et al. (2020) also impose a strict upper linewidth selection to ensure linearity of the TFR, we relax this constraint and allow for TFR curvature, resulting in a final sample of 2010 galaxies.

2.2 Pantheon+ supernovae sample

Type Ia SNs are widely used as standardizable candles in cosmology. The SALT2 model standardizes their light curve (J. Guy et al. 2007), yielding a standardized apparent magnitude via the Tripp formula (R. Tripp 1998):

$$m_{\rm standard} = m_{\rm obs} + \mathcal{A}x_1 - \mathcal{B}c, \tag{3}$$

where $m_{\rm obs}$ is the observed SN apparent magnitude, x_1 characterizes the light curve stretch, and c the colour. The global parameters $\mathcal{A}$ and $\mathcal{B}$ quantify the standardization with respect to stretch and colour, respectively. Combined with the absolute magnitude $M_{\rm SN}$, the standardized magnitude $m_{\rm standard}$ defines the distance modulus.

The Pantheon+ data set is a compilation of 1701 spectroscopically confirmed Type Ia SNs spanning redshifts from $z \sim 0.001$ to ~ 2.3 (D. Brout et al. 2022; A. Carr et al. 2022; E. R. Peterson et al. 2022; D. Scolnic et al. 2022). For consistency with our peculiar velocity modelling and to match the redshift range of the CF4 sample, we restrict the data to $z_{\rm CMB} \leq 0.05$, yielding a subset of 525 SNs. Pantheon+ combines the original Pantheon+sample (D. M. Scolnic et al. 2018) with updated low- and high-redshift SNs, incorporating improved photometric calibration, light-curve standardization, and systematic uncertainty modelling. Distance moduli are derived using the SALT2 light-curve fitter and corrected for selection effects via the BEAMS with Bias Corrections method (R. Kessler & D. Scolnic 2017), which introduces an additive bias correction term in equation (3). We use such bias-corrected apparent magnitudes $m_{\rm corr}$ from the Pantheon+ sample, which include a fiducial Tripp calibration. We therefore sample only $M_{\rm SN}$, not the light curve stretch and colour coefficients of the Tripp calibration.

Uncertainties in the distance moduli (or equivalently, bias-corrected apparent magnitude), including both statistical and systematic contributions, are encoded in a covariance matrix that incorporates uncertainty in the Tripp parameters $\mathcal{A}$ and $\mathcal{B}$, held fixed at fiducial values. The Pantheon+ data release therefore provides a set of standardized bias-corrected apparent magnitudes and an associated covariance matrix, with the only global parameter varied in our analysis being $M_{\rm SN}$, which is inferred jointly with the flow model.

2.2.1 *Lane et al. reanalysis*

A reanalysis of the Pantheon+ data was presented by Z. G. Lane et al. (2024), who constructed an SN covariance matrix designed to minimize dependence on the assumed cosmology and peculiar velocities. Unlike the standard Pantheon+ covariance matrix and bias-corrected apparent magnitudes, which assume a fiducial cosmology and Tripp parameters, the Z. G. Lane et al. covariance matrix does not include any contribution from the Tripp parametrization. This enables simultaneous inference of the Tripp coefficients, with the standardized apparent magnitudes expressed as in equation (3), and no further bias corrections. We treat this as a variant of the Pantheon+ data, denoted 'Pantheon+ Lane' henceforth.

2.3 Peculiar velocity data

In C15, the luminosity-weighted density field is reconstructed from the redshift–space positions of galaxies in the 2M++ catalogue using the iterative method of A. Yahil et al. (1991). 2M++ is a whole-sky redshift compilation of 69160 galaxies (G. Lavaux & M. J. Hudson 2011), derived from 2MASS photometry (M. F. Skrutskie et al. 2006) and redshifts from 2MRS (J. P. Huchra et al. 2012), 6dF (D. H. Jones et al. 2009), and Sloan Digital Sky Survey (SDSS) DR7 (K. N. Abazajian et al. 2009). Apparent magnitudes are corrected for Galactic extinction, k-corrections, evolution, and surface brightness dimming. The catalogue is magnitude-limited to $K < 11.5$ in the 2MRS region and $K < 12.5$ in the 6dF and SDSS regions.

The velocity field is derived from the galaxy density field using linear theory and must be scaled to match peculiar velocities via a parameter $\beta^\star$, which we treat as a free parameter of the model. The field is generated on a 256^3 grid with a box size of 400 h^{-1} Mpc, assuming $\Omega_{\rm m} = 0.3$. $\beta^\star$ is defined as

$$\beta^\star \equiv \frac{f\sigma_{8,\rm NL}}{\sigma_8^b}, \tag{4}$$

where f is the dimensionless growth rate, with $f \approx \Omega_{\rm m}^{0.55}$ in ΛCDM (F. R. Bouchet et al. 1995; L. Wang & P. J. Steinhardt 1998). The terms σ_8^b and $\sigma_{8,\rm NL}$ represent the fluctuation amplitude in the biased galaxy field and in the non-linear matter field, respectively. The value of σ_8^g in 2M++ was measured as 0.98 ± 0.07 (M. Westover 2007) and 0.99 ± 0.04 (J. Carrick et al. 2015). C15, combined with peculiar velocity samples, has been extensively used to constrain the growth of structure and the S_8 parameter (e.g. S. S. Boruah et al. 2020; K. Said et al. 2020; P. Boubel et al. 2024a; R. Stiskalek et al. 2025a).

3 METHODOLOGY

We adopt a forward modelling framework, predicting the observables from the model parameters and comparing them to the observed values through a likelihood function. This has the advantage that it exploits all the information in the data (as no summary statistics are used), and makes it straightforward to model systematics and other effects that impact the observed quantities. It also lends itself naturally to a Bayesian inference by determining the probability distributions of the parameters implied by the observational data.

In particular our method jointly calibrates the TFR relation (or SNs) with the galaxy bias and the peculiar velocity field, following the methodology developed in S25. While B25 also largely adopts a forward-modelling approach, one key step does not fall within that framework: by querying peculiar velocity in redshift space, they use the observed redshift to assign peculiar velocity to a galaxy, approximately updating the fiducial velocity field calibration of C15, and then use that to estimate the cosmological redshift. That is a backward-modelling step (going from the redshift observation to a model parameter, the cosmological redshift or analogously distance) that introduces complications in triple-valued zones where a single observed redshift can map to multiple line-of-sight distances (M. A. Strauss & J. A. Willick 1995). We instead model the observed redshift directly by querying the peculiar velocity field in real space. Here, we provide a brief summary of the model; a more detailed explanation is given in Section A, with further discussion in S25.

3.1 Tully–Fisher model

Each galaxy is described by its observed redshift, apparent magnitude, and linewidth, with which we infer the distance. The velocity field is modelled as a combination of a reconstructed internal flow, scaled by a parameter $\beta^{\star}$, and an external flow $\boldsymbol{V}_{\rm ext}$. In the absence of a velocity field model, the flow reduces to a constant $\boldsymbol{V}_{\rm ext}$ term. We also include a Gaussian dispersion parameter σ_v to account for small-scale velocities not captured by the reconstruction.

The distances and true linewidths of each galaxy are latent model parameters over which we marginalize. The true linewidths are drawn from an inferred Gaussian distribution (D. J. Bartlett & H. Desmond 2023) and compared to the observations through a likelihood function. Via the TFR, the true linewidths determine the absolute magnitudes, which, combined with the distances, yield predictions of the apparent magnitudes to be compared with the data. Likewise, the true model distances, together with the peculiar velocity model, provide predictions of the observed redshifts for comparison with the measurements. The forward model accounts for both homogeneous and inhomogeneous Malmquist bias. The homogeneous case is an assumption that sources are uniformly distributed in volume, yielding a distance prior $p(r) \propto r^2$ in the absence of selection effects. The inhomogeneous Malmquist bias is set by the density field of C15 and modelled as

$$n(r) = 1 + b_1 \delta(r), \tag{5}$$

where $\delta(r)$ is the density contrast at the galaxy's position. To ensure non-negative values, $n(r)$ is clipped from below at zero. We adopt this linear bias model to be consistent with the reconstruction of C15, which uses linear theory to relate the galaxy density field, effectively smoothed over scales of $4\,h^{-1}\,$Mpc, to a peculiar velocity field. We treat b_1 as a model parameter and infer it with a wide uniform prior.

In general, TFR samples are subject to a complex selection function combining HI flux, linewidth, optical magnitude, and possibly redshift selection. Following the empirical approach of G. Lavaux (2016), we model the prior distribution of galaxy distances as

$$\pi(r \mid \boldsymbol{\theta}) = \frac{n(r,\, \boldsymbol{\theta})\, f(r,\, \boldsymbol{\theta})}{\int {\rm d}r'\, n(r',\, \boldsymbol{\theta})\, f(r',\, \boldsymbol{\theta})}, \tag{6}$$

where $n(r,\, \boldsymbol{\theta})$ accounts for the inhomogeneous Malmquist bias through the large-scale density field, $\boldsymbol{\theta}$ collectively denotes the model parameters, and

$$f(r,\, \boldsymbol{\theta}) = r^p \exp\left[-\left(\frac{r}{R}\right)^q\right]. \tag{7}$$

Here p, q, and R are free parameters: $p \approx 2$ recovers the homogeneous Malmquist bias, R sets the characteristic scale of sample incompleteness, and q controls how sharply the completeness falls off. The normalization in equation (6) is computed explicitly since it depends on $\boldsymbol{\theta}$. This treatment remains phenomenological, approximating the selection rather than modelling it directly. A rigorous forward modelling of the survey selection, which requires knowledge of the selection function, is presented in B. C. Kelly, X. Fan & M. Vestergaard (2008) and we recently applied it to H_0 inference in R. Stiskalek et al. (2025b).

We extend the TFR zero-point $a_{\rm TFR}$ by introducing a dipole, such that

$$a_{\rm TFR} \rightarrow a_{\rm TFR} + \boldsymbol{\Delta}_{\rm ZP} \cdot \hat{\boldsymbol{r}}, \tag{8}$$

where $\boldsymbol{\Delta}_{\rm ZP}$ is the zero-point dipole vector and $\hat{\boldsymbol{r}}$ is the unit vector in the direction of each galaxy. The dipole amplitude is denoted $\Delta_{\rm ZP}$, with direction specified in Galactic coordinates as $(\ell_{\Delta_{\rm ZP}},\, b_{\Delta_{\rm ZP}})$. The zero-point $a_{\rm TFR}$ is perfectly degenerate with the Hubble constant H_0, since the distance modulus depends on the logarithm of the luminosity distance. This leads to a degenerate parameter combination $a_{\rm TFR} + 5\log h$, where $h \equiv H_0/(100\ {\rm km\,s^{-1}\,Mpc^{-1}})$. Following S25 and similar works, we express distances in units of $h^{-1}\,$Mpc, rendering the analysis independent of h and constraining only the degenerate combination.

As explored by B25, one may allow for a dipole in this degenerate parameter across the sky. Provided sufficient sky coverage, such a dipole can be constrained. This introduces four possible interpretations:

(i) An anisotropic zero-point with isotropic H_0;
(ii) An isotropic zero-point with an anisotropic H_0;
(iii) A combination of both aforementioned effects;
(iv) Spurious H_0 anisotropy arising from flows not captured by C15.

B25 found no evidence for significant spatial variation in linewidths within the CF4 W1 subsample, by comparing the linewidth distribution of sources in the northern and southern hemispheres of the ALFALFA survey (M. P. Haynes et al. 2018). Moreover, the use of WISE photometry minimizes systematic variation in magnitude calibration across the sky. Based on these considerations, they interpreted their measured dipole in $a_{\rm TFR} + 5\log h$ as a potential signature of anisotropy in H_0. Assuming the dipole in $a_{\rm TFR} + 5\log h$ arises from a dipole in H_0, the corresponding fractional variation in the Hubble constant is given by

$$\Delta H_0 / H_0 = 10^{\Delta_{\rm ZP}/5} - 1, \tag{9}$$

where $\Delta_{\rm ZP}$ is the magnitude of the dipole in the degenerate combination $a_{\rm TFR} + 5\log h$ and H_0 is the monopole term of the Hubble constant.

As a second extension, we introduce a radial dependence of $\boldsymbol{V}_{\rm ext}$ to test for potential signatures of H_0 anisotropy. We uniformly sample the magnitude and sky direction of $\boldsymbol{V}_{\rm ext}$ at $N_{\rm knots}$ radial distance knots (set to $0,\ 20,\ 40,\ 60,\ 80$ and $100\,h^{-1}\,$Mpc for CF4), and apply cubic interpolation of its Cartesian components to evaluate $\boldsymbol{V}_{\rm ext}$ at each galaxy position. This procedure allows for a smoothly varying $\boldsymbol{V}_{\rm ext}$. Since a dipole in H_0 would imply a linearly increasing $\boldsymbol{V}_{\rm ext}$ with distance, this extension serves as a consistency check of the inferred $\boldsymbol{\Delta}_{\rm ZP}$ dipole. It also lets us infer the radial dependence of the bulk flow more generally.

3.2 Pantheon+ supernova model

To model the Pantheon+ data, we adopt an approach analogous to that used for the TFR, with the primary difference being that the predicted apparent magnitude is a function of the SN properties. The Pantheon+ sample accounts for both statistical and systematic uncertainties, encapsulated in a covariance matrix for the apparent magnitudes. These uncertainties arise from the Tripp standardization, photometric calibration, and the heterogeneity of the contributing surveys. While the full covariance matrix provided in the Pantheon+ release includes contributions from peculiar velocities, our model accounts for these explicitly. We therefore use a reduced version of the covariance matrix with the peculiar velocity terms removed, as provided to us by Anthony Carr (private communication).

As before, galaxy distances are sampled from the prior of equation (6). These are converted to distance moduli μ, which combined with $M_{\rm SN}$ give the predicted apparent magnitudes as $m_{\rm pred} = \mu + M_{\rm SN}$ (the Tripp parameters are calibrated to fiducial values in the Pantheon+ sample). The likelihood is then evaluated against the standardized, bias-corrected magnitudes,

$$\mathcal{L}(\boldsymbol{m}_{\rm corr} \mid \boldsymbol{m}_{\rm pred}) = \mathcal{N}(\boldsymbol{m}_{\rm corr}; \boldsymbol{m}_{\rm pred}, \, \mathbf{C}), \tag{10}$$

where $\boldsymbol{m}_{\rm corr}$ is a vector of the standardized, bias-corrected apparent magnitudes, $\boldsymbol{m}_{\rm pred}$ is a vector of predicted apparent magnitudes, and $\mathbf{C}$ is the reduced Pantheon+ covariance matrix. After this, the rest of the inference follows the same steps as the TFR analysis. Similarly as for the TFR, we extend the standardized SN absolute magnitude to include a dipole term:

$$M_{\rm SN} \to M_{\rm SN} + \boldsymbol{\Delta}_{\rm ZP} \cdot \hat{\boldsymbol{r}}. \tag{11}$$

On the other hand, when using the covariance matrix of Pantheon+ SNs from Z. G. Lane et al. (2024), we explicitly sample the Tripp parameters $\mathcal{A}$ and $\mathcal{B}$ to predict the SN apparent magnitude (see equation 3), since their data does not assume a fiducial Tripp standardization or its contribution to the covariance matrix. We evaluate the likelihood in the apparent magnitudes as in equation (5) of A. Seifert et al. (2025), except that we explicitly sample the galaxy distances, rather than setting them deterministically from the observed redshift which is equivalent to assuming no redshift uncertainty and no peculiar velocities.

3.3 Galactic extinction

A potential systematic effect in inferring the dipole in the zero-point, or equivalently in the apparent magnitude, arises from the treatment of Galactic extinction. We test this on the CF4 subsample, for which the applied Galactic extinction corrections are available in the public data release. The reported magnitudes in a given waveband λ are given by (see section 2.4 of E. Kourkchi et al. 2019):

$$\overline{m}^{\lambda} = m^{\lambda}_{\rm total} - A^{\lambda}_{b} - A^{\lambda}_{k} - A^{\lambda}_{a}, \tag{12}$$

where $m^{\lambda}_{\rm total}$ is the measured total magnitude, A^{λ}_{b} is the Milky Way extinction, computed as

$$A^{\lambda}_{b} = R_{\lambda} \, E(B-V), \tag{13}$$

A^{λ}_{k} is the k-correction, A^{λ}_{a} is the total flux aperture correction, and $E(B-V)$ is the colour excess (reddening). As described by E. Kourkchi et al. (2019), the CF4 catalogue uses the D. J. Schlegel et al. (1998) 100 μm cirrus maps to extract the Milky Way $E(B-V)$ values. For wavebands $\lambda \in (u, g, r, i)$, the extinction coefficients R_{λ} are taken from E. F. Schlafly & D. P. Finkbeiner (2011), while for the infrared bands, $R_{\rm W1} = 0.186$ and $R_{\rm W2} = 0.123$, as determined by E. L. Fitzpatrick (1999).

To assess potential systematics from Galactic extinction corrections, we use the `dustmaps` package[1] (G. M. Green 2018) to extract $E(B-V)$ values at the angular positions of CF4 galaxies. We adopt the extinction maps of Y.-K. Chiang (2023) and Planck Collaboration XLVIII (2016). The former is a dust reddening map on the plane of the sky, derived from a reanalysis of E. F. Schlafly & D. P. Finkbeiner (2011), which in turn is based on D. J. Schlegel et al. 1998. It uses tomographically constrained templates from WISE galaxy density fields to remove contamination from the cosmic infrared background (CIB). The latter map, from Planck Collaboration XLVIII (2016), employs the generalized needlet internal linear combination (GNILC) method to separate Galactic dust emission from CIB anisotropies, yielding an alternative all-sky extinction estimate. We consider these approaches for modelling Galactic extinction:

(i) Use the $E(B-V)$ values from D. J. Schlegel et al. (1998), jointly sampling the extinction coefficient $R_{\rm W1}$. We adopt a Gaussian prior on $R_{\rm W1}$ centred at 0.19, with either a standard deviation of 0.01, consistent with the measurement of H. B. Yuan, X. W. Liu & M. S. Xiang 2013, or a broader, more conservative prior with standard deviation 0.05.

(ii) Use the $E(B-V)$ values from Y.-K. Chiang (2023) or Planck Collaboration XLVIII (2016), also jointly sampling $R_{\rm W1}$ under the same prior choices as above.

The other TFR samples, 2MTF and SFI++, are compiled in the optical and hence may be more susceptible to dust extinction. We cannot test this explicitly as the dust corrections they used are not publicly available, but note that the very small effect between different dust models found for CF4 suggests that the effect still would not be significant. Similarly, we do not consider Galactic extinction variations in the Pantheon+ samples.

3.4 Mock data generation

We use mock data to estimate the sample size required for a CF4-like survey to yield a significant detection of a dipole of given strength. The mock catalogue is designed to replicate the CF4 TFR subsample with W1 photometry and full-sky coverage, excluding the Galactic Zone of Avoidance. The injected parameter values, listed in Table 1, correspond to the posterior mean from the CF4 inference.

For N sources, we draw sky positions uniformly over the sphere, excluding a mock ‘Zone of Avoidance’ defined by $|b| < 7.5°$. Each line of sight is then used to query the C15 field for the density and peculiar velocity. The radial distance r is sampled from the prior in equation (6), which incorporates the linear galaxy bias of equation (5). At this distance we evaluate the line-of-sight peculiar velocity $\boldsymbol{v}(\boldsymbol{r})$ and compute the true redshift,

$$1 + z_{\rm true} = \big(1 + z_{\rm cosmo}(r)\big) \left(1 + \frac{[\boldsymbol{V}_{\rm ext} + \beta \boldsymbol{v}(\boldsymbol{r})] \cdot \hat{\boldsymbol{r}}}{c}\right), \tag{14}$$

where $z_{\rm cosmo}$ is the cosmological redshift at distance r and $\hat{\boldsymbol{r}}$ is the unit vector towards the source. The observed redshift is then drawn with scatter σ_v,

$$cz_{\rm obs} \hookleftarrow \mathcal{N}(cz_{\rm true}, \, \sigma_v), \tag{15}$$

assuming the ‘measurement’ error of $z_{\rm obs}$ to be subdominant to σ_v. The apparent magnitude is obtained by first sampling $\eta_{\rm true}$ from a Gaussian hyperprior,

$$\eta_{\rm true} \hookleftarrow \mathcal{N}(\hat{\eta}, \, w_\eta), \tag{16}$$

and then the observed linewidth as

$$\eta_{\rm obs} \hookleftarrow \mathcal{N}(\eta_{\rm true}, \, \sigma_\eta). \tag{17}$$

The true apparent magnitude is

$$m_{\rm true} = \mu(r) + M(\eta_{\rm true}), \tag{18}$$

---

[1] https://dustmaps.readthedocs.io/en/latest/

Table 1. Injected parameter values used to generate mock TFR catalogues mimicking the CF4 W1-band sample, assuming the C15 density and peculiar velocity field. The parameters correspond to the posterior mean from the dipole inference on the observed sample. A comparison of the resulting redshift distribution for one mock catalogue with the CF4 subsample is shown in Fig. 1.

Parameter	Injected Value	Description
a_{TFR}	-19.95	TFR zero-point
b_{TFR}	-9.6	TFR slope
c_{TFR}	10.5	TFR curvature for $\eta > 0$
σ_{int}	0.34	Intrinsic TFR scatter
σ_η	0.02	Linewidth uncertainty
σ_m	0.05	Magnitude uncertainty
b_{min}	7.5°	Galactic latitude cut: $\lvert b\rvert > b_{\mathrm{min}}$
V_{ext}	$160\,\mathrm{km\,s^{-1}}$	External flow magnitude
ℓ_{ext}	302°	External flow Galactic longitude
b_{ext}	-17°	External flow Galactic latitude
Δ_{ZP}	0.087	Zero-point dipole amplitude
$\ell_{\Delta_{\mathrm{ZP}}}$	127°	Zero-point dipole Galactic longitude
$b_{\Delta_{\mathrm{ZP}}}$	10°	Zero-point dipole Galactic latitude
σ_v	$280\,\mathrm{km\,s^{-1}}$	Redshift uncertainty including small-scale velocity dispersion
$\beta^\star$	0.43	Velocity scaling of C15
p	1.8	Distance prior low-distance exponent
R	$32\,h^{-1}\,\mathrm{Mpc}$	Distance prior incompleteness distance
q	1.4	Distance prior incompleteness rate

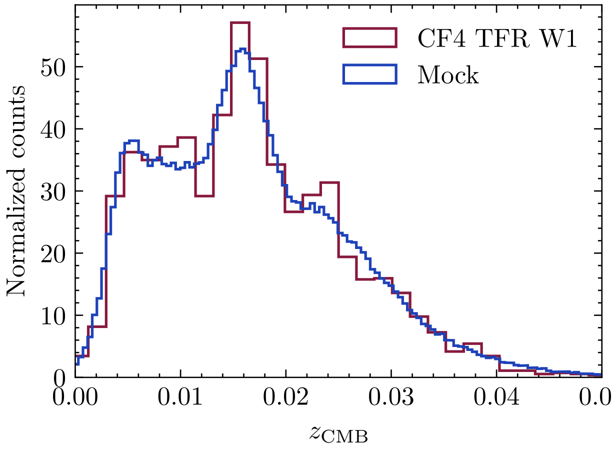

Figure 1. Comparison of the redshift distribution of the CF4 TFR W1 sample with a mock catalogue generated using the parameters in Table 1. The mock is designed to replicate the CF4 subsample and assess how the dipole detection significance depends on sample size.

where $\mu(r)$ is the distance modulus and $M(\eta_{\mathrm{true}})$ is defined in equation (2). The observed magnitude is drawn as

$$m_{\mathrm{obs}} \hookleftarrow \mathcal{N}\left(m_{\mathrm{true}},\ \sqrt{\sigma_m^2 + \sigma_{\mathrm{int}}^2}\right). \tag{19}$$

The resulting catalogue contains observed redshifts, magnitudes, linewidths, and sky positions. An example redshift distribution compared with the CF4 data is shown in Fig. 1.

3.5 Inference procedure

To sample the posterior distribution we use the No-U-Turn Sampler (NUTS; M. D. Hoffman & A. Gelman 2011) method of Hamiltonian Monte Carlo (HMC) as implemented in the `numpyro`[2] package (E. Bingham et al. 2019; D. Phan, N. Pradhan & M. Jankowiak 2019), ensuring a Gelman–Rubin statistic $\hat{R} - 1 \leq 0.001$ for convergence (A. Gelman & D. B. Rubin 1992). For each model, we run four independent chains with 1000 warm-up steps and 2000 sampling steps, typically yielding more than 4000 effective samples for the dipole and other parameters. We adopt a uniform prior on the zero-point dipole amplitude from 0 to 0.3 and sample its direction uniformly over the surface of a sphere. For all remaining parameters, we adopt uniform priors over sufficiently large ranges, except for σ_v, σ_{int} for which we adopt reference (and scale invariant) priors $\pi(x) \propto 1/x$. All these parameters are shared between both the isotropic and dipole models so their choice does not affect the Bayes factors between the models.

We calculate the Bayesian evidence to quantify models' relative goodness of fit. This is the probability of the data D given the model M:

$$\mathcal{Z} \equiv p(D \mid M) = \int \mathrm{d}\boldsymbol{\theta}\, \mathcal{L}(D \mid \boldsymbol{\theta},\, M)\, \pi(\boldsymbol{\theta}), \tag{20}$$

where $\boldsymbol{\theta}$ is the set of free parameters of the model, $\mathcal{L}$ is the likelihood function, and $\pi(\boldsymbol{\theta})$ is the prior on the parameters. We estimate this using the `harmonic` package,[3] which fits a normalizing flow to the posterior samples to compute evidence using the harmonic estimator (J. D. McEwen et al. 2021; A. Polanska et al. 2025). For the TFR samples we numerically marginalize over both galaxy latent parameters (r and η_{true}) at each Markov Chain Monte Carlo (MCMC) step, yielding a low-dimensional posterior suitable for evidence computation with `harmonic`. In contrast, for the Pantheon+samples the covariance prevents straightforward numerical marginalization of the distances, requiring them to be sampled explicitly. This results in an approximately 600-dimensional posterior, which makes evidence calculation challenging. While `harmonic` is designed to compute the Bayesian evidence directly from posterior samples of moderately high dimension, it has typically been tested on $\mathcal{O}(10)$-dimensional problems. For example, D. Piras et al. (2024) applied `harmonic` to a 39-dimensional posterior, validating their results against nested sampling, and also to a 159-dimensional posterior for which nested sampling was computationally infeasible. `harmonic` uses a normalizing flow within a harmonic-mean framework to perform density estimation for evidence computation, but for a ~ 600-dimensional posterior the required number of samples becomes prohibitive, convergence cannot be guaranteed, and its use has not been previously explored in such high dimensionality. Accordingly, we do not apply `harmonic` in this high-dimensional setting. We therefore restrict evidence computation to the TFR samples. For these, we compare models using the Bayes factor, defined as the ratio of pieces of evidence for two models M_1 and M_2,

$$\mathcal{B} \equiv \frac{p(D \mid M_1)}{p(D \mid M_2)}, \tag{21}$$

which quantifies the relative probability of the data under the two models. Assuming equal prior probabilities for the isotropic and dipole models, the Bayes factor is equal to the model odds $p(M_1 \mid D)/p(M_2 \mid D)$.

---

[2] https://num.pyro.ai/en/latest/
[3] https://github.com/astro-informatics/harmonic

Table 2. Summary of zero-point dipole and external flow parameters, as well as the fractional anisotropy in H_0, and Bayes factors in favour of the anisotropic model. Positive values of $\log \mathcal{B}$ indicate preference for the anisotropic model, negative values for the isotropic model. Only values with $|\log \mathcal{B}| \geq 2$ are considered as 'decisive' evidence (H. Jeffreys 1939).

Sample	$\Delta_{\rm ZP}$	$\ell_{\Delta_{\rm ZP}}$ [deg]	$b_{\Delta_{\rm ZP}}$ [deg]	$V_{\rm ext}$ [km s^{-1}]	$\ell_{\rm ext}$ [deg]	$b_{\rm ext}$ [deg]	$\Delta H_0/H_0$	$\log \mathcal{B}$
CF4 W1	0.087 ± 0.019	128 ± 12	9 ± 10	158 ± 20	303 ± 8	-17 ± 5	0.041 ± 0.009	$+2.94$
SFI++	0.051 ± 0.018	84 ± 25	-5 ± 17	168 ± 23	311 ± 7	4 ± 6	0.024 ± 0.008	$+0.47$
2MTF	0.041 ± 0.023	130 ± 83	-37 ± 29	150 ± 21	317 ± 9	0 ± 7	0.019 ± 0.011	-0.38
Pantheon+	0.049 ± 0.014	135 ± 56	-49 ± 18	212 ± 39	312 ± 10	-9 ± 7	0.023 ± 0.006	–
Pantheon+ Lane	0.060 ± 0.014	84 ± 50	-52 ± 14	220 ± 34	312 ± 8	-21 ± 6	0.028 ± 0.006	–

4 RESULTS

We present results from applying our flow model to various data sets, focusing primarily on the CF4 TFR W1 subsample. In Section 4.1, we report constraints on the zero-point dipole magnitude and direction under different Galactic dust corrections and extinction priors, and assess the model preference based on Bayesian evidence. In Section 4.2, we apply the same model to the 2MTF and SFI++ TFR samples. We then examine the Pantheon+ SN sample in Section 4.3. In Table 2 we summarize the recovered dipole, external flow, and Bayes factors for the CF4 W1 band, 2MTF, SFI++, and the Pantheon+samples, while in Fig. 4 we compare the inferred dipole direction with literature estimates, finding broad consistency. In Section 4.4 we investigate the robustness of these results to assumptions about peculiar velocity corrections and the prior on the dipole amplitude. In Section 4.5, we use mock TFR catalogues to assess the detectability of a zero-point dipole as a function of survey sample size, evaluating the level of anisotropy that could be robustly inferred given realistic distance uncertainties. Finally, in Section 4.6 we examine the bulk flow equivalent to the H_0 dipole, and compare it with both the bulk flow inferred when $\boldsymbol{V}_{\rm ext}$ is allowed a radial dependence and with ΛCDM expectation.

4.1 CosmicFlows-4 TFR W1 dipole

Applying our model to the CF4 TFR W1 sample, we infer a zero-point dipole with magnitude of 0.087 ± 0.019 mag and direction of $(\ell, b) = (127 \pm 11^\circ, 10 \pm 8^\circ)$ in Galactic coordinates, corresponding to a dipole in H_0 of 4.1 ± 0.9 per cent. We now focus on the marginal posterior for selected parameters; the full posterior is discussed in Section B. Fig. 2 shows the recovered posteriors for the zero-point dipole amplitude and direction, as well as for σ_v and $\sigma_{\rm int}$, which represent the velocity dispersion and intrinsic TFR scatter, respectively. Both σ_v and $\sigma_{\rm int}$ quantify the residuals of the model: a model that describes the data better is likely to yield smaller values of σ_v and $\sigma_{\rm int}$. We consider three cases:

(i) Isotropic TFR zero-point.

(ii) Dipole in the zero-point with fixed extinction from the CF4 catalogue.

(iii) Dipole in the zero-point with the extinction coefficient jointly inferred from the data, using one of three dust maps: D. J. Schlegel et al. (1998), Y.-K. Chiang (2023), or Planck Collaboration XLVIII (2016).

We find that neither the choice of dust map nor the sampling of the Galactic extinction coefficient has an appreciable effect on the inferred zero-point dipole. This indicates no significant degeneracy between the dipole and the Galactic extinction correction. We also consider two prior choices for the extinction coefficient $R_{\rm W1}$ (a narrow prior consistent with H. B. Yuan et al. (2013), and a broader, less informative alternative) and find that the results are similarly insensitive to this choice. All three models that include a dipole in the zero-point yield values of σ_v and $\sigma_{\rm int}$ that are near-identical to those recovered by the isotropic model without a dipole. This would suggest that introducing a dipole does not appreciably improve the fit or capture additional trends in the data. However, upon considering a model comparison we find that the Bayesian evidence favours the anisotropic model with $\log \mathcal{B} = 2.938$, or odds of 877 : 1. We also find that the pieces of evidence of the models with different extinction treatments are nearly identical, indicating that the CF4 TFR W1 data is insufficient to distinguish between the dust maps.

It is important to note that, since the dipole magnitude is a strictly positive quantity, even if the Cartesian components of the dipole vector are consistent with zero, the magnitude of the dipole can still be significantly positive (for example, if the components follow a zero-mean Gaussian, the magnitude of the vector follows a Maxwell–Boltzmann distribution). The fact that the posterior of the dipole amplitude is a few σ inconsistent with zero is therefore not a reliable indicator of the presence of a dipole, hence the necessity of computing the Bayes factor.

Our conclusions are consistent with those of B25, who report a dipole magnitude of 0.063 ± 0.016 mag, slightly lower than our inferred value and with a smaller quoted uncertainty. They find only a very weak support for the anisotropic model, reporting $\log \mathcal{B} = 0.43$ in favour of the model including a zero-point dipole. However, they derive the Bayes factor using an approximation based on the Bayesian Information Criterion. We compare our results to B25 further in Section 5.

4.2 2MTF and SFI++ dipole

We next consider the two alternative TFR samples: 2MTF and SFI++. While subsets of these were used in constructing the CF4 TFR sample, we now analyse them independently. We apply the same flow model as for the CF4 TFR W1 subsample, using identical priors on the zero-point dipole parameters (uniform in magnitude and direction on the sphere). We show the results in Fig. 3(a). For 2MTF, we infer a dipole magnitude of 0.040 ± 0.024 mag with the evidence marginally favouring the isotropic model, with $\log \mathcal{B} = -0.381$. For SFI++, the recovered dipole magnitude is 0.051 ± 0.018 mag, though this time the evidence ratio marginally favours the inclusion of the dipole with $\log \mathcal{B} = 0.465$. Compared to SFI++, 2MTF covers a smaller redshift range and contains fewer galaxies, so it is unsurprising that it is less informative regarding the zero-point dipole.

4.3 Pantheon+ dipole

Now we test the presence of a dipole in the Pantheon+ data, repeating the procedure described above with the SN flow model described in Section 3.2. Owing to correlated uncertainties in SN

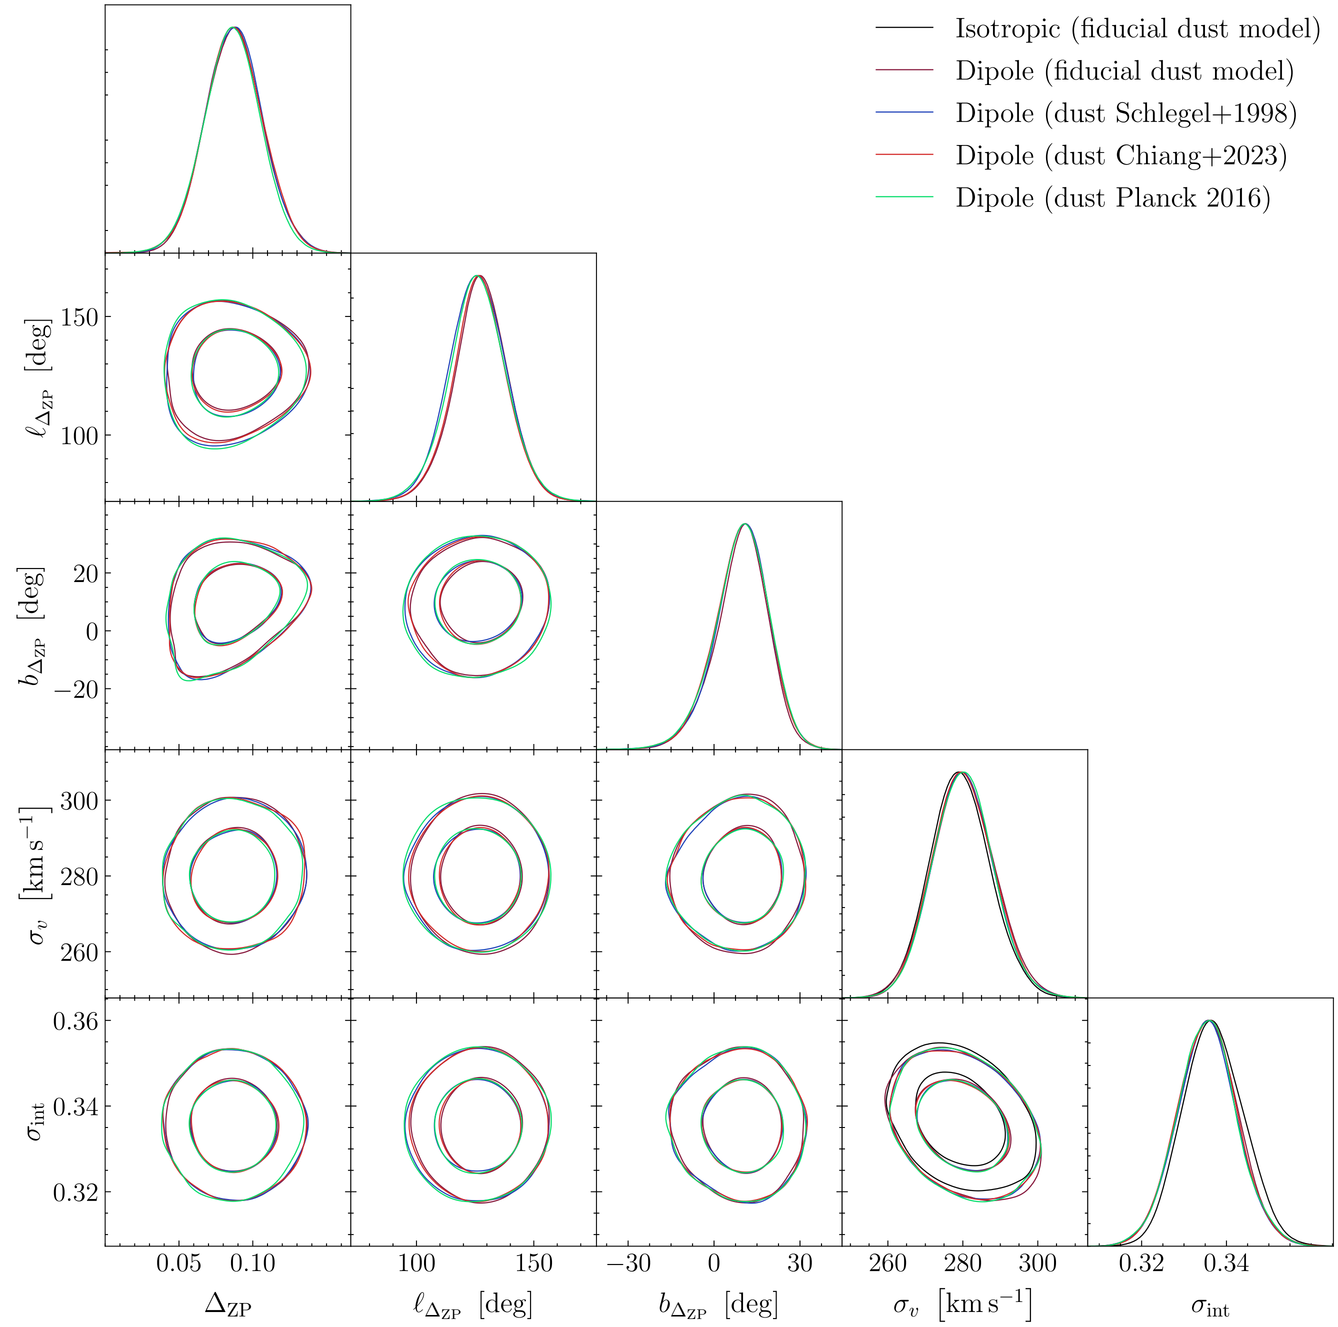

Figure 2. Comparison of the inferred TFR zero-point dipole under different Galactic dust corrections using the CF4 TFR W1 data: (a) fixed extinction from the CF4 catalogue, which is based on D. J. Schlegel, D. P. Finkbeiner & M. Davis (1998), (b) $E(B-V)$ from D. J. Schlegel et al. (1998) with a sampled extinction coefficient $R_{\rm W1}$, (c) $E(B-V)$ from the Y.-K. Chiang (2023) map, and (d) $E(B-V)$ from the Planck Collaboration XLVIII (2016) map, both also with sampled $R_{\rm W1}$. The recovered zero-point dipole is unaffected by the extinction treatment and its magnitude is $\Delta_{\rm ZP} = 0.087 \pm 0.19$ mag. The values of redshift scatter σ_v and TFR scatter $\sigma_{\rm int}$ show no improvement over the isotropic model without a zero-point dipole, though the Bayesian evidence favours the anisotropic models with odds of 877 : 1.

standardization, the sampled distances cannot be (easily) numerically marginalized, resulting in a high-dimensional posterior and we do not compute the Bayesian evidence. In Fig. 3(b) we show that for the Pantheon+ sample the recovered dipole amplitude is 0.049 ± 0.013 mag, which is nearly 4σ inconsistent with zero. In the Pantheon+ Lane sample the dipole amplitude is only marginally higher: 0.060 ± 0.13 mag. Since the dipole magnitudes are significantly different from zero, the data appear to favour the inclusion of a dipole even though we cannot compute the Bayes factor. The inferred dipole directions of the two samples are mutually consistent but show some tension with the TFR samples, possibly due to their different redshift depths.

Furthermore, we find that the intrinsic SN apparent magnitude scatter $\sigma_{\rm int}$ inferred by the model is 0.01 ± 0.01 for Pantheon+ and 0.14 ± 0.01 mag for Pantheon+ Lane(both with and without the zero-point dipole). D. Brout et al. (2022) propagate an additional

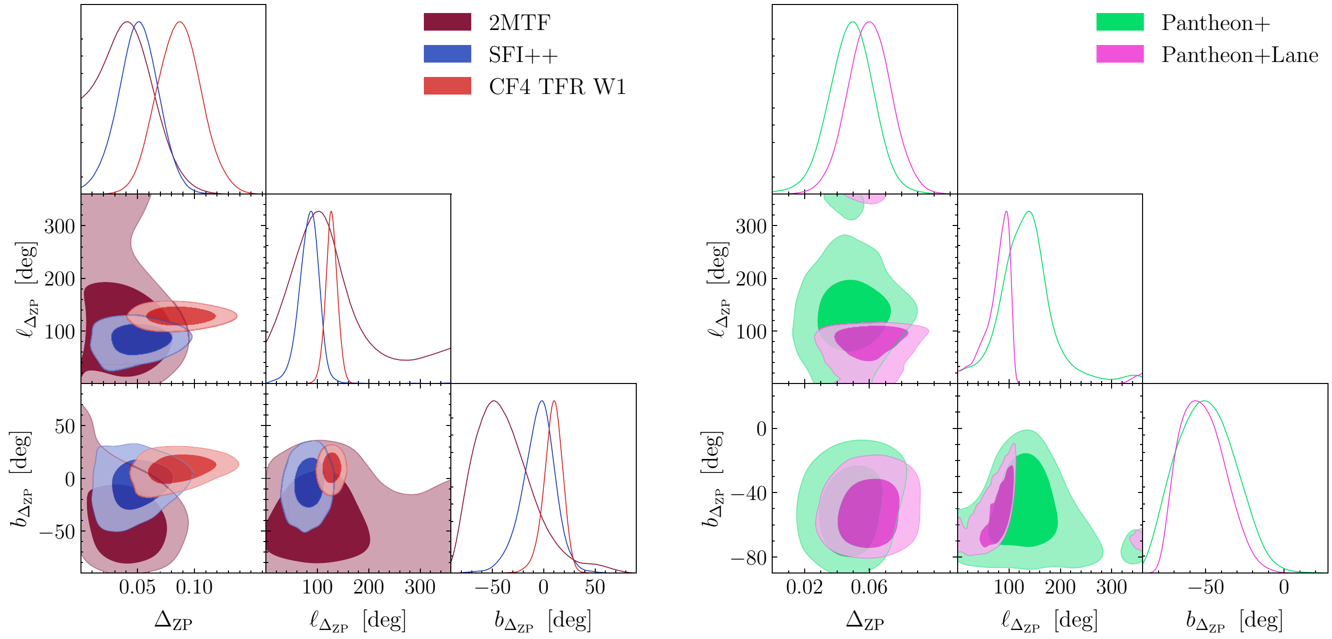

(a) (b)

Figure 3. (a) Inferred zero-point dipole in the TFR samples. Except for 2MTF (the least informative sample), the dipole magnitude is significantly offset from zero. (b) Inferred zero-point dipole in the SN sample. The dipole directions from the TFR and SN samples differ only slightly in the $b_{\Delta_{\rm ZP}}$ coordinate. Comparison of the inferred zero-point dipole in the TFR and SN samples. The contours show the 1σ and 2σ confidence regions.

intrinsic scatter term directly into the covariance matrix, which explains why no further scatter is required (indeed the propagated scatter may be overestimated). A. Sah et al. (2025) argue that the Pantheon+ sample may suppress deviations from ΛCDM through this treatment, contrasting it with the covariance matrices of the Joint Light-Curve Analysis (JLA; M. Betoule et al. 2014) and Pantheon+ Lane(Z. G. Lane et al. 2024), which have a much smaller diagonal scatter. However, our inference of an additional intrinsic scatter of 0.14 ± 0.01 mag indicates that the extra scatter propagated into the Pantheon+ covariance is justified.

4.4 Model variations

Here we briefly consider how the results change on altering two aspects of the model. First, we consider the case in which the C15 reconstruction is not used to model the underlying density field and peculiar velocities. In this scenario, we omit the inhomogeneous Malmquist bias and model the velocity field with a constant flow $\boldsymbol{V}_{\rm ext}$. We then repeat the inference using the CF4 TFR W1 subsample. In this case, the recovered dipole amplitude is 0.179 ± 0.023 mag, significantly larger than the value obtained when accounting for peculiar velocities. The model selection in this 'no peculiar velocity' scenario favours the inclusion of a zero-point dipole, effectively compensating for the unmodelled peculiar velocities, with a Bayes factor of $\log \mathcal{B} = 6.73$. On the other hand, for Pantheon+ we only infer a significantly larger zero-point dipole uncertainty when not accounting for peculiar velocities, with the dipole magnitude being 0.051 ± 0.035 mag. This highlights the importance of robustly accounting for peculiar velocities when investigating anisotropy, although it is important to bear in mind that these corrections may also absorb some anisotropy that may be present (as we discuss in Section 1).

Secondly, we consider the impact of the assumed prior on the dipole amplitude. Previously we have taken this to be a uniform prior on the magnitude of the dipole vector; we now make this uniform on the Cartesian components (the choice made by B25). This implicitly induces a prior on the dipole magnitude proportional to the square of the magnitude, pushing the posterior towards larger amplitudes. We repeat the analysis on the CF4 TFR W1 subsample using this component-based prior, and compare the results to those obtained with a uniform prior on the magnitude in Fig. 5. As expected, the recovered dipole magnitude is larger under the component prior, while the inferred direction is essentially unchanged. We find that the dipole magnitude is 0.095 ± 0.019 mag with the flat-component prior, compared to 0.087 ± 0.019 mag with the flat-magnitude prior, a $\sim 0.5\sigma$ difference.

4.5 Dipole mock analysis

Above we reported a significant preference for a dipole in the CF4 TFR W1 sample. We now inject a dipole of magnitude 0.087 mag, corresponding to the posterior mean value inferred from the CF4 sample, and examine how the evidence ratio between anisotropic and isotropic models depends on sample size. The remaining injected parameters are listed in Table 1, and an example redshift distribution of a mock sample is shown in Fig. 1. Following Section 3.4, we generate mock catalogues with 500, 1000, 2000, 4000, 8000, 16 000, and 32 000 sources, drawing ten random realizations for each size. We then apply our flow model to the mock TFR data, fitting both an isotropic model and one including a zero-point dipole, and compute the corresponding evidence ratio.

We present these results in Fig. 6. The Bayes factor becomes decisive ($\log \mathcal{B} > 2$ on the Jeffreys' scale, H. Jeffreys 1939) in favour of the dipole model only at sample sizes of order 2000. At smaller sizes it remains inconclusive and does not reliably

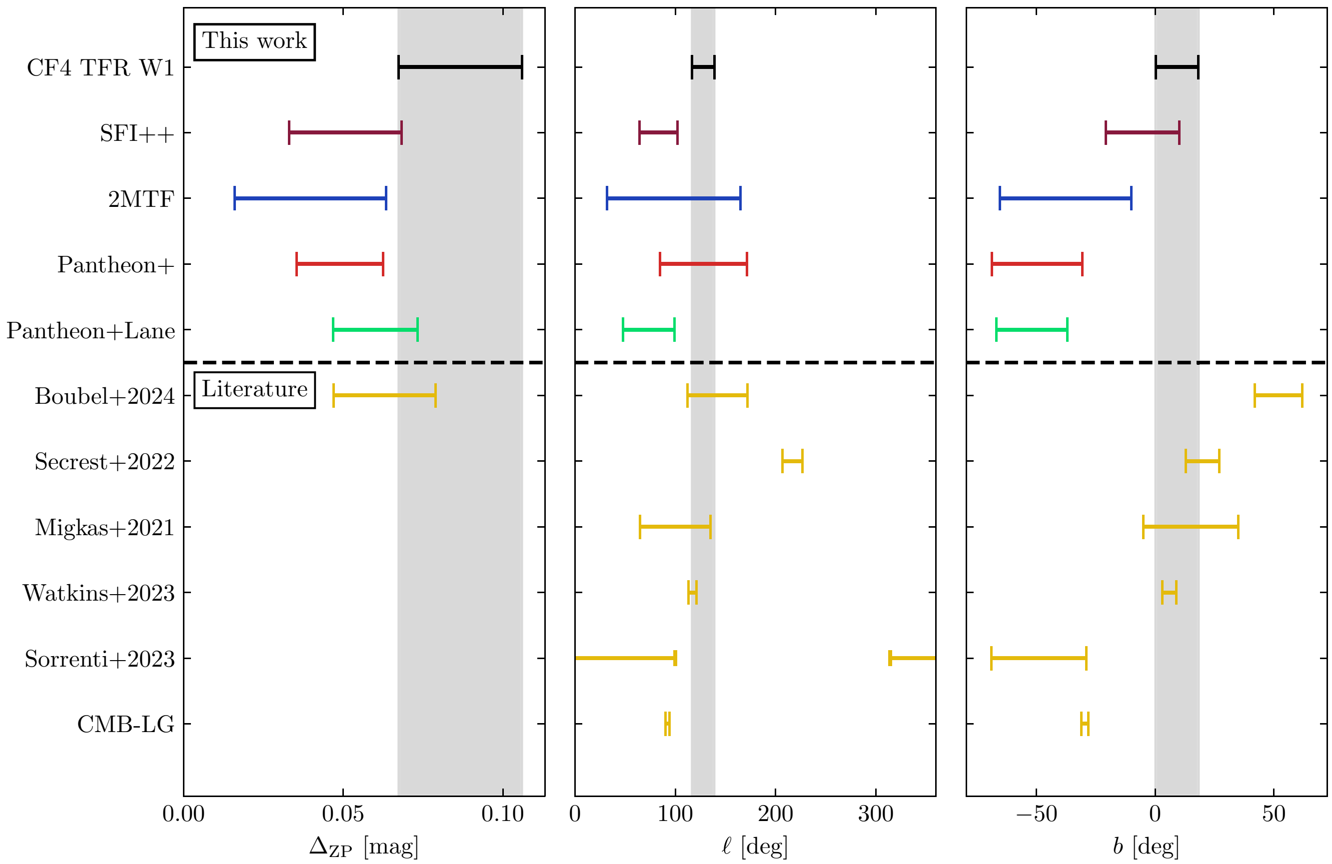

Figure 4. Summary of the inferred zero-point dipole for each catalogue, given by its magnitude $\Delta_{\rm ZP}$ and direction in Galactic coordinates (ℓ, b), compared with literature estimates: the CF4 TFR measurement (P. Boubel et al. 2025), the quasar dipole, the cluster scaling-relation dipole (K. Migkas et al. 2021), the CF4 bulk flow (R. Watkins et al. 2023), the Pantheon+ dipole (F. Sorrenti, R. Durrer & M. Kunz 2023b, their $z < 0.05$ sample), and the Local Group velocity in the CMB frame (Planck Collaboration I 2020a). For the latter four, we plot the opposite direction to that reported, as our dipole is defined in the zero-point rather than in peculiar velocities. We show the inferred magnitude only for P. Boubel et al. (2025), as the remaining works measure dipoles in different quantities that we do not convert. All results are reported in the CMB frame. Error bars indicate the 16th and 84th percentiles.

discriminate between the models. Since the W1 sample after our quality cuts contains 3246 galaxies, it is reassuring that the Bayes factor favours the zero-point dipole model near-decisively at that scale, assuming the injected dipole magnitude. The expected Bayes factor preference of mock samples of similar size to CF4 is likewise consistent with our results on the data. Fig. 6 is also consistent with the absence of a significant dipole detection in the smaller 2MTF and SFI++ samples. For mock samples without an injected dipole the Bayes factor is always negative, favouring the isotropic model, although the strength of this preference increases only weakly with sample size. We note that these results are specific to a CF4 TFR W1-like survey: while deeper surveys could in principle provide stronger constraints through larger samples, potential gains may be offset by increased absolute distance uncertainties at higher redshift.

4.6 Bulk flow measurement

The inferred dipole in $\boldsymbol{\Delta}_{\rm ZP}$ can be interpreted as a dipole in H_0 (see equation 9), or alternatively as a bulk flow linearly rising with distance under a cosmological model with isotropic expansion, which in turn corresponds to a bulk flow that increases linearly with distance. The bulk flow is typically defined as the average velocity within a volume V of radius R centred on the observer, assuming a top-hat filter (R. Watkins, H. A. Feldman & M. J. Hudson 2009; H. A. Feldman, R. Watkins & M. J. Hudson 2010; A. Nusser & M. Davis 2011; Y. Hoffman, H. M. Courtois & R. B. Tully 2015; R. Watkins & H. A. Feldman 2015; A. Nusser 2016; M. I. Scrimgeour et al. 2016; M. Feix, E. Branchini & A. Nusser 2017; W. A. Hellwing, M. Bilicki & N. I. Libeskind 2018; S. Peery, R. Watkins & H. A. Feldman 2018; R. Watkins et al. 2023; A. M. Whitford, C. Howlett & T. M. Davis 2023). In our case, three possible contributions enter:

$$\boldsymbol{B}(R) = \frac{1}{V}\int_V \left[\beta^\star \boldsymbol{v}(\boldsymbol{r}) + \boldsymbol{V}_{\rm ext}(\boldsymbol{r}) - H_0 r \left(10^{\Delta_{\rm ZP}/5} - 1\right)\frac{\boldsymbol{\Delta}_{\rm ZP}}{\Delta_{\rm ZP}}\right] \mathrm{d}V, \tag{22}$$

where the first term is the velocity field of C15 scaled by $\beta^\star$, the second is the external bulk flow $\boldsymbol{V}_{\rm ext}$ (optionally position-dependent), and the third is the bulk-flow equivalent of the zero-point dipole (or equivalently H_0 dipole), here approximated with a linear Hubble law (the full expression yields negligible differences). We assume here $H_0 = 73.04\ \mathrm{km\,s^{-1}\,Mpc^{-1}}$ (A. G. Riess et al. 2022) for illustrative purposes: the exact results depend minimally on the assumed H_0. Sources along the zero-point dipole direction are intrinsically fainter (i.e. have a larger zero-point), inducing a dipole in their apparent magnitudes but not in their redshifts, since they are assumed to be uniformly distributed in volume. If the model accounts only for a velocity dipole and not the zero-point dipole, it assigns all sources the same zero-point and hence absolute magnitude on average.

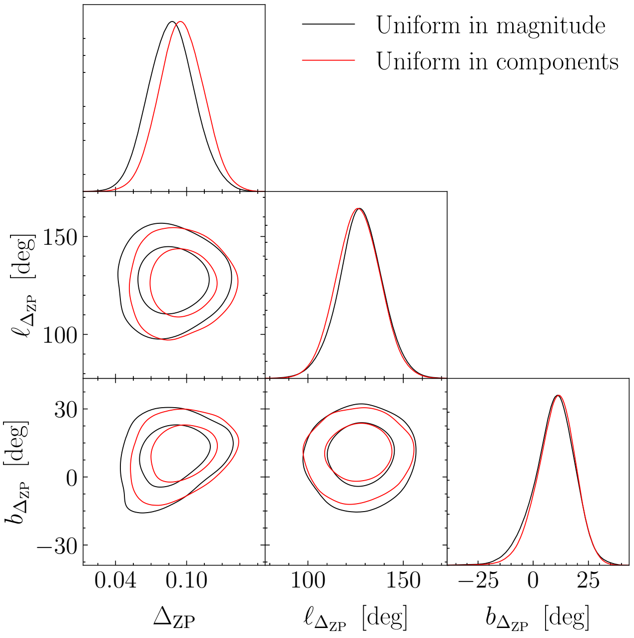

Figure 5. Comparison of the inferred dipole in the CF4 TFR W1 subsample under two different priors: uniform in the dipole magnitude or uniform in Cartesian components of the dipole. The latter induces a prior that favours larger magnitudes, resulting in a higher posterior dipole amplitude. The inferred direction is largely insensitive to the prior choice. Contours show the 1σ and 2σ confidence regions.

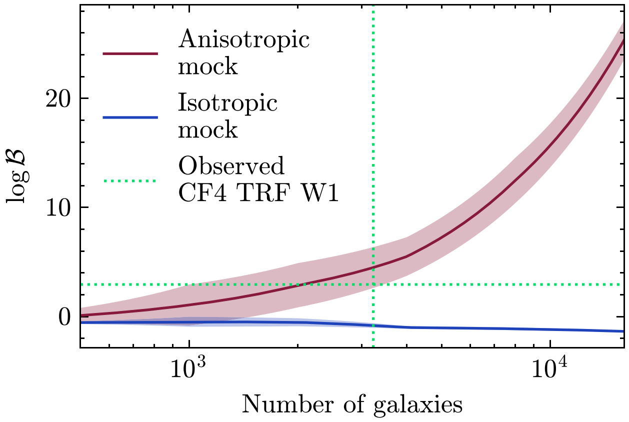

Figure 6. Model comparison using CF4 TFR W1-like mock data sets of varying size (see Section 3.4), testing an anisotropic zero-point dipole model against an isotropic alternative. Mocks are generated either with a zero-point dipole of magnitude 0.087 (equal to the posterior mean of our constraints; red) or without it (blue). Larger values indicate stronger preference for the anisotropic model. The green lines indicate the sample size and Bayes factor of the CF4 TFR W1 data, which are sufficient to yield 'decisive' evidence on the Jeffreys' scale ($\log \mathcal{B} > 2$, H. Jeffreys 1939). The shaded band shows the 16th–84th percentile range of the Bayes factor $\log \mathcal{B}$ across mock realizations of the same size.

Consequently, sources along the dipole direction are inferred as more distant. To match the isotropy of redshifts, the model compensates by introducing a negative velocity dipole along the zero-point direction, hence the minus sign in equation (22).

We compare the magnitude of this bulk flow to the ΛCDM expectation, computed with a top-hat filter of radius R (see e.g. section 6.1.2 of S. S. Boruah et al. 2020), assuming *Planck* 2018 cosmology (Planck Collaboration VI 2020b). This comparison is essential: while the data may show a strong preference for a dipole model (or more generally the existence of a bulk flow), it could still be entirely consistent with ΛCDM.

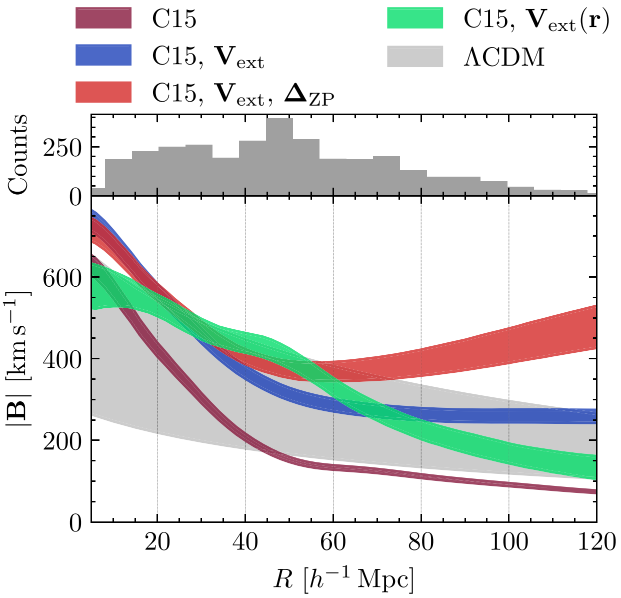

Figure 7. Bulk flow as a function of radius R inferred from the CF4 W1 sample. We first show the bulk flow from J. Carrick et al. (2015) only and then introduce three extensions of it: (i) a constant external flow $\boldsymbol{V}_{\rm ext}$, (ii) a constant $\boldsymbol{V}_{\rm ext}$ combined with the zero-point dipole $\mathbf{\Delta}_{\rm ZP}$, and (iii) a radially varying $\boldsymbol{V}_{\rm ext}(R)$. For reference, the shaded band shows the ΛCDM expectation computed with a top-hat filter of radius R, and the top panel displays the distribution of redshift distances in the CF4 W1 sample. All bands show the range between the 16th and 84th percentile (the width of the C15 is due to the $\beta^{\star}$ posterior.) All cases are consistent with ΛCDM expectations except the zero-point dipole model at large radii, where the (extrapolated) linearly growing dipole dominates. This is however disfavoured relative to the more flexible radially varying $\boldsymbol{V}_{\rm ext}$ model. For $\boldsymbol{V}_{\rm ext}(\boldsymbol{r})$, we place spline knots at 0, 20, 40, 60, 80 and 100 h^{-1} Mpc, indicated by vertical dashed lines.

We show the resulting bulk flow in Fig. 7 for three extensions of C15: (i) a constant $\boldsymbol{V}_{\rm ext}$ as in our main analysis, (ii) a constant $\boldsymbol{V}_{\rm ext}$ combined with $\mathbf{\Delta}_{\rm ZP}$, and (iii) a radially varying $\boldsymbol{V}_{\rm ext}(R)$. (In all cases the velocities contained within C15 are included.) As expected, the first case asymptotes to a constant above $R \approx 60\,h^{-1}$ Mpc, where the constant $\boldsymbol{V}_{\rm ext}$ dominates over the decaying contribution from C15. The second case behaves similarly at small radii, since the zero-point dipole contribution scales with distance, but yields a larger bulk flow above $R \approx 40\ h^{-1}$ Mpc and then grows linearly once the dipole contribution dominates. Finally, the radially varying $\boldsymbol{V}_{\rm ext}(R)$ produces a bulk flow slightly lower than the constant $\boldsymbol{V}_{\rm ext}$ case at small radii, higher at $R \approx 40$–$60\ h^{-1}$ Mpc, and decaying again at larger distances.

In Fig. 8 we show the magnitude of $\boldsymbol{V}_{\rm ext}(\boldsymbol{r})$ that is superimposed on the C15 velocity field. The magnitude rises from 200 km s^{-1} at small radii to about 300 km s^{-1} at 50 h^{-1} Mpc, after which it rapidly decays. $\boldsymbol{V}_{\rm ext}$ is interpolated between knots at 0, 20, 40, 60, 80 and 100 h^{-1} Mpc, with the magnitude at each knot sampled uniformly between 0 and 500 km s^{-1}. Fig. 8 demonstrates that this more flexible

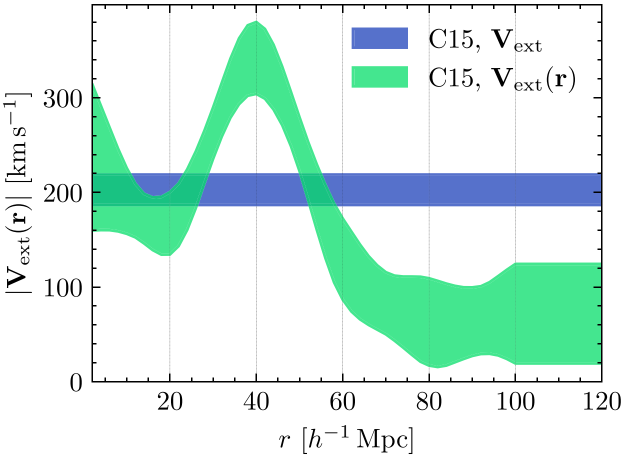

Figure 8. Magnitude of $\boldsymbol{V}_{\rm ext}(\boldsymbol{r})$ that is superimposed on the C15 velocity field for a radially varying external flow, compared to the constant $\boldsymbol{V}_{\rm ext}$ case. The shaded band indicates the 16th to 84th percentile range. The magnitude rises and then decays towards zero, inconsistent with the linearly growing bulk flow expected from H_0 anisotropy. For $\boldsymbol{V}_{\rm ext}(\boldsymbol{r})$, we place spline knots at 0, 20, 40, 60, 80 and 100 h^{-1} Mpc, indicated by vertical dashed lines. Above 100 h^{-1} Mpc the model uses constant extrapolation.

model is inconsistent with a linearly growing bulk flow, which would signal H_0 anisotropy. Furthermore, above 80 h^{-1} Mpc inferred $\boldsymbol{V}_{\rm ext}$ is nearly consistent with zero, while the uncertainty band broadens largely due to the lack of data and the posterior transitioning to the prior. The posterior on the magnitude does not appreciably broaden beyond 100 h^{-1} Mpc, as sufficient data remain to constrain it at that distance; beyond this, we assume a constant extrapolation. The $\boldsymbol{V}_{\rm ext}(\boldsymbol{r})$ model has higher evidence than the $\boldsymbol{V}_{\rm ext} + \boldsymbol{\Delta}_{\rm ZP}$ model, and we verify that these conclusions are robust to the interpolation scheme and number and locations of the knots.

5 DISCUSSION AND CONCLUSION

We revisited recent claims of Hubble constant anisotropy from direct distance tracers, made with the CF4 TFR W1 subsample by B25, and extended the analysis to the 2MTF and SFI++ TFR samples, as well as to the Pantheon+ and Pantheon+ Lane SN samples (the latter being a reanalysis of Pantheon+ by Z. G. Lane et al. 2024), all restricted to $z < 0.05$. Using a forward-modelling framework that simultaneously fits the distance relation, marginalizes over distances, and models peculiar velocities, we tested for a dipole in the zero-point of the TFR and in the standardized absolute magnitude of SNs. Since we do not apply an absolute distance calibration, this inferred zero-point represents a degenerate combination of the underlying distance relation zero-point and the Hubble constant.

5.1 Summary of results

In the CF4 W1 TFR sample we infer a zero-point dipole amplitude of 0.087 ± 0.019 mag, with more than 2σ preference for a non-zero dipole magnitude in SFI++, Pantheon+, and Pantheon+ Lane as well. Only in the 2MTF sample, which is also the least constraining, do we not infer a significant zero-point dipole. Under the assumption of the underlying zero-point being isotropic and the dipole arising from variation in H_0 alone (see equation 9), the CF4 result corresponds to anisotropies in H_0 of approximately 4.1 per cent; see Table 2 for a summary of the main results. We have confirmed that these results are robust to variations in the Galactic extinction correction. Specifically, we tested both the original $E(B-V)$ values used in CF4 (D. J. Schlegel et al. 1998) and the maps of Y.-K. Chiang (2023) or Planck Collaboration XLVIII (2016), jointly sampling the extinction coefficient $R_{\rm W1}$ in each case.

In the CF4 sample, Bayesian evidence ratios (with a uniform prior on dipole magnitude and isotropic orientation) strongly favour the anisotropic model, with odds of 877 : 1. The SFI++ sample provides only very weak support for anisotropy (3 : 1), while the 2MTF sample very weakly prefers isotropy (2 : 1). For the two Pantheon+ variants, the magnitude covariance matrix prevents a straightforward evidence computation, but the dipole magnitude is more than 3.5σ discrepant from zero.

Such a dipole in H_0 can equivalently be interpreted as a bulk flow that grows linearly with distance. To test if this is the preferred radial dependence we introduce a non-parametric model of radially varying $\boldsymbol{V}_{\rm ext}$ on top of the C15 field. We find that its magnitude increases roughly linearly up to 50 h^{-1} Mpc but then rapidly decays towards zero (see Fig. 8), inconsistent with any H_0 anisotropy. The inferred bulk flow from this radially varying $\boldsymbol{V}_{\rm ext}$ remains consistent with ΛCDM at all radii, decaying with r as expected (see Fig. 7). Thus, while the data favours an anisotropic model over one with a constant $\boldsymbol{V}_{\rm ext}$, a non-linearly-rising bulk flow is preferred which does not challenge ΛCDM. The preference for a dipole in H_0 is caused by the relative inflexibility of that model and limited radial range of the data.

Moreover, we caution that beyond the Galactic extinction already considered, other angular selection effects could mimic a dipole-like signal. Since our analysis constrains only a degenerate combination of the dipole in the distance-tracer zero-point and the Hubble constant, the detected dipole could in principle reflect the former rather than a cosmological origin. To mitigate this, we tested several independent data sets with different calibrations and found consistent hints of a dipole in both the TFR and SNs. For instance, in the case of the CF4 W1 sample, B25 found no evidence of anisotropy in the linewidths, while the use of consistent WISE photometry argues against spurious calibration dipoles. Nevertheless, because the CP is a foundational principle of the standard cosmological model, extraordinary evidence is required for a robust detection, warranting further tests with future data.

Using CF4-like mock catalogues, we showed that a dipole of amplitude 0.087 (the posterior mean from our CF4 inference) is decisively favoured ($\log \mathcal{B} > 2$) once the sample size exceeds ~ 2000. Smaller (mock) samples lack the statistical power to distinguish the models, consistent with the absence of decisive evidence in the 2MTF and SFI++ samples. By contrast, after quality cuts the CF4 W1 sample contains ~ 3200 galaxies, for which the evidence strongly favours the anisotropic model, in line with the mock expectations. As shown in Fig. 6, the evidence in favour of the anisotropic model increases steeply with sample size, highlighting the need for larger samples of well-calibrated distance tracers to achieve higher-significance detections. Extending to higher redshifts provides one route: peculiar velocities become subdominant at larger distances, reducing contamination from local flows. This gain, however, is offset by increased absolute distance uncertainties – typically ~ 20 per cent for the TFR and fundamental plane, and ~ 10 per cent for Type Ia SNs. Our work highlights however that one must be extremely careful in claiming evidence for a fundamental anisotropy, because this may seem to be preferred even if the data really only contains a bulk flow consistent with ΛCDM. Higher redshift data is crucial for breaking this degeneracy because the ΛCDM bulk flow goes to zero at large r

while an anisotropic H_0 causes an effective bulk flow that continues to rise.

5.2 Comparison to previous work

We compare our results for the CF4 W1 TFR sample with those of B25, who also employ the C15 reconstruction to account for peculiar velocities and report a dipole of magnitude 0.063 ± 0.016 mag in the direction of $(\ell, b) = (142 \pm 30^\circ, 52 \pm 10^\circ)$, in good agreement with our measurement, though slightly lower and with smaller uncertainty. They find only weak support for the anisotropic model, reporting odds of 3 : 1 in favour of the zero-point dipole. The discrepancy in the dipole magnitude likely stems from a combination of the following factors:

(i) *Velocity assignment*. While our model queries the velocity field in real space, assigning velocities based on inferred real-space distances, B25 work in redshift space, first mapping the C15 field into redshift space using the prescription of A. Carr et al. (2022), and then assigning peculiar velocities as a function of the observed redshift of each galaxy (while approximately undoing the fiducial $\boldsymbol{V}_{\rm ext}$ and $\beta^\star$ of C15 and applying new ones). The observed redshift together with the peculiar velocity then yields an estimate of the cosmological redshift with a scatter of σ_v. This approach introduces complications in the so-called triple-valued zones (M. A. Strauss & J. A. Willick 1995), where a given observed redshift corresponds to multiple possible distances along the line of sight. In such regions, the mapping from redshift to peculiar velocity becomes non-unique, whereas the real-space assignment used in our framework retains a one-to-one correspondence between distance and peculiar velocity. B25 use the observed data to infer latent model parameters (cosmological redshift), which are then used to predict the apparent magnitude. Our method more straightforwardly forward-models the observables and hence avoids these complications.

(ii) *Lower redshift cut applied to the W1 sample*. While B25 impose $cz_{\rm CMB} > 3000\ {\rm km\,s^{-1}}$, we do not apply such a cut in our main analysis. We verify that applying the same cut yields a dipole magnitude of 0.062 ± 0.022 in the direction $(\ell, b) = (137 \pm 19^\circ, 24 \pm 12^\circ)$, both of which are in a better agreement with B25 than our fiducial result.

(iii) *Photometry-quality cut*. B25 do not apply a photometry-quality cut, thus retaining approximately 1300 more galaxies in the CF4 TFR W1 sample. We test repeating our fiducial analysis without applying the quality cut and recover a dipole of magnitude of 0.066 ± 0.018 mag in the direction $(\ell, b) = (128 \pm 14^\circ, 18 \pm 10^\circ)$. This magnitude is in good agreement with B25, though the direction remains slightly off.

Thus, after matching the lower redshift or photometry-quality cut, we find a dipole in reasonable agreement with B25, though the direction remains mildly inconsistent. The residual differences with B25 are most likely due to the velocity-assignment scheme.

Our work is part of the larger programme of searching for deviations from the CP through anisotropy, which is receiving increasing attention. On the TFR side, besides B25 the studies most directly comparable to ours examine bulk flows in the local Universe. These are strongly correlated with anisotropic Hubble expansion, and indeed if the bulk flow is linearly rising, then locally they are indistinguishable. Claims for such behaviour include R. Watkins et al. (2009), R. Watkins & H. A. Feldman (2015), R. Watkins et al. (2023), the latter in particular using the same CF4 data as us and claiming a bulk flow of $\sim$400 ${\rm km\,s^{-1}}$ at 250 h^{-1} Mpc, in $\sim 5\sigma$ tension with cosmic variance in ΛCDM. However, in contrast to R. Watkins et al. (2023), we find a decaying bulk flow, consistent with ΛCDM expectations.

In our analysis we infer a dipole in the degenerate quantity $a_{\rm TFR} + 5\log h$ (or $M_{\rm SN} + 5\log h$). Interpreted as a dipole in H_0, this implies that along the dipole direction H_0 is enhanced, requiring a negative peculiar velocity to reproduce the same observed redshift. Since the aforementioned studies report a dipole in peculiar velocity, we compare the opposite of their inferred directions with our result. Working in the CMB frame, we compare our inferred dipole with literature estimates in Fig. 4. The dipole direction we find in CF4 agrees with both K. Migkas et al. (2021) and R. Watkins et al. (2023), the latter using a larger set of CF4 data. By contrast, the Pantheon+ dipole direction disagrees with our TFR sample but matches well with the Pantheon+ analysis of F. Sorrenti et al. (2023b), which may reflect the different redshift ranges probed.

On the SN side, several studies find significant or semi-significant anisotropies, (e.g. N. Horstmann et al. 2022; C. Krishnan et al. 2022; W. Rahman et al. 2022; Z. Zhai & W. J. Percival 2022; J. A. Cowell et al. 2023; R. Mc Conville & E. Ó Colgáin 2023; F. Sorrenti et al. 2023a; J. P. Hu et al. 2024; A. Sah et al. 2025; F. Sorrenti et al. 2025). W. Rahman et al. (2022) investigated a possible dipole in the distance moduli of Type Ia SNs using the JLA sample (M. Betoule et al. 2014). They applied the reconstruction of C15 to correct for peculiar velocities and found no evidence for anisotropic expansion. However, other than this study, typically these SNs studies are at higher redshift than we consider; we have cut our sample to be consistent with the TFR data and to remain within the volume covered by the C15 reconstruction ($z < 0.05$). In addition the methodologies differ: instead of fitting dipoles (or multipoles) some studies perform separate analyses in different parts of the sky, which reduces statistical constraining power but enables investigation of more general types of anisotropy. A limiting case of this is the 'hemispherical fitting' method which compares results in one half of the sky to the other. Sometimes the fitting is done in frames other than the CMB frame (e.g. A. Sah et al. 2025), sometimes anisotropy specifically in the direction of the CMB dipole is searched for (e.g. C. Krishnan et al. 2022; Z. Zhai & W. J. Percival 2022), and sometimes more sophisticated frameworks are employed (e.g. A. Heinesen 2023; B. Kalbouneh et al. 2023). The situation is therefore still unclear as to whether SNs evince a bona fide anisotropy (though we find evidence for it with the Pantheon+ samples). We caution however that such results are sensitive to the peculiar velocity field, which is not well known beyond $z \approx 0.05$.

5.3 Disentangling the cosmological dipole from coherent flows

Accounting for peculiar velocities assuming the C15 reconstruction may (partially) absorb a potential cosmological H_0 dipole, making the disentanglement of a cosmological dipole in H_0 from coherent flows challenging. In principle, anisotropies in the expansion rate would imprint on the density field inferred from redshift-space positions, thereby biasing the reconstructed peculiar velocities. While the C15 reconstruction assumes linear theory and ΛCDM to relate the density and velocity fields via the continuity equation and growth rate, it does not constrain the amplitude or anisotropy of the density fluctuations beyond the data itself. Consequently, any non-standard large-scale anisotropy in expansion could propagate into the reconstructed velocity field and be inadvertently subtracted. Nevertheless, the measured clustering amplitude of the 2M++ galaxies remains consistent with ΛCDM expectations. The present analysis therefore searches for any dipolar excess beyond that predicted by C15 based on the 2M++ galaxy positions.

To conclude, using both TFR galaxy samples and Type Ia SN samples, we find (apparent) evidence for anisotropy in the zero-point calibration of the distance-tracer relation. In the CF4 and Pantheon+ sample, the implied variation in the Hubble constant is at the level of 4.1 ± 0.9 and 2.3 ± 0.06 per cent, respectively. Only in the less constraining TFR samples (2MTF and SFI++) do we find no preference between the isotropic and anisotropic model. These results are robust to variations in Galactic extinction but are sensitive to the treatment of peculiar velocities, which are non-negligible at $z < 0.05$. Without accounting for peculiar velocities, the CF4 dipole would be even larger, while the Pantheon+ remains largely unaffected. We have also tested the reanalysis of Pantheon+ by Z. G. Lane et al. (2024), which yields a dipole larger by 1σ compared to the fiducial Pantheon+ data, corresponding to a variation in the Hubble constant of 2.8 ± 0.6 per cent.

While the preference for the zero-point dipole could in principle indicate a genuine anisotropy in the Hubble constant, a more plausible explanation is that it reflects local peculiar velocities not captured by C15 (or systematics) rather than a true cosmological anisotropy. Allowing $\boldsymbol{V}_{\rm ext}$ to vary radially, we find its magnitude rises initially and then rapidly decays to zero, yielding a bulk flow fully consistent with ΛCDM.

These results highlight the importance of continued tests of the local Hubble flow while also underscoring the challenge of disentangling genuine cosmological anisotropies from local peculiar velocities. Properly accounting for peculiar velocities and systematic effects is crucial when using direct distance tracers in cosmological analyses. It is only through such close scrutiny with robust statistical methodology that deviations from the CP may be reliably detected.

ACKNOWLEDGEMENTS

We thank Indranil Banik, Deaglan J. Bartlett, Nicholas Choustikov, Julien Devriendt, Pedro Ferreira, Sebastian von Hausegger, Mike Hudson, Mohamed Rameez, Animesh Sah, Subir Sarkar, and Adrianne Slyz for useful inputs and discussions. We thank Anthony Carr for providing a version of the Pantheon+ covariance matrix with the peculiar velocity contributions removed. We thank Jonathan Patterson for smoothly running the Glamdring Cluster hosted by the University of Oxford, where a part of the data processing was performed. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 693024). The authors would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work.[4]

RS acknowledges financial support from STFC Grant No. ST/X508664/1, the Snell Exhibition of Balliol College, Oxford, and the CCA Pre-doctoral Program. HD is supported by a Royal Society University Research Fellowship (grant no. 211046). GL acknowledges support from the Simons Foundation through the Simons Collaboration on 'Learning the Universe'.

DATA AVAILABILITY

The J. Carrick et al. (2015) reconstruction is available at https://cosmicflows.iap.fr. The public CF4 data are available at https://edd.ifa.hawaii.edu/dfirst.php. The code underlying the article is available at https://github.com/Richard-Sti/CANDEL and all other data will be made available on reasonable request to the authors.

REFERENCES

Abazajian K. N. et al., 2009, ApJS, 182, 543
Aluri P. K., et al., 2023, Class. Quant. Grav., 40, 094001
Bartlett D. J., Desmond H., 2023, Open J. Astrophys., 6, 42
Betoule M. et al., 2014, A&A, 568, A22
Bingham E. et al., 2019, J. Mach. Learn. Res., 20, 1
Boruah S. S., Hudson M. J., Lavaux G., 2020, MNRAS, 498, 2703
Boubel P., Colless M., Said K., Staveley-Smith L., 2024a, MNRAS, 531, 84
Boubel P., Colless M., Said K., Staveley-Smith L., 2024b, MNRAS, 533, 1550
Boubel P., Colless M., Said K., Staveley-Smith L., 2025, J. Cosmol. Astropart. Phys., 2025, 066
Bouchet F. R., Colombi S., Hivon E., Juszkiewicz R., 1995, A&A, 296, 575
Brout D. et al., 2022, ApJ, 938, 110
Carr A., Davis T. M., Scolnic D., Said K., Brout D., Peterson E. R., Kessler R., 2022, Publ. Astron. Soc. Aust., 39, e046
Carrick J., Turnbull S. J., Lavaux G., Hudson M. J., 2015, MNRAS, 450, 317
Chiang Y.-K., 2023, ApJ, 958, 118
Cowell J. A., Dhawan S., Macpherson H. J., 2023, MNRAS, 526, 1482
Dam L., Lewis G. F., Brewer B. J., 2023, MNRAS, 525, 231
Di Valentino E. et al., 2025, Phys. Dark Universe, 49, 101965
Feix M., Branchini E., Nusser A., 2017, MNRAS, 468, 1420
Feldman H. A., Watkins R., Hudson M. J., 2010, MNRAS, 407, 2328
Fitzpatrick E. L., 1999, PASP, 111, 63
Gaztañaga E., Sravan Kumar K., 2024, J. Cosmol. Astropart. Phys., 2024, 001
Gelman A., Rubin D. B., 1992, Stat. Sci., 7, 457
Green G. M., 2018, J. Open Source Softw., 3, 695
Guy J. et al., 2007, A&A, 466, 11
Haynes M. P. et al., 2018, ApJ, 861, 49
Heinesen A., 2023, Phys. Rev. D, 108, 103530
Hellwing W. A., Bilicki M., Libeskind N. I., 2018, Phys. Rev. D, 97, 103519
Hoffman M. D., Gelman A., 2011, preprint (arXiv:1111.4246)
Hoffman Y., Courtois H. M., Tully R. B., 2015, MNRAS, 449, 4494
Hogg D. W., 1999, preprint(astro–ph/9905116)
Hong T. et al., 2019, MNRAS, 487, 2061
Horstmann N., Pietschke Y., Schwarz D. J., 2022, A&A, 668, A34
Hu J. P., Wang Y. Y., Hu J., Wang F. Y., 2024, A&A, 681, A88
Huchra J. P. et al., 2012, ApJS, 199, 26
Jeffreys H., 1939, Theory of Probability. Oxford Univ. Press, Oxford
Jones D. H. et al., 2009, MNRAS, 399, 683
Jones J., Copi C. J., Starkman G. D., Akrami Y., 2023, preprint (arXiv:2310.12859)
Kalbouneh B., Marinoni C., Bel J., 2023, Phys. Rev. D, 107, 023507
Kalbouneh B., Marinoni C., Maartens R., 2024, J. Cosmol. Astropart. Phys., 2024, 069
Kalbouneh B., Marinoni C., Maartens R., Bel J., Santiago J., Clarkson C., Sarma M., Virey J.-M., 2025, preprint (arXiv:2510.02510)
Kelly B. C., Fan X., Vestergaard M., 2008, ApJ, 682, 874
Kessler R., Scolnic D., 2017, ApJ, 836, 56
Kourkchi E., Tully R. B., Neill J. D., Seibert M., Courtois H. M., Dupuy A., 2019, ApJ, 884, 82
Kourkchi E., Tully R. B., Anand G. S., Courtois H. M., Dupuy A., Neill J. D., Rizzi L., Seibert M., 2020a, ApJ, 896, 3
Kourkchi E. et al., 2020b, ApJ, 902, 145
Krishnan C., Mohayaee R., Colgáin E. O., Sheikh-Jabbari M. M., Yin L., 2022, Phys. Rev. D, 105, 063514
Lane Z. G., Seifert A., Ridden-Harper R., Wiltshire D. L., 2024, MNRAS, 536, 1752
Lavaux G., 2016, MNRAS, 457, 172
Lavaux G., Hudson M. J., 2011, MNRAS, 416, 2840
Luongo O., Muccino M., Colgáin E. Ó., Sheikh-Jabbari M. M., Yin L., 2022, Phys. Rev. D, 105, 103510

---

[4] https://doi.org/10.5281/zenodo.22558

Maartens R., Santiago J., Clarkson C., Kalbouneh B., Marinoni C., 2024, J. Cosmol. Astropart. Phys., 2024, 070

Masters K. L., Springob C. M., Haynes M. P., Giovanelli R., 2006, ApJ, 653, 861

Masters K. L., Springob C. M., Huchra J. P., 2008, AJ, 135, 1738

Mc Conville R., Ó Colgáin E., 2023, Phys. Rev. D, 108, 123533

McEwen J. D., Wallis C. G. R., Price M. A., Spurio Mancini A., 2021, preprint (arXiv:2111.12720)

Migkas K., Pacaud F., Schellenberger G., Erler J., Nguyen-Dang N. T., Reiprich T. H., Ramos-Ceja M. E., Lovisari L., 2021, A&A, 649, A151

Nusser A., 2016, MNRAS, 455, 178

Nusser A., Davis M., 2011, ApJ, 736, 93

Pandya A. et al., 2024, A&A, 691, A355

Peery S., Watkins R., Feldman H. A., 2018, MNRAS, 481, 1368

Peterson E. R. et al., 2022, ApJ, 938, 112

Phan D., Pradhan N., Jankowiak M., 2019, preprint (arXiv:1912.11554)

Piras D., Polanska A., Mancini A. S., Price M. A., McEwen J. D., 2024, Open J. Astrophys., 7, 73

Planck Collaboration XLVIII, 2016, A&A, 596, A109

Planck Collaboration I, 2020a, A&A, 641, A1

Planck Collaboration VI, 2020b, A&A, 641, A6

Polanska A., Price M. A., Piras D., Spurio Mancini A., McEwen J. D., 2025, Open J. Astrophys., 8, 156

Rahman W., Trotta R., Boruah S. S., Hudson M. J., van Dyk D. A., 2022, MNRAS, 514, 139

Riess A. G. et al., 2022, ApJ, 934, L7

Sah A., Rameez M., Sarkar S., Tsagas C., 2025, Eur Phys. J. C, 85, 596

Said K., Colless M., Magoulas C., Lucey J. R., Hudson M. J., 2020, MNRAS, 497, 1275

Sarma M., Marinoni C., Kalbouneh B., Clarkson C., Maartens R., 2025, preprint (arXiv:2510.03517)

Schlafly E. F., Finkbeiner D. P., 2011, ApJ, 737, 103

Schlegel D. J., Finkbeiner D. P., Davis M., 1998, ApJ, 500, 525

Scolnic D. et al., 2022, ApJ, 938, 113

Scolnic D. M. et al., 2018, ApJ, 859, 101

Scrimgeour M. I. et al., 2016, MNRAS, 455, 386

Secrest N., von Hausegger S., Rameez M., Mohayaee R., Sarkar S., 2025, preprint (arXiv:2505.23526)

Secrest N. J., von Hausegger S., Rameez M., Mohayaee R., Sarkar S., 2022a, ApJ, 937, L31

Seifert A., Lane Z. G., Galoppo M., Ridden-Harper R., Wiltshire D. L., 2025, MNRAS, 537, L55

Skrutskie M. F. et al., 2006, AJ, 131, 1163

Sorrenti F., Durrer R., Kunz M., 2023a, J. Cosmol. Astropart. Phys., 2023, 054

Sorrenti F., Durrer R., Kunz M., 2023b, J. Cosmol. Astropart. Phys., 2023, 054

Sorrenti F., Durrer R., Kunz M., 2025, J. Cosmol. Astropart. Phys., 2025, 013

Springob C. M., Masters K. L., Haynes M. P., Giovanelli R., Marinoni C., 2007, ApJS, 172, 599

Staveley-Smith L. et al., 1998, AJ, 116, 2717

Stiskalek R., Desmond H., Devriendt J., Slyz A., Lavaux G., Hudson M. J., Bartlett D. J., Courtois H. M., 2025a, preprint (arXiv:2502.00121)

Stiskalek R., Desmond H., Tsaprazi E., Heavens A., Lavaux G., McAlpine S., Jasche J., 2025b, preprint (arXiv:2509.09665)

Strauss M. A., Willick J. A., 1995, Phys. Rep., 261, 271

Tripp R., 1998, A&A, 331, 815

Tully R. B., Fisher J. R., 1977, A&A, 54, 661

Tully R. B. et al., 2023, ApJ, 944, 94

Wang L., Steinhardt P. J., 1998, ApJ, 508, 483

Watkins R., Feldman H. A., 2015, MNRAS, 447, 132

Watkins R., Feldman H. A., Hudson M. J., 2009, MNRAS, 392, 743

Watkins R. et al., 2023, MNRAS, 524, 1885

Westover M., 2007, PhD thesis, Harvard Univ.

Whitford A. M., Howlett C., Davis T. M., 2023, MNRAS, 526, 3051

Yahil A., Strauss M. A., Davis M., Huchra J. P., 1991, ApJ, 372, 380

Yuan H. B., Liu X. W., Xiang M. S., 2013, MNRAS, 430, 2188

Zhai Z., Percival W. J., 2022, Phys. Rev. D, 106, 103527

APPENDIX A: FLOW MODEL OF S25

A1 Tully–Fisher flow model

In this section, we summarize the flow model used in this analysis, which is based on our previous work in S25. For a more detailed description, we refer the reader to that study. We employ a forward model that jointly calibrates the TFR (and its dipole) while also calibrating the local Universe reconstruction (which accounts for smaller scale peculiar velocity modes) and other model parameters. For each galaxy, we have observations of its redshift, $z_{\rm CMB}$ (converted to the CMB frame) along with its uncertainty, σ_z. We define its value multiplied by the speed of light as $\sigma_{cz} \equiv c\sigma_z$. The total redshift of a galaxy in the CMB frame is given by

$$1 + z_{\rm CMB} = (1 + z_{\rm cosmo})\left(1 + z_{\rm pec}\right), \tag{A1}$$

where $z_{\rm cosmo}$ is the redshift due to cosmic expansion, and $z_{\rm pec} = V_{\rm pec}/c$ represents the redshift contribution from the galaxy's line-of-sight peculiar velocity, $V_{\rm pec}$. In a flat ΛCDM universe dominated by non-relativistic matter and dark energy, the cosmological redshift, $z_{\rm cosmo}$, is related to the comoving distance, r, by (e.g. D. W. Hogg 1999)

$$r(z_{\rm cosmo}) = \frac{c}{H_0}\int_0^{z_{\rm cosmo}} \frac{{\rm d}z'}{\sqrt{\Omega_{\rm m}(1+z')^3 + 1 - \Omega_{\rm m}}}, \tag{A2}$$

where $\Omega_{\rm m}$ is the matter density parameter. The velocity field, $\boldsymbol{v}(\boldsymbol{r})$, is modelled under the single-flow approximation, so that the line-of-sight peculiar velocity is

$$V_{\rm pec} = (\beta\boldsymbol{v}(\boldsymbol{r}) + \boldsymbol{V}_{\rm ext})\cdot\hat{\boldsymbol{r}}, \tag{A3}$$

where $\hat{\boldsymbol{r}}$ is the galaxy line-of-sight unit vector, and $\boldsymbol{V}_{\rm ext}$ accounts for external flows originating beyond the reconstruction volume. The parameter $\beta^{\star}$ is a calibration factor to scale the velocities predicted by C15 and is a function of both cosmology and the galaxy bias. When assuming no underlying velocity field, we effectively set $\boldsymbol{v} = 0$, modelling the flow solely as $\boldsymbol{V}_{\rm ext}$. A third parameter, σ_v, which we introduce later, captures small-scale velocity dispersion not accounted for by the reconstruction.

For each galaxy we observe the apparent magnitude $m_{\rm obs}$ with uncertainty $\sigma_{\rm m}$ and the linewidth $\eta_{\rm obs}$ with uncertainty σ_η (which is a distance-independent observable). These quantities constrain the galaxy distance. In S25 we introduce two latent parameters per galaxy to be inferred: the distance r and the true linewidth $\eta_{\rm true}$. The TFR relates $\eta_{\rm true}$ to the absolute magnitude M via equation (2). Converting r to a distance modulus μ yields a predicted apparent magnitude, $m_{\rm pred} = \mu(r) + M(\eta_{\rm true})$, with the distance modulus being

$$\mu = 5\log\frac{d_{\rm L}}{\rm Mpc} + 25, \tag{A4}$$

where the luminosity distance $d_{\rm L}$ is related to the comoving distance r by

$$d_{\rm L} = (1 + z_{\rm cosmo})r, \tag{A5}$$

assuming a flat ΛCDM universe.

We jointly infer the velocity field calibration parameters, $(\boldsymbol{V}_{\rm ext}, \beta, \sigma_v)$, the distance indicator parameters, $(a_{\rm TFR}, b_{\rm TFR}, c_{\rm TFR}, \sigma_{\rm int})$, the mean $\hat{\eta}$ and standard deviation w_η of the $\eta_{\rm true}$ prior, and potentially the zero-point dipole $\boldsymbol{\Delta}_{\rm ZP}$. Assuming independent sources, S25 formulate the likelihood of the observed redshift, magnitude, and linewidth given the distance, true linewidth, and the set of model parameters $\boldsymbol{\theta}$ as

$$\begin{aligned}
&\mathcal{L}(z_{\rm CMB}, m_{\rm obs}, \eta_{\rm obs} \mid \boldsymbol{\theta}, r, \eta_{\rm true}) \\
&\quad = \mathcal{N}\left(cz_{\rm CMB}; cz_{\rm pred}, \sqrt{\sigma_v^2 + \sigma_{cz_{\rm CMB}}^2}\right) \\
&\quad \times \mathcal{N}\left(m_{\rm obs}; m_{\rm pred}, \sqrt{\sigma_{\rm int}^2 + \sigma_m^2}\right) \\
&\quad \times \frac{\mathcal{N}(\eta_{\rm obs}; \eta_{\rm true}, \sigma_\eta)}{p(S = 1 \mid \hat{\eta}, w_\eta)},
\end{aligned} \tag{A6}$$

where $p(S = 1 \mid \hat{\eta}, w_\eta)$ is the expected fraction of retained sources given a truncation in $\eta_{\rm obs}$. Strictly speaking, $p(S = 1 \mid \hat{\eta}, w_\eta)$ is not part of the data likelihood. It should multiply the product of the per-sample likelihoods and the prior as $\left[p(S = 1 \mid \hat{\eta}, w_\eta)\right]^{-n}$, where n is the number of galaxies in the sample. For notational convenience, however, we absorb it into the likelihood (see B. C. Kelly et al. 2008 for more details). The first term in equation (A6) is the likelihood of the observed redshift given the predicted value, which depends on the inferred distance and peculiar velocity, while the second term is the likelihood of the observed magnitude given the predicted apparent magnitude. Furthermore, we have that

$$\begin{aligned}
p(S = 1 \mid \hat{\eta}, w_\eta) = &\iint {\rm d}\eta_{\rm obs}\, {\rm d}\eta_{\rm true}\; p(S = 1 \mid \eta_{\rm obs}) \\
&\times \mathcal{L}(\eta_{\rm obs} \mid \eta_{\rm true})\, \pi(\eta_{\rm true} \mid \hat{\eta}, w_\eta),
\end{aligned} \tag{A7}$$

where $p(S = 1 \mid \eta_{\rm obs})$ is a binary detection indicator between $\eta_{\rm min}$ and $\eta_{\rm max}$,

$$p(S = 1 \mid \eta_{\rm obs}) = \begin{cases} 1 & \text{if} \quad \eta_{\rm min} < \eta_{\rm obs} < \eta_{\rm max}, \\ 0 & \text{otherwise.} \end{cases} \tag{A8}$$

$\mathcal{L}(\eta_{\rm obs} \mid \eta_{\rm true})$ denotes the Gaussian likelihood of the observed given the true linewidth, and $\pi(\eta_{\rm true} \mid \hat{\eta}, w_\eta)$ is the Gaussian prior on the true linewidth. Given these assumptions, it can be shown that

$$p(S = 1 \mid \eta_{\rm obs}) = \Phi\left(\frac{\eta_{\rm max} - \hat{\eta}}{\sqrt{\sigma_\eta^2 + w_\eta^2}}\right) - \Phi\left(\frac{\eta_{\rm min} - \hat{\eta}}{\sqrt{\sigma_\eta^2 + w_\eta^2}}\right), \tag{A9}$$

where $\Phi(x)$ is the cumulative density function of the standard normal distribution, defined as

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-t^2/2} {\rm d}t. \tag{A10}$$

We assume only a lower threshold in $\eta_{\rm obs}$ in CF4 and SFI++ , while in 2MTF we impose both a lower and upper threshold. We marginalize over r and $\eta_{\rm true}$ as

$$\begin{aligned}
&\mathcal{L}(z_{\rm CMB}, m_{\rm obs}, \eta_{\rm obs} \mid \boldsymbol{\theta}) \\
&\quad = \iint \mathcal{L}(z_{\rm CMB}, m_{\rm obs}, \eta_{\rm obs} \mid \boldsymbol{\theta}, r, \eta_{\rm true}) \\
&\quad \times \pi(r \mid \boldsymbol{\theta})\, \pi(\eta_{\rm true} \mid \boldsymbol{\theta})\, {\rm d}r\, {\rm d}\eta_{\rm true},
\end{aligned} \tag{A11}$$

where $\pi(r \mid \boldsymbol{\theta})$ is defined in equation (6). Rather than modelling the full HI selection of the TFR sample, which is the primary reason why the TFR samples do not extend to higher redshifts, we follow G. Lavaux (2016) and adopt an effective treatment by setting the distance prior to equation (6), which incorporates both the homogeneous and inhomogeneous Malmquist bias. We may either sample r and $\eta_{\rm true}$ for each galaxy directly with an HMC sampler, or marginalize over them numerically at each MCMC step. The former is computationally faster, while the latter yields a lower dimensional parameter space suitable for evidence computation. Because evidence values are central to this work, we adopt the latter approach and evaluate the two-dimensional integral numerically on a grid in r and $\eta_{\rm true}$ at each MCMC step. We define a fixed radial distance grid ranging from 0.001 to 201 h^{-1} Mpc with a step size of 0.5 h^{-1} Mpc, which is sufficient given that the C15 field is smoothed on scales of 4 h^{-1} Mpc. For $\eta_{\rm true}$, we adopt an adaptive binning scheme. The Gaussian likelihood term $\mathcal{L}(\eta_{\rm obs} \mid \eta_{\rm true})$ determines the region of $\eta_{\rm true}$ that carries non-negligible probability mass. Accordingly, for each source we define a grid spanning $\eta_{\rm obs} \pm 5\sigma_\eta$, discretized into 31 equally spaced steps, and evaluate the likelihood over a 402 × 31 grid at each MCMC iteration. We then marginalize over this two-dimensional grid using Simpson's rule. This computation is performed on GPUs using our `JAX`-based implementation.[5] To verify that the numerical integration does not introduce bias, we compared the results to a model in which the latent parameters $\eta_{\rm true}$ were explicitly sampled and found the posteriors to be consistent.

We compute the model evidence, $\mathcal{Z}$, defined as the integral of the likelihood weighted by the prior over the parameter space,

$$\mathcal{Z} \equiv \int {\rm d}\boldsymbol{\theta} \mathcal{L}(D \mid \boldsymbol{\theta}) \pi(\boldsymbol{\theta}), \tag{A12}$$

where D represents the data and $\boldsymbol{\theta}$ the model parameters. The ratio of pieces of evidence between two models, known as the Bayes factor, quantifies the relative statistical support for one model over another, assuming equal prior model probabilities.

APPENDIX B: CF4 TFR W1 FULL POSTERIOR

In Fig. B1, we present the posterior distribution over model parameters inferred from the CF4 TFR W1 data, with and without a dipole in the zero-point. The parameter set includes the isotropic TFR zero-point, slope, and curvature ($a_{\rm TFR}$, $b_{\rm TFR}$, $c_{\rm TFR}$); the velocity field calibration factor $\beta^\star$; the intrinsic TFR scatter ($\sigma_{\rm int}$); the redshift scatter (σ_v); the external flow parameters ($V_{\rm ext}$, $\ell_{\rm ext}$, $b_{\rm ext}$); the zero-point dipole parameters ($\Delta_{\rm ZP}$, $\ell_{\Delta_{\rm ZP}}$, $b_{\Delta_{\rm ZP}}$); the linewidth hyperprior mean and standard deviation ($\hat{\eta}$, w_η); and the distance prior parameters (R, n, p).

We find a mild degeneracy between the zero-point dipole amplitude and the magnitude of the external flow: a larger external flow can compensate for a smaller zero-point dipole, and vice versa. Comparing the model with a zero-point dipole to the isotropic case, we find a slight shift in the zero-point monopole, while the posteriors for $b_{\rm TFR}$, $c_{\rm TFR}$, α, $\beta^\star$, $\sigma_{\rm int}$, $\hat{\eta}$, and w_η remain unchanged. As expected, introducing the dipole shifts the posterior on $\boldsymbol{V}_{\rm ext}$, affecting both its magnitude and Galactic latitude.

---

[5] https://github.com/jax-ml/jax

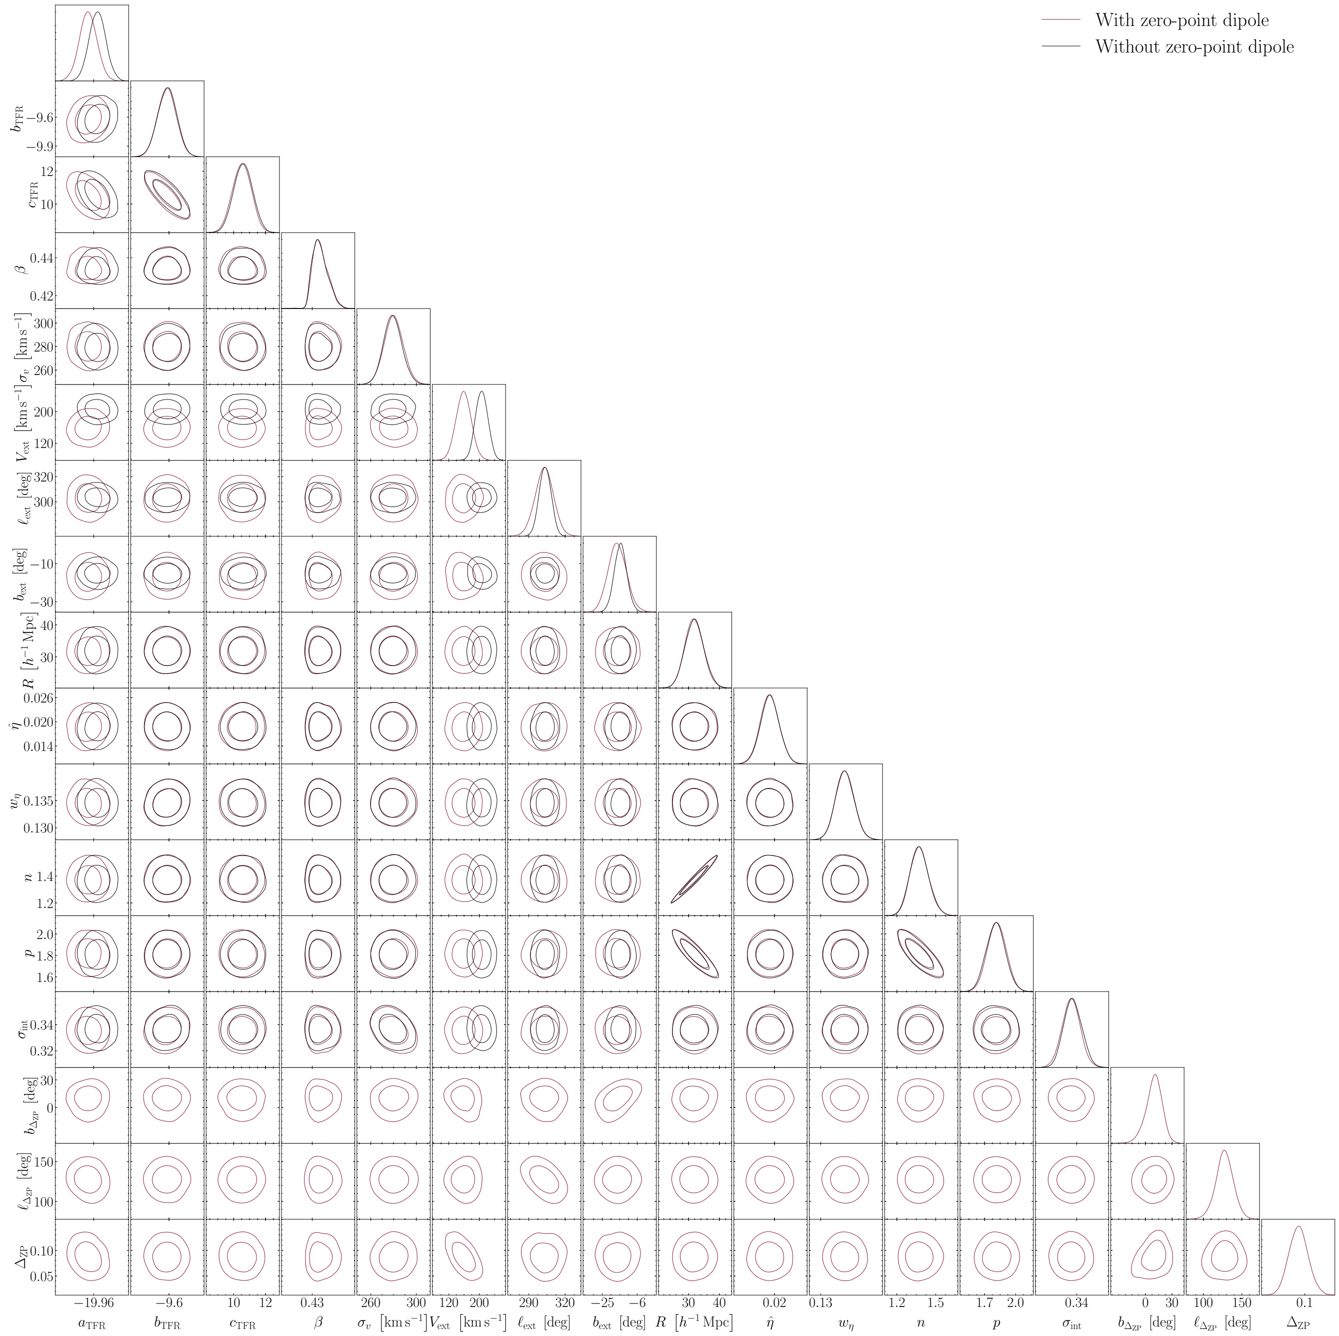

Figure B1. Posterior distributions for the model parameters inferred from the CF4 TFR W1 data, comparing the zero-point dipole model (red) to the isotropic model (black). Parameters shown include the TFR calibration (a_{TFR}, b_{TFR}, c_{TFR}), velocity field calibration factor ($\beta^\star$), scatter terms (σ_{int}, σ_v), external flow parameters (V_{ext}, ℓ_{ext}, b_{ext}), linewidth hyperparameters ($\hat{\eta}$, w_η), distance prior hyperparameters (R, n, p), and zero-point dipole parameters (Δ_{ZP}, $\ell_{\Delta_{\mathrm{ZP}}}$, $b_{\Delta_{\mathrm{ZP}}}$). We observe a mild degeneracy between the magnitude of the zero-point dipole and the external flow, and find that introducing a dipole shifts the inferred $\boldsymbol{V}_{\mathrm{ext}}$ while leaving other parameters largely unchanged. Contours denote 1σ and 2σ credible regions.

This paper has been typeset from a TEX/LATEX file prepared by the author.