MNRAS 000, 1–21 (2026) Preprint 23 July 2026 Compiled using MNRAS LATEX style file v3.3

We benchmark interpretable halo summaries against the full 6D information carried by the LMC's disturbance.

LMC-induced Perturbations in the Milky Way Halo II: Bridging Field-level Inference and Summary-level Simulation-Based Inference

Yanjun Sheng,[1]★ Yuan-Sen Ting,[2,3,4] Tomasz Różański,[1] Xiang-Xiang Xue,[5,6] Roland M. Crocker,[1] Jiashu Pan[7]

[1]Research School of Astronomy and Astrophysics, Australian National University, Canberra, ACT 2611, Australia
[2]Department of Astronomy, The Ohio State University, Columbus, OH 43210, USA
[3]Center for Cosmology and AstroParticle Physics (CCAPP), The Ohio State University, Columbus, OH 43210, USA
[4]Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany
[5]National Astronomical Observatories, Chinese Academy of Sciences, Beijing 100101, People's Republic of China
[6]Institute for Frontiers in Astronomy and Astrophysics, Beijing Normal University, Beijing 102206, People's Republic of China
[7]School of Engineering, Westlake University, Hangzhou, Zhejiang, 310030, People's Republic of China

Accepted XXX. Received YYY; in original form ZZZ

arXiv:2607.20006v1 [astro-ph.GA] 22 Jul 2026

ABSTRACT

The gravitational interaction between the Milky Way (MW) and the Large Magellanic Cloud (LMC) drives the outer halo into dynamical disequilibrium, imprinting the masses and structural parameters of both galaxies onto the 6D phase-space distribution of halo tracers. This signal has been characterised with summary statistics ranging from low-order velocity moments to basis function expansions, yet how much information these summaries discard, and whether they are complementary, remains unclear. We address these questions by comparing a field-level likelihood benchmark with physically interpretable summaries for constraining $(M_{\rm MW}, M_{\rm LMC}, c, q)$, where c and q are the MW halo concentration and flattening. A Conditional Flow Matching (CFM) model trained on the HaloDance N-body suite provides an exact likelihood at a held-out fiducial point; for 5,000 tracers in 30–120 kpc it tightens marginal constraints by factors of 2.5–9.9 over an all-sky velocity-moment forecast. We then expand the halo density and velocity fields in a multipole basis-function expansion (BFE) and compress the coefficients with the Massive Optimised Parameter Estimation and Data compression (MOPED) algorithm into four parameter-sensitive summaries that preserve their Fisher information. A variational mutual-information analysis shows that the BFE+MOPED summaries and the velocity moments are complementary, so we combine them into a joint 19-dimensional vector as our primary inference pipeline: it tightens the marginal constraints by up to 15 per cent over BFE+MOPED alone and by 30–71 per cent over velocity moments alone, reaching within a factor of 1.3–2.9 of the field-level benchmark. We thus establish a physically interpretable summary-level route to MW+LMC inference alongside the field-level benchmark that bounds its information content.

Combining two imperfect summaries works because they discard different parts of the signal.

Key words: Galaxy: halo – Galaxy: kinematics and dynamics – Magellanic Clouds – methods: statistical – methods: numerical

1 INTRODUCTION

The Large Magellanic Cloud (LMC) is the most massive satellite galaxy of the Milky Way (MW), with a virial mass estimated at 0.8–2.5×10^{11} $M_\odot$, representing 10–25 per cent of the MW's total mass (Bland-Hawthorn & Gerhard 2016; Shipp et al. 2021; Fushimi et al. 2024). Independent constraints on this mass come from its rotation curve (van der Marel & Kallivayalil 2014), three-dimensional kinematics of its globular clusters (Watkins et al. 2024), deflections of MW stellar streams (Erkal et al. 2019; Shipp et al. 2021; Vasiliev et al. 2021; Koposov et al. 2023), the cosmic-web timing argument (Peñarrubia et al. 2016; Chamberlain et al. 2025), and MW halo stars entrained in the LMC's gravitational wake (Petersen et al. 2024; Cavieres et al. 2025). The MW–LMC interaction produces dynamical signatures whose amplitudes depend on the mass distributions of both systems (Gómez et al. 2015; Laporte et al. 2018a,b; Garavito-Camargo et al. 2019; Petersen & Peñarrubia 2020; Erkal et al. 2021; Vasiliev 2024; Conroy et al. 2021; Sheng et al. 2024, 2025; Brooks et al. 2025, 2026a,b; Darragh-Ford et al. 2025).

all-sky moments

same halo catalogue

angular location lost

BFE + MOPED

some angular structure retained

The orbital history of the LMC has been fundamentally revised over the past two decades. Hubble Space Telescope proper motions, which show the LMC's tangential velocity approaching the MW escape speed (Kallivayalil et al. 2006, 2013; Besla et al. 2007, 2018), replaced the once-canonical picture of short-period, multi-orbit trajectories (e.g., Tremaine 1976; Murai & Fujimoto 1980; Lin & Lynden-Bell 1982) with a first-infall scenario in which the LMC is near pericentre at ~50 kpc moving faster than 300 km s^{-1} (Kallivayalil et al. 2013; Vasiliev et al. 2021; Shipp et al. 2024; Lucchini et al. 2025). A second passage ~100 kpc roughly 1 Gyr ago nonetheless remains viable given uncertainties in the MW potential (Vasiliev 2024; Suzuki et al. 2026). Independent of orbital history, the LMC's infall produces two principal dynamical effects in the MW halo: (i) a reflex motion of the inner Galaxy relative to the outer halo, manifesting as a dipole in the stellar density and radial velocity distribution

★ E-mail: Yanjun.Sheng@anu.edu.au

© 2026 The Authors

The LMC leaves both a global reflex dipole and a localized wake, so no single simple statistic captures everything.

(Petersen & Peñarrubia 2020, 2021; Erkal et al. 2021), and (ii) a dynamical friction wake, which is a localised overdensity of dark matter and stars trailing the LMC's past trajectory (Chandrasekhar 1943; Garavito-Camargo et al. 2019; Tamfal et al. 2021; Foote et al. 2023). These effects depend not only on the mass of the LMC but also on the shape, concentration, and velocity anisotropy of the MW halo, making the inference problem intrinsically coupled.

Observations have established both signatures in the MW outer halo. The reflex motion appears as a systematic Galactocentric velocity dipole in halo tracers beyond 50 kpc (Erkal et al. 2021; Petersen & Peñarrubia 2021; Conroy et al. 2021; Byström et al. 2025; Chandra et al. 2025; Li et al. 2026), while the dynamical friction wake appears as a localized stellar overdensity trailing the LMC's past orbit, including the Pisces Plume (Belokurov et al. 2019; Conroy et al. 2021; Fushimi et al. 2024; Cavieres et al. 2025). These detections show that the MW's outer halo is far from equilibrium and that its 6D phase-space structure carries a coupled imprint of the LMC's gravitational influence.

The dipole and wake depend differently on MW structure, which makes the inference coupled rather than merely noisy.

On the theoretical side, idealised N-body simulations have characterised both the local wake and the global density dipole, and understood the origin of the latter (Gómez et al. 2015; Garavito-Camargo et al. 2019, 2021; Tamfal et al. 2021; Foote et al. 2023; Vasiliev 2023; Sheng et al. 2024). Cosmological zoom-in runs confirm that the reflex motion and wake signatures persist in a full ΛCDM context with realistic merger histories (Darragh-Ford et al. 2025), and dedicated N-body suites have begun to systematically explore the multi-dimensional parameter space of MW mass, LMC mass, halo shape, and orbital configuration (Sheng et al. 2025; Brooks et al. 2025; Garver et al. 2026).

global reflex response

outer halo lags

LMC

inner MW shifts

To extract physical insight from these simulations, multipole expansions of the density and potential fields (typically implemented as biorthogonal basis function expansions, or BFEs) decompose the perturbed halo into distinct angular and radial channels, separating the $\ell = 1$ reflex motion from the $\ell \geq 2$ local wake (Garavito-Camargo et al. 2021; Lilleengen et al. 2023; Darragh-Ford et al. 2025; Foote et al. 2026). The same biorthogonal projection underpins the matrix method of linear response theory, which has shown analytically that the reflex motion depends only on the MW potential while the wake is sensitive to the halo velocity anisotropy (Rozier et al. 2022). Spherical harmonic decomposition of halo tracers offers a complementary route, applied to simulations (Cunningham et al. 2020). These frameworks together show that the MW's response to the LMC decomposes into interpretable multipole components, each encoding different aspects of the underlying mass distributions and orbital configuration of halo stars.

In the moving MW frame, the outer halo appears to move oppositely.

Two limitations remain. First, many analyses compare data with a small number of fiducial simulations rather than sampling the multi-dimensional parameter posterior. Second, the summaries used to compress this signal, most commonly low-order velocity moments (all-sky mean velocities and dispersions in broad radial shells), but also basis function expansions of the density and velocity fields, are adopted without knowing how much phase-space information each retains or whether different summaries are complementary. Low-order velocity moments in particular are interpretable but discard the localized structure of the reflex motion and wake. This gap has persisted partly because no absolute upper limit on the phase-space information was available against which any summary could be measured. The HaloDance suite (Sheng et al. 2025) and neural posterior estimation (Brooks et al. 2026a) address the parameter-space search, but still rely on low-order velocity moments without a benchmark for the total information content.

In this work we use field-level inference as a benchmark for the information content of the 6D phase-space distribution of halo tracers, and compare it against summary statistics whose physical channels can be inspected. We train a Conditional Flow Matching (CFM) generative model directly on the raw particle coordinates of the HaloDance simulations, this yields a likelihood for a tracer catalogue at any proposed parameter point and defines a reference information limit for this simulation setup. We then quantify how much of that information is retained by standard velocity moments and by a multipole BFE of the halo density and velocity fields compressed with the Massive Optimised Parameter Estimation and Data compression (MOPED) algorithm, and use a variational mutual-information analysis to show that the two summary sets are complementary. Combining them yields our primary parameter constraints, while comparison against the field-level likelihood bounds how much information the interpretable summaries leave unrecovered.

2 DATASETS AND EMULATION

The inference setup combines the HaloDance suite of high-resolution N-body simulations, which define the MW–LMC parameter space, with a Conditional Flow Matching (CFM) generative model trained on them. The CFM interpolates between discrete simulation points and provides a likelihood for 6D phase-space samples.

local wake

2.1 HaloDance simulation suite

The analysis is built on the HaloDance suite of N-body simulations, which models the gravitational interaction between the MW and LMC and the resulting perturbation to the MW stellar halo. The MW model consists of an NFW dark matter halo (Navarro et al. 1997), a Miyamoto–Nagai stellar disk (Miyamoto & Nagai 1975), and a Hernquist bulge (Hernquist 1990). The LMC is modelled as a spherical Hernquist dark matter halo. Initial galaxy models are generated with the IC code (Yurin & Springel 2014) and evolved in isolation for 3 Gyr with GADGET-4 (Springel et al. 2021) to ensure dynamical stability before the interaction begins.

MW halo stars

motion

local wake

The local wake is a trailing excess of MW halo stars.

The suite comprises 2,976 high-resolution simulations, each with 10^7 particles and a spatial resolution of 0.5 kpc. The simulations span a four-dimensional parameter space using Latin Hypercube Sampling:

- MW virial mass, $M_{\rm MW} \in [0.5, 2.0] \times 10^{12}$ M$_\odot$ (fiducial: 0.7×10^{12} M$_\odot$),
- LMC virial mass, $M_{\rm LMC} \in [0.7, 2.1] \times 10^{11}$ M$_\odot$ (fiducial: 1.5×10^{11} M$_\odot$),
- MW halo concentration, $c \in [5, 15]$ (fiducial: 9.415),
- MW halo flattening, $q \in [0.5, 1.5]$ (fiducial: 1.0).

We consider two tracer velocity anisotropy profiles β:

- Isotropic, $\beta = 0$ (1,848 models; adopted as fiducial),
- Radially varying, $\beta(r) = -0.15 - 0.2\,\alpha(r)$ (1,000 models), where $\alpha(r) = d\ln\rho/d\ln r$ is the local density slope (Hansen & Moore 2006).

Each simulation includes the LMC's infall and completes a single pericentric passage, and parameters are chosen to span a range of virial radii while the final 5 Gyr are retained, ensuring all models follow a first-infall scenario without imposing it by hand. Initial orbital coordinates for each simulation are determined via a multi-layer perceptron neural network trained to map initial phase-space coordinates to the observed present-day configuration (Kallivayalil et al. 2013). Each system is evolved for 2 Gyr with GADGET-4.

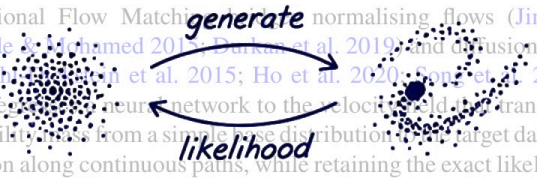
HaloDance varies four parameters systematically; that breadth is essential for learning degeneracies rather than one preferred simulation.

For each simulation, we extract the 6D phase-space coordinates

A conditional flow learns the parameter-dependent 6D halo distribution and turns each tracer catalogue into an evaluable likelihood.

of dark matter particles within 30–120 kpc of the MW centre (i.e., the disk centre) as the raw data for our inference pipeline. We refer the reader to Sheng et al. (2025) for a detailed description of the simulation setup, parameter grid, and convergence tests.

2.2 Conditional Flow Matching model



Conditional Flow Matching (CFM) belongs to the family of simulation-based inference methods that learn a flexible density model directly from forward simulations, without requiring explicit likelihood evaluations during training. It is related to continuous normalising flows (Jimenez Rezende & Mohamed 2015; Durkan et al. 2019) and diffusion models (Sohl-Dickstein et al. 2015; Ho et al. 2020; Song et al. 2020). CFM regresses a neural network to the velocity field that transports probability mass from a simple base distribution to the target data distribution along continuous paths, while retaining the exact likelihood computation of continuous normalising flows. In astrophysics, flow matching has recently been applied to galaxy property inference from imaging (Yunus et al. 2025) and inference of the Galactic potential from stellar streams (Viterbo & Buck 2026). Here the CFM model learns the conditional phase-space distribution of MW halo particles, $p(\mathbf{x} \mid \boldsymbol{\theta})$, where $\mathbf{x}$ denotes the 6D phase-space vector. Specifically, CFM parameterises a time-dependent vector field with a neural network to define a continuous trajectory mapping a simple 6D Gaussian base distribution to the parameter-dependent phase-space distribution of the halo tracers. Because this mapping is continuous and invertible, it allows both mock stellar distributions to be generated by flowing noise particles forward and exact probability densities to be evaluated for observed stellar phase-space coordinates under the model.

Figure 1 summarises this construction, from the Gaussian base particles to HaloDance phase-space particles and the forward and backward probability-flow operations. The learned velocity field $\mathbf{u}_t(\mathbf{x} \mid \boldsymbol{\theta})$ is conditioned on the MW–LMC parameters $\boldsymbol{\theta} = (M_{\rm MW}, M_{\rm LMC}, c, q)$, and its input is a 6D phase-space vector $\mathbf{x} = (x, y, z, v_x, v_y, v_z)$ representing position and velocity in Galactocentric coordinates, augmented with radial features to provide rotation-invariant spatial information.

The model is trained on a 90 per cent split of the isotropic HaloDance simulations, with 10 per cent held out for validation. The training process uses optimal-transport conditional flow matching (OT-CFM) to learn a straight-line probability path mapping Gaussian noise to the physical phase space. To ensure training stability and accuracy, we implement a curriculum weighting over the parameter space. Once trained, the likelihood of a given particle sample is computed using the continuous change-of-variables formula. The full technical specifications of the neural network architecture, parameter embeddings, training hyperparameters, and probability flow integration are detailed in Appendix A.

We first verify that the CFM model generalises beyond its training set. Figure 2 compares the phase-space distribution of the six coordinates (x, y, z, v_x, v_y, v_z) for 5×10^5 CFM-generated particles (blue) and an equal number of particles from a held-out fiducial simulation ($M_{\rm MW} = 0.7 \times 10^{12}\,{\rm M}_\odot$, $M_{\rm LMC} = 1.5 \times 10^{11}\,{\rm M}_\odot$, $c = 9.415$, $q = 1.0$, $\beta(r) = 0$) that was not included in HaloDance training (grey). Both samples are restricted to the radial range 30–120 kpc, and CFM sampling uses a fixed-step RK4 ODE solver with 128 steps. The one-dimensional marginals on the diagonal and the pairwise distributions below them overlap between the two samples, with no visible systematic shifts or shape discrepancies, showing that the CFM matches the bulk phase-space distribution at a parameter point outside its training set. This check probes only low-order structure, and the benchmark in Section 2.3 shows that the CFM also captures the localised, higher-order structure, such as the LMC wake, that carries much of the parameter information.

Holding out the fiducial simulation prevents the benchmark from becoming a memorization test.

The same model can also assign a likelihood to an observed particle catalogue at any proposed $\boldsymbol{\theta}$. For a single particle, the CFM log-likelihood follows from the continuous normalising-flow interpretation and is computed by integrating the probability flow ODE backward from the data to the Gaussian base distribution, keeping track of the instantaneous change of variables. For a sample of N independent particles, the total log-likelihood is the sum of the per-particle log-likelihoods. The full mathematical formulation of the change-of-variables ODE, coordinate normalisation, divergence estimation via the Hutchinson trace estimator, and the ODE solver settings are detailed in Appendix A. This likelihood is what allows the CFM to serve as a benchmark for the information content of the raw 6D phase-space distribution.

The 11-by-11-by-11-by-11 grid is expensive, but it avoids sampler-convergence ambiguity for this local benchmark.

2.3 Field-level inference benchmark

We evaluate the CFM log-likelihood on a grid around the fiducial simulation using the following setup:

- Fiducial point, $\boldsymbol{\theta}_\star = (M_{\rm MW}, M_{\rm LMC}, c, q) = (0.70 \times 10^{12}\,{\rm M}_\odot,\ 1.50 \times 10^{11}\,{\rm M}_\odot,\ 9.415,\ 1.0)$.
- Grid bounds, $M_{\rm MW} \in [0.68, 0.73] \times 10^{12}\,{\rm M}_\odot$, $M_{\rm LMC} \in [1.44, 1.68] \times 10^{11}\,{\rm M}_\odot$, $c \in [8.0, 10.6]$, and $q \in [0.97, 1.03]$.
- Mock data, $N = 5{,}000$ particles drawn from the held-out fiducial simulation and restricted to 30–120 kpc.

Because the benchmark is local and four-dimensional, we evaluate the likelihood on an 11-point grid per parameter ($11^4 = 14{,}641$ points), which yields the normalised posterior without sampler-convergence choices and parallelises trivially across GPUs. At each grid point we sum the per-particle log-likelihoods, $\log \mathcal{L}(\boldsymbol{\theta}) = \sum_{i=1}^{N} \log p(\mathbf{x}_i \mid \boldsymbol{\theta})$. With a uniform prior over the grid, the discrete posterior is $p(\boldsymbol{\theta} \mid \mathbf{x}) \propto \exp[\log \mathcal{L}(\boldsymbol{\theta})]$, normalised by the Riemann sum over the grid.

We compare this particle-level posterior with a Fisher forecast from the velocity moments used in Sheng et al. (2025). The moment vector $\mathbf{v}$ has 15 components: the mean velocities $\langle v_r \rangle$ and $\langle v_b \rangle$ together with the dispersions σ_{v_r}, σ_{v_b}, and σ_{v_ℓ}, measured in each of three Galactocentric distance shells (30–60, 60–90, and 90–120 kpc). Here v_r, v_b, and v_ℓ denote the radial and the two tangential (Galactic latitude and longitude) velocity components. This baseline compresses each halo to all-sky first and second velocity moments in radial bins, without using the angular structure of the density field, the wake, or the full 6D particle distribution.

Assuming a Gaussian likelihood for the summary vector $\mathbf{v}$, the Fisher information matrix $\mathbf{F} \in \mathbb{R}^{P \times P}$ for the $P = 4$ model parameters is defined at the fiducial parameter point $\boldsymbol{\theta}_\star$ as:

$$F_{ij} = \frac{\partial \boldsymbol{\mu}_{\mathbf{v}}^{\top}}{\partial \theta_i} \mathbf{C}_{\mathbf{v}}^{-1} \frac{\partial \boldsymbol{\mu}_{\mathbf{v}}}{\partial \theta_j}, \tag{1}$$

where $\boldsymbol{\mu}_{\mathbf{v}}(\boldsymbol{\theta}) \equiv \langle \mathbf{v} \rangle_{\boldsymbol{\theta}}$ is the mean summary vector at parameter point $\boldsymbol{\theta}$, and $\mathbf{C}_{\mathbf{v}}$ is the summary covariance matrix. The covariance matrix $\mathbf{C}_{\mathbf{v}}$ is estimated from $N_{\rm real} = 2{,}000$ independent realisations of the fiducial simulation:

$$\mathbf{C}_{\mathbf{v}} = \frac{1}{N_{\rm real} - 1} \sum_{k=1}^{N_{\rm real}} (\mathbf{v}^{(k)} - \bar{\mathbf{v}})(\mathbf{v}^{(k)} - \bar{\mathbf{v}})^{\top}, \tag{2}$$

where $\mathbf{v}^{(k)}$ is the summary vector of the k-th realisation and $\bar{\mathbf{v}}$ is the sample mean. The parameter derivatives are computed by central finite differences:

$$\frac{\partial \boldsymbol{\mu}_{\mathbf{v}}}{\partial \theta_i} \approx \frac{\boldsymbol{\mu}_{\mathbf{v}}(\boldsymbol{\theta}_\star + \Delta\theta_i \hat{\mathbf{e}}_i) - \boldsymbol{\mu}_{\mathbf{v}}(\boldsymbol{\theta}_\star - \Delta\theta_i \hat{\mathbf{e}}_i)}{2\Delta\theta_i},$$

Velocity moments deliberately omit angular structure, making them a clean, conventional baseline.

At the fiducial point, the full phase-space likelihood is 2.5–9.9 times sharper than conventional velocity moments.

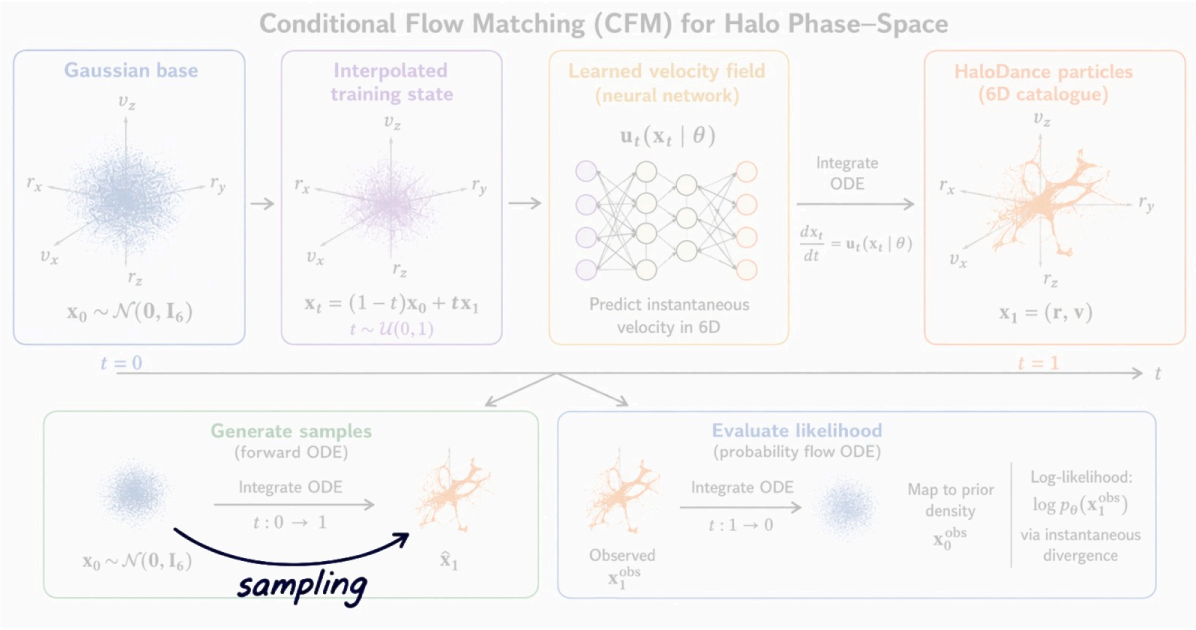


Figure 1. Schematic of the conditional flow matching setup. Gaussian base particles $\mathbf{x}_0$ are connected to HaloDance phase-space particles $\mathbf{x}_1 = (\mathbf{r}, \mathbf{v})$ through interpolated states $\mathbf{x}_t$. The network learns the conditional velocity field $\mathbf{u}_t(\mathbf{x}_t \mid \theta)$, forward integration generates mock halo tracers, while backward integration maps observed 6D catalogues to the base distribution for likelihood evaluation.

Table 1. Fiducial marginal 1σ constraints from the full phase-space CFM posterior and the velocity-moment Fisher forecast.

Parameter	CFM	Velocity moments	$\sigma_{\mathrm{vel}}/\sigma_{\mathrm{CFM}}$
M_{MW}	0.066	0.655	9.9
M_{LMC}	0.36	2.02	5.6
c	0.52	2.26	4.3
q	0.0084	0.021	2.5

information left on the table

Notes. Mass constraints are in units of $10^{11}\,\mathrm{M}_\odot$ for M_{MW} and $10^{10}\,\mathrm{M}_\odot$ for M_{LMC}; c and q are dimensionless. The final column reports the ratio of the uncertainties $\sigma_{\mathrm{vel}}/\sigma_{\mathrm{CFM}}$. The CFM constraints are half the 16th to 84th percentile range of the one-dimensional marginal posteriors.

where $\hat{\mathbf{e}}_i$ is the unit vector along the i-th parameter direction, $\Delta\theta_i$ is a small step size, and the mean summary vectors $\mu_{\mathrm{v}}(\theta_\star \pm \Delta\theta_i\hat{\mathbf{e}}_i)$ are evaluated by averaging over $N_{\mathrm{deriv}} = 100$ independent paired realisations per parameter direction. Under this Gaussian summary-likelihood approximation, the inverse Fisher matrix gives the Cramér–Rao bound on the parameter covariance, and the forecast marginalised 1σ constraints follow as $\sigma_\alpha = \sqrt{[F^{-1}]_{\alpha\alpha}}$.

Figure 3 shows the one- and two-dimensional marginal posteriors from the CFM grid search. The CFM contours are the credible regions of the grid posterior after smoothing the two-dimensional marginals with a Gaussian kernel of one-bin width, while the velocity-moment contours are analytic Gaussian ellipses at the matching radii. The fiducial truth $\theta_\star$ lies within the 68 per cent contour in all panels, showing that the particle-level CFM likelihood recovers the input MW–LMC parameters in this test. Table 1 gives the comparison to the velocity-moment Fisher forecast. The last column reports the ratio of the velocity-moment marginal uncertainty to the CFM marginal uncertainty, so it gives the constraint-tightening factor directly. The CFM posterior is tighter than the velocity-moment forecast by factors of 9.9, 5.6, 4.3, and 2.5 for $(M_{\mathrm{MW}}, M_{\mathrm{LMC}}, c, q)$, respectively. Thus, at this fiducial point, the raw catalogue carries information that is not captured by radial-bin means and dispersions alone.

MW mass
LMC mass
concentration
flattening
CFM
moments

Two limitations motivate the summary-level analysis that follows. First, evaluating the exact field-level likelihood requires a dense CFM grid and is computationally demanding, so this benchmark is restricted to the single fiducial point; validation across the broader parameter space instead uses the summary-level pipeline developed in Section 3. The fiducial comparison is nevertheless not atypical: $\theta_\star$ lies near the centre of the HaloDance ranges in M_{LMC}, c, and q, and across all 185 validation grids the summary-level posteriors preserve the same qualitative ordering, with both the BFE+MOPED and joint-summary posteriors tighter than the velocity-moment posteriors for every parameter, by factors of up to about two. The gain over velocity moments is therefore not tied to this particular fiducial choice. Second, the benchmark does not identify which structures carry that information. The neural transport field provides a likelihood for the catalogue, but its learned representation is not directly interpretable in terms of velocity channels. Because anisotropic tides, shape selection effects, and wake modelling can affect these channels differently, the phase-space distribution must be compressed into summaries that can be mapped back to physical halo perturbations.

This large gap is the reason we search for richer summaries rather than simply refining radial bins.

3 BFE DECOMPOSITION + MOPED COMPRESSION

The BFE+MOPED compression is applied to the phase-space distribution to retain information in summaries whose physical content can be inspected. This interpretability matters because simulation systematics can affect the reflex dipole, wake-like multipoles, first-moment velocities, and second-moment velocity channels differently. In the BFE representation, each coefficient, and any linear combination of coefficients, corresponds to a spatial and kinematic mode of the halo. The compression can therefore be traced back to the harmonic and velocity-moment channels that drive the constraints on each parameter.

The benchmark tells us how much information exists, but not which physical structures carry it.

The learned flow reproduces held-out one- and two-dimensional phase-space structure before we trust its likelihood.

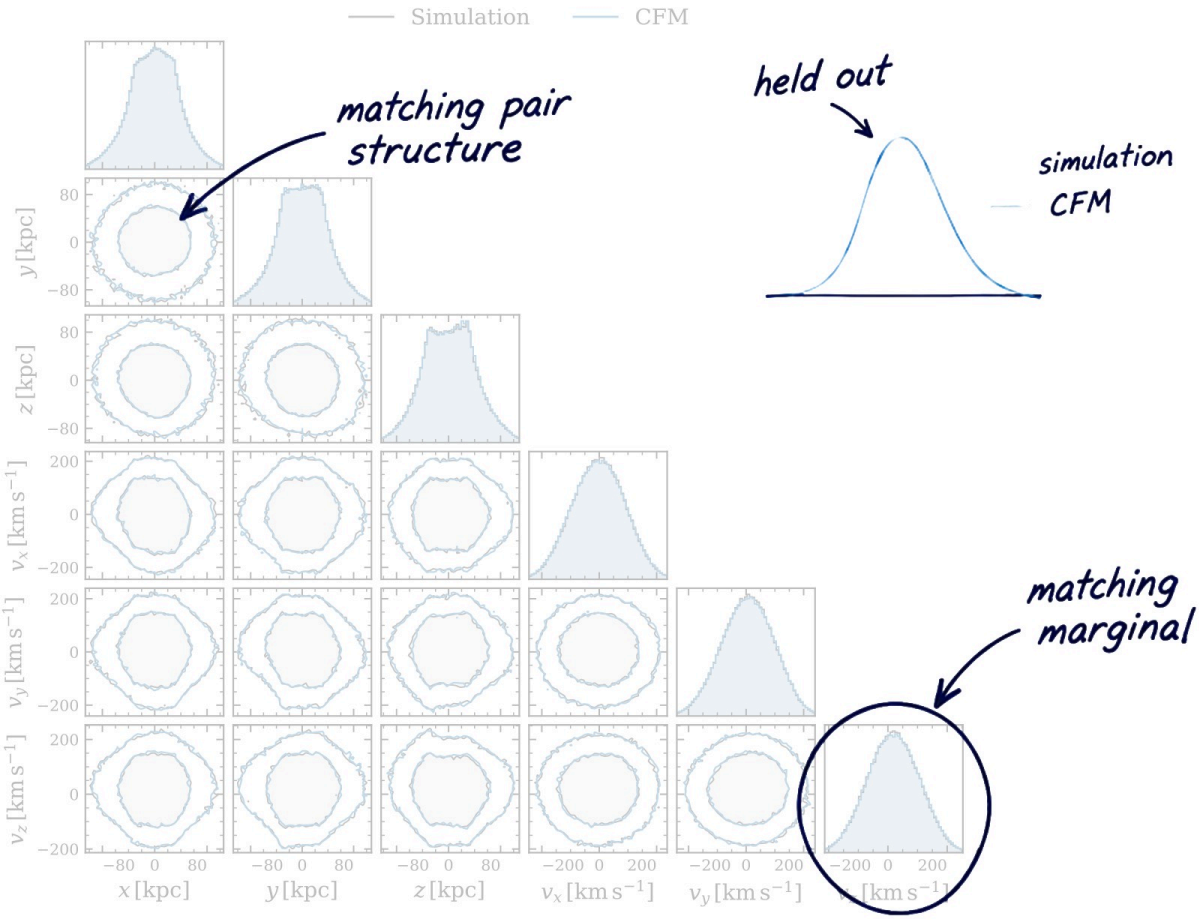


Figure 2. Pairwise comparison of the six Galactocentric phase-space coordinates for CFM-generated particles (blue) and particles from a held-out fiducial simulation (grey), both restricted to 30–120 kpc. The fiducial point is $(M_{\rm MW}, M_{\rm LMC}, c, q) = (0.7\times10^{12}\,{\rm M}_\odot, 1.5\times10^{11}\,{\rm M}_\odot, 9.415, 1.0)$ with $\beta(r) = 0$. The diagonal panels show the one-dimensional marginals of x, y, z (kpc) and v_x, v_y, v_z (km s^{-1}), and the lower triangle shows the pairwise distributions. The two samples overlap across all panels, showing that the CFM matches both the one-dimensional marginals and the pairwise structure at a parameter point outside its training set.

Agreement here is necessary, not sufficient; localized wake structure is a harder test than smooth marginals.

3.2 Decomposition of the phase space

A BFE represents a self-gravitating system with paired density and potential basis functions. The term "biorthogonal" means that the density basis and the potential basis are orthogonal to each other as a pair. Projecting the particle distribution onto the potential basis gives the coefficient of the matching density basis function. This pairing matters because each density basis function and its potential partner satisfy Poisson's equation together.

We implement the expansion using the publicly available EXP code (Petersen & Weinberg 2025). The density and potential fields of a discrete system of N equal-mass particles are both expanded over the same amplitude coefficients a_μ,

$$\rho(\mathbf{x}) = \sum_\mu a_\mu\,\rho_\mu(\mathbf{x}),$$
$$\Phi(\mathbf{x}) = \sum_\mu a_\mu\,\phi_\mu(\mathbf{x}),$$
$$a_\mu = \frac{1}{N}\sum_{i=1}^{N}\phi_\mu(\mathbf{x}_i), \qquad (4)$$

The BFE will provide that spatially resolved check while retaining a physical interpretation.

where $\{\rho_\mu, \phi_\mu\}$ are biorthogonal density–potential pairs satisfying Poisson's equation $\nabla^2\phi_\mu = 4\pi G\,\rho_\mu$ and the biorthogonality condition $\int_0^\infty \phi_\mu\,\rho_\nu\,w(r)\,{\rm d}r = \delta_{\mu\nu}$, with $w(r) \equiv r^2/4\pi G$ (Darragh-Ford et al. 2025). In spherical coordinates, Poisson's equation is separable, so the basis functions factorize as

$$\rho_\mu(\mathbf{x}) = \rho_{n\ell}(r)\,Y_\ell^m(\theta,\phi), \qquad \phi_\mu(\mathbf{x}) = \phi_{n\ell}(r)\,Y_\ell^m(\theta,\phi), \qquad (5)$$

with multi-index $\mu = (n, \ell, m)$. The spherical harmonics Y_ℓ^m carry the angular structure of each mode. The radial functions $\{\rho_{n\ell}(r), \phi_{n\ell}(r)\}$ are the eigenfunctions of the radial Sturm–Liouville problem set by Poisson's equation and the biorthogonality condition. In practice, EXP computes these radial functions for a user-specified input density profile.

The angular order ℓ gives each density coefficient a physical interpretation. The $\ell = 0$ terms encode the spherically symmetric mass distribution, or NFW monopole. The $\ell = 1$ terms encode the density dipole produced when the MW inner halo is displaced relative to the outer halo by the LMC's gravitational pull. Higher orders, $\ell \geq 2$, encode halo morphology beyond the dipole, including the LMC dynamical-friction wake. The $n = 0$ eigenfunction reproduces the input profile, and higher-n eigenfunctions add radial nodes. In this study, we adopt an NFW profile with scale radius $r_s = 16$ kpc as the zeroth-order basis and expand to $\ell_{\rm max} = 5$ and $n_{\rm max} = 40$ radial orders over the radial range 0.5–300 kpc.

To verify that these expansion orders can describe perturbed MW halos in simulation, we apply the BFE to a sample of 5×10^5 particles drawn from the fiducial simulation within 30–120 kpc in Galacto-

The BFE converts the halo into density, mean-velocity, and velocity-dispersion fields while preserving angular structure.

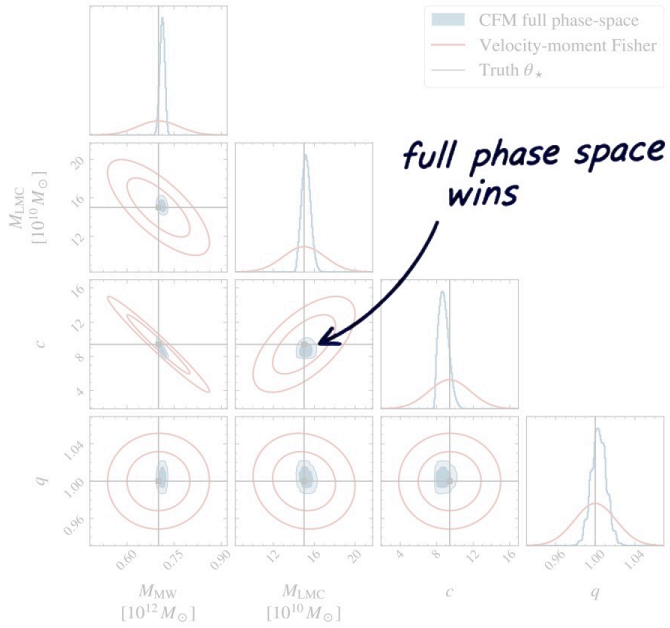


Figure 3. CFM-based full-phase-space likelihood benchmark for the held-out fiducial simulation at $\theta_\star$ with $\beta(r) = 0$. The corner plot overlays two contour sets in the same parameter planes: the CFM full-phase-space posterior from the likelihood grid for $N = 5{,}000$ particles in 30–120 kpc (blue, filled), and the velocity-moment Fisher forecast (red, Gaussian contours). Both sets show the 68 and 95 per cent credible regions, and black lines mark the true parameters $\theta_\star$.

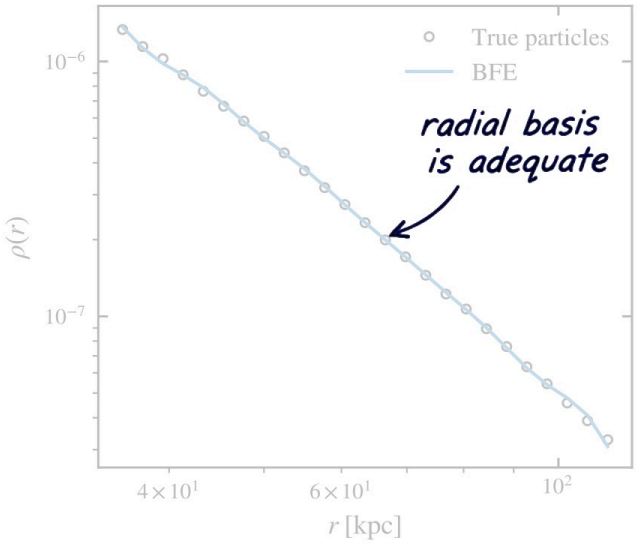
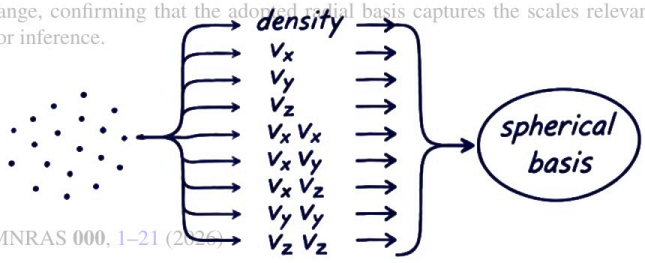


Figure 4. Spherically averaged density profile of particles drawn from the fiducial simulation within 30–120 kpc. Open circles show the particle density estimate, and the solid blue line shows the BFE reconstruction with $n_{\rm max} = 40$ and $\ell_{\rm max} = 5$. The BFE traces the particle data across the full radial range, confirming that the adopted radial basis captures the scales relevant for inference.



centric coordinates, the same particle set used for the comparison in Figure 2. Figure 4 shows that $n_{\rm max} = 40$ radial orders reconstruct the spherically averaged density profile of this sample. The BFE line traces through the particle density estimate across the full radial range, indicating that the basis captures the radial structure of the halo over the scales relevant for inference.

To go beyond the density distribution and capture the velocity structure, we extend the BFE to velocity-weighted projections. Since the BFE projection is linear in particle weights, replacing the mass weight m_i with $m_i v_{k,i}$ yields coefficients that project the momentum density field $j_k(\mathbf{x}) \equiv \rho(\mathbf{x})\langle v_k\rangle(\mathbf{x})$ onto the same biorthogonal basis,

Velocity weighting reuses the same spatial basis; it is not an expansion in velocity coordinates.

$$b_\mu^k = \frac{1}{N}\sum_{i=1}^{N} v_{k,i}\,\phi_\mu(\mathbf{x}_i), \tag{6}$$

for each Cartesian component $k \in \{x, y, z\}$. Here $\phi_\mu(\mathbf{x}_i)$ is evaluated at the particle position, while $v_{k,i}$ is only the velocity weight applied to that spatial projection. The expansion is therefore a velocity-weighted spatial BFE, not a basis expansion over velocity coordinates. The mean velocity field is then recovered point-wise as $\langle v_k\rangle(\mathbf{x}) = j_k(\mathbf{x})/\rho(\mathbf{x})$. Similarly, replacing weights with $m_i v_{k,i} v_{l,i}$ yields coefficients c_μ^{kl} for the second velocity-moment field $\rho\langle v_k v_l\rangle$, from which the velocity dispersion tensor $\sigma^2_{kl}(\mathbf{x}) = \langle v_k v_l\rangle(\mathbf{x}) - \langle v_k\rangle(\mathbf{x})\langle v_l\rangle(\mathbf{x})$ is obtained. For particles restricted to 30–120 kpc, these projections give one density channel, three momentum-density channels, and six independent second-moment channels. Together they form the high-dimensional phase-space feature vector used for compression.

Figure 5 is a reconstruction test for the BFE representation, not an inference result. The left column shows the particle-estimated field, the middle column shows the BFE reconstruction, and the right column shows the particle field minus the BFE reconstruction. The density row contains the large-scale north–south asymmetry and the local wake. The radial-velocity row contains the strongest first-moment kinematic response, a coherent reflex dipole with amplitude $|v_r| \sim 25$ km s^{-1}. The v_b and v_ℓ rows have lower-amplitude angular structure. The residual maps are small compared with the large-scale patterns that are later compressed by MOPED. Across the three velocity components the absolute residual is comparable, about 2.9 km s^{-1}, and does not grow for the lower-amplitude v_b and v_ℓ fields (with RMS amplitudes of 16.7 and 7.3 km s^{-1}, against 26.2 km s^{-1} for v_r). This component-independent residual reflects finite-sampling (shot) noise rather than an inadequate angular basis: the residual falls with particle number at close to the expected $1/\sqrt{N}$ rate (fitted slope magnitudes of 0.31–0.40, with a small residual floor), whereas increasing $\ell_{\rm max}$ beyond 5–6 does not improve the reconstruction.

Ten physical channels are richer than all-sky moments but still traceable back to recognizable halo responses.

3.2 Linear compression with MOPED

With the full $\ell_{\rm max} = 5$, $n_{\rm max} = 40$ BFE basis, the 10 channels contain 14,400 active coefficients per halo (36 real harmonics per radial order and channel, after discarding the redundant $m = 0$ sine terms of the raw storage); the adopted truncation at $n_{\rm max} = 30$, chosen below by the Fisher-volume sweep, retains 10,800. Even after this masking, the feature vector $\mathbf{f}$ is too high-dimensional for robust likelihood estimation from the available simulations, so we apply the MOPED algorithm, which reduces $\mathbf{f}$ to exactly $P = 4$ linear combinations, one per parameter, while preserving the Fisher information of the original data. When the noise covariance matrix $\mathbf{C}$ of $\mathbf{f}$ is independent of the parameters, the Fisher matrix computed from the P compressed summaries is identical to that from the full data vector.

With five angular orders and forty radial orders, the BFE recovers the large-scale dipole and wake while leaving mostly shot noise.

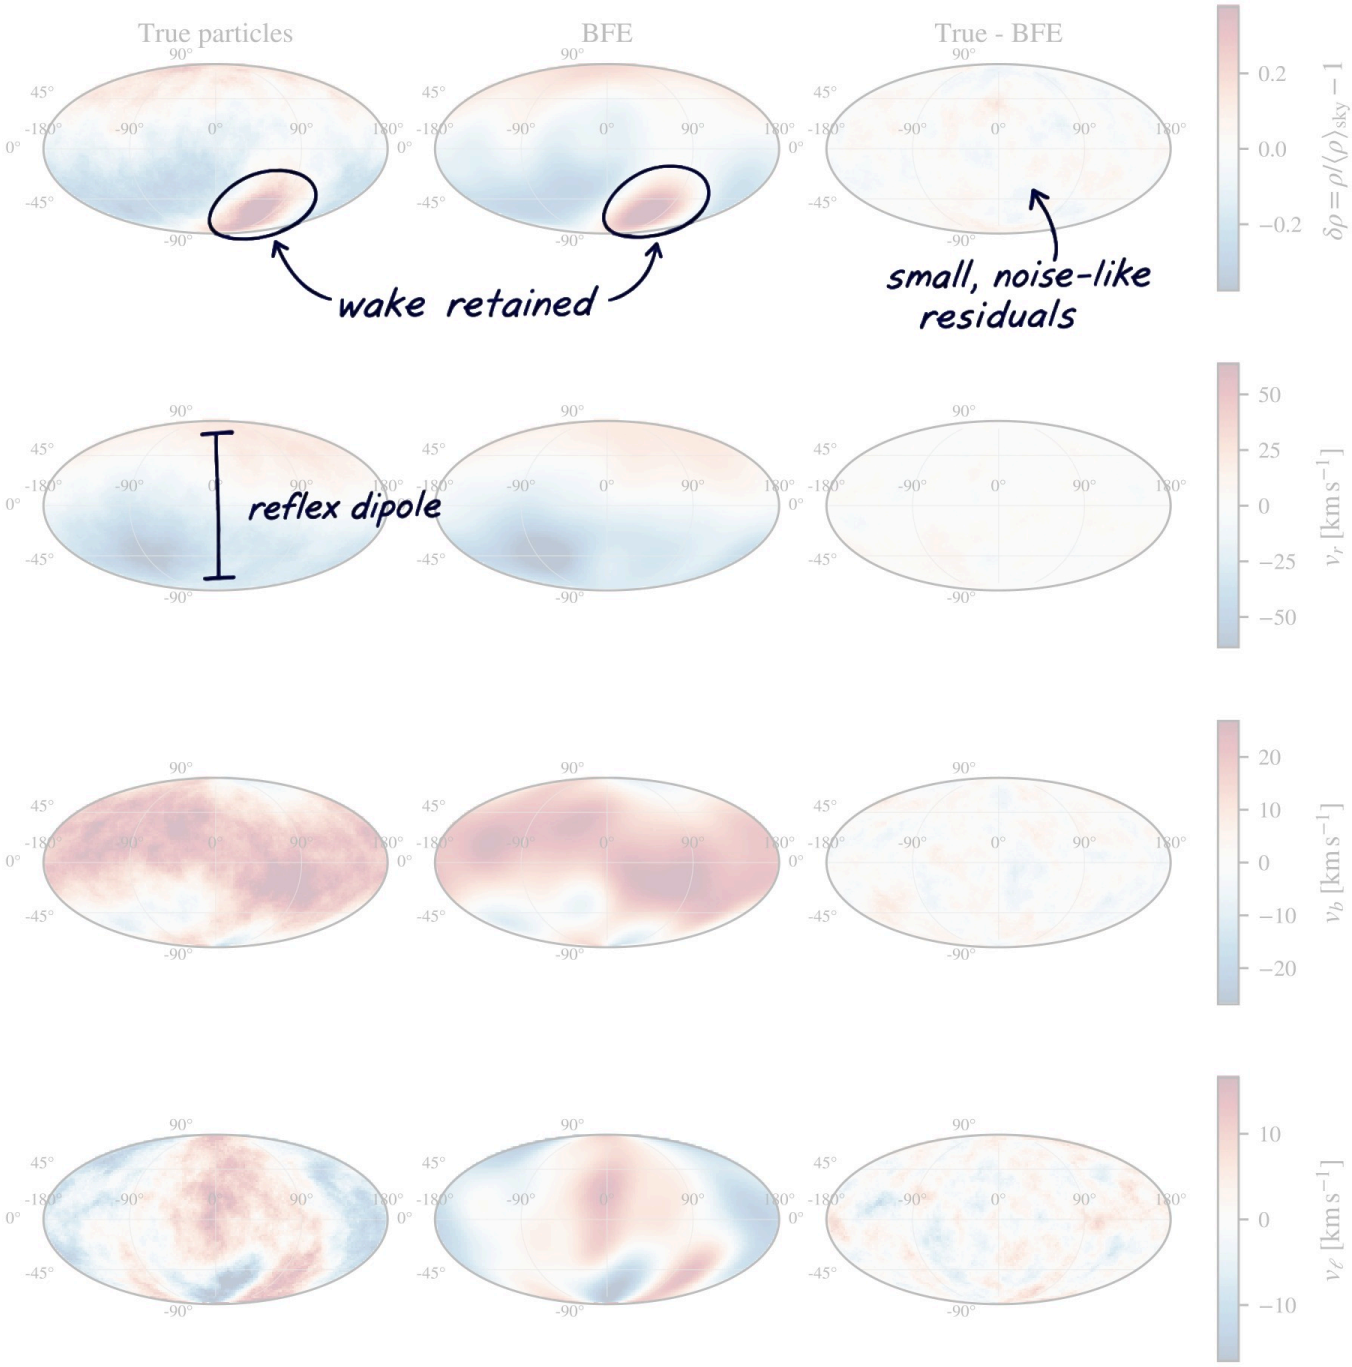


Figure 5. BFE reconstruction of the fiducial density and velocity fields in the 60–90 kpc shell. Rows show the fractional density contrast $\delta\rho = \rho/\langle\rho\rangle_{\rm sky} - 1$ and the Galactocentric velocities v_r, v_b, and v_ℓ. Columns show the particle field, the BFE reconstruction with $n_{\rm max} = 40$ and $\ell_{\rm max} = 5$, and the residual. This figure checks that the BFE basis represents the large-scale density asymmetry, radial-velocity reflex dipole, and smaller angular velocity structure before MOPED compression is applied.

The residual texture is visually busy, but its amplitude is small compared with the coherent reconstructed fields.

Increasing angular order further did not help; adding basis complexity cannot remove finite-particle noise.

MOPED compresses roughly ten thousand BFE coefficients into four locally parameter-sensitive numbers.

In the original MOPED formulation the optimal weight vector for parameter θ_α is $\mathbf{b}_\alpha \propto \mathbf{C}^{-1}\boldsymbol{\mu}_{,\alpha}$, where $\boldsymbol{\mu}_{,\alpha} \equiv \partial\langle\mathbf{f}\rangle/\partial\theta_\alpha$ is the gradient of the mean feature vector with respect to θ_α, and successive vectors $\{\mathbf{b}_\alpha\}$ are made mutually orthogonal by Gram–Schmidt construction (Heavens et al. (2000), eq. 14).

We implement the MOPED projection using a standard three-step construction at the fiducial point $\boldsymbol{\theta}_\star$. First, BFE features are standardised using 2,000 independent realisations at $\boldsymbol{\theta}_\star$, adopting a diagonal covariance approximation for the high-dimensional feature vector. Second, the parameter derivatives (scores) of the feature means are estimated by central finite differences using realisations drawn from the trained flow. A MOPED projection matrix is constructed via QR-decomposition of the score matrix, which acts as a Gram–Schmidt orthogonalisation of the parameter-sensitive directions in the standardised feature space. The standardisation procedure, and tests of its robustness, are detailed in Appendix B.

Fisher preservation is an exact-limit result; finite-simulation covariance and derivative estimates can degrade compression.

We denote the resulting four-dimensional BFE+MOPED summary vector by $\mathbf{m}$. Its components are written as $s_{M_{\rm MW}}$, $s_{M_{\rm LMC}}$, s_c, and s_q, where each component is the MOPED projection optimised for the corresponding parameter direction at $\boldsymbol{\theta}_\star$. This notation separates the BFE+MOPED summaries $\mathbf{m}$ from the generic summary variable $\mathbf{s}$ used later in the mutual-information definition.

As a diagnostic of the fiducial MOPED construction, Figure 6 shows the response matrix $R_{\alpha\beta} = \partial\langle\tilde{s}_\alpha\rangle/\partial\theta_\beta$ of the whitened summaries, in units of the response per 1σ parameter shift. The matrix is diagonally dominant: each summary mainly tracks its intended parameter, and the $M_{\rm LMC}$, c, and q summaries are cleanly isolated, with off-diagonal terms below about half of their diagonals. The two largest off-diagonal terms both appear in the MW-mass summary: its response to c reflects the familiar NFW mass–concentration degeneracy, and its response to $M_{\rm LMC}$, about a quarter of its self-response and of the same sign, means a heavier LMC shifts the MW-mass summary in the same direction as a heavier MW. These couplings match the correlated mass constraints in the Fisher forecast and reflect physical degeneracies rather than shortcomings of the compression.

We use this fiducial transfer matrix for all BFE+MOPED summaries, so the channel-level physical decomposition should be interpreted as a local diagnostic. We assess the transfer-matrix stability using 10 independent, non-fiducial anchor simulations drawn broadly across the prior volume, with $M_{\rm MW} \in [0.50, 1.80] \times 10^{12}$ $\mathrm{M}_\odot$, $M_{\rm LMC} \in [7.9, 18.9]\times10^{10}$ $\mathrm{M}_\odot$, $c \in [6.2, 14.8]$, and $q \in [0.51, 1.31]$. At each anchor, we recompute the MOPED transfer matrix and compare its parameter-sensitive directions with the fiducial ones through their principal angles. The best-matched direction stays close to the fiducial one across the prior (principal cosine 0.70–0.84), while the remaining directions rotate more, reaching principal angles of 64–79°. Projected onto the $M_{\rm MW}$–$M_{\rm LMC}$ plane, $s_{M_{\rm MW}}$ is the most stable, with mean cosine to the fiducial direction of 0.74 and 0.68, followed by s_c at 0.64, while the $M_{\rm LMC}$ direction is the least stable, at 0.48. The channel decomposition is therefore most reliable for c and $M_{\rm MW}$, and weakest for $M_{\rm LMC}$.

The LMC-mass direction rotates most across parameter space, so its channel interpretation is the least globally stable.

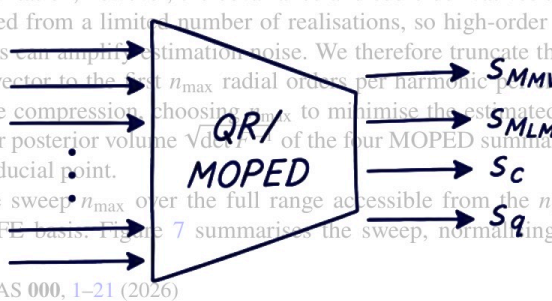
In the exact Fisher-information limit, adding features cannot reduce the information in the data vector. In our finite-simulation implementation, however, the covariance and score derivatives are estimated from a limited number of realisations, so high-order radial modes can amplify estimation noise. We therefore truncate the feature vector to the first $n_{\rm max}$ radial orders per harmonic multipole channel before compression, choosing $n_{\rm max}$ to minimise the estimated joint Fisher posterior volume $\sqrt{\det \mathbf{F}}$ of the four MOPED summaries at the fiducial point.

We sweep $n_{\rm max}$ over the full range accessible from the $n_{\rm max} = 40$ BFE basis. Figure 7 summarises the sweep, normalising each marginal constraint and the joint Fisher volume by the mean over the plateau values $n_{\rm max} = 30$, 35, and 40. The marginal constraint on $M_{\rm LMC}$ improves sharply between $n_{\rm max} = 20$ and $n_{\rm max} = 25$, from $\sigma_{M_{\rm LMC}} = 1.36 \times 10^{10}$ $\mathrm{M}_\odot$ to 1.11×10^{10} $\mathrm{M}_\odot$, indicating that radial orders $n \approx 21$–25 capture the dominant LMC-sensitive structure at 30–120 kpc.

Beyond $n_{\rm max} = 30$ the estimated constraints plateau and then loosen; the joint Fisher volume reaches its minimum at $n_{\rm max} = 30$, 6.0 per cent tighter than the full $n_{\rm max} = 40$ expansion. This loosening is finite-sample degradation of the estimated compression, not a loss of physical information. The score covariance of the four summaries stays well conditioned across the sweep (condition number 4.6–6.2), but the radial modes of order $n > 30$, though comparable in amplitude to lower orders, have finite-difference derivatives that are markedly less stable across random seeds. These modes let the MOPED projection fit derivative noise and generalise worse to independent validation data.

Modes beyond radial order thirty began fitting derivative noise rather than adding robust information.

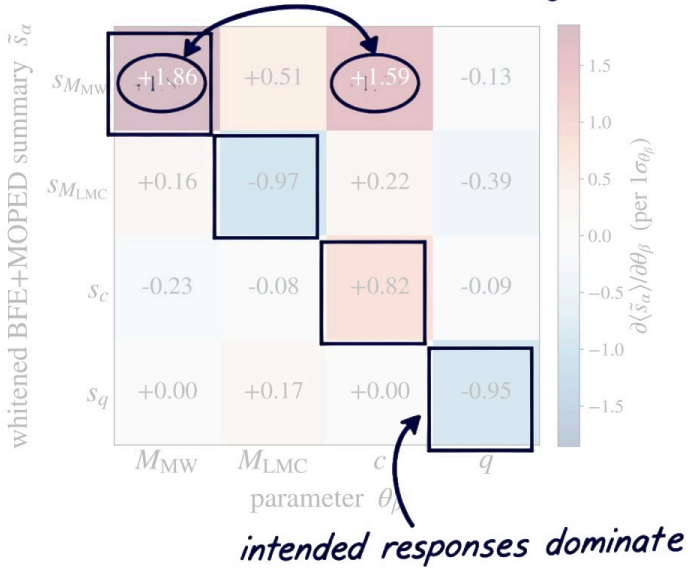
We test the same 10 independent, non-fiducial test simulations from the anchor study and find that $n_{\rm max} = 30$ is specific to the fiducial point. Across these validation runs, truncating at $n_{\rm max} = 30$ gives 0.7–3.1 per cent tighter marginal constraints than the full $n_{\rm max} = 40$ expansion (1.9% for $M_{\rm MW}$, 1.7% for $M_{\rm LMC}$, 0.7% for c, 3.1% for q). This suggests that the degradation from modes beyond the knee is not confined to the fiducial simulation.



3.3 Information retained by BFE+MOPED alone

We now compare the BFE+MOPED summaries alone against the field-level CFM and velocity-moment benchmarks. The local Fisher matrix of $\mathbf{m}$ is computed at $\boldsymbol{\theta}_\star$ exactly as for the velocity moments (equation 1), with the mean and 4×4 covariance of $\mathbf{m}$ estimated from the fiducial realisations and the mean derivatives evaluated by finite differences using the derivative simulations. For the BFE+MOPED summaries, however, the Gaussian-likelihood assumption is empiri-

mass-concentration coupling



intended responses dominate

Figure 6. Response matrix of the four whitened BFE+MOPED summaries to the model parameters, $R_{\alpha\beta} = \partial\langle\tilde{s}_\alpha\rangle/\partial\theta_\beta$, evaluated at the fiducial point and shown in units of the response per 1σ shift of each parameter. Rows are the summaries $s_{M_{\rm MW}}$, $s_{M_{\rm LMC}}$, s_c, and s_q, and columns are the model parameters $(M_{\rm MW}, M_{\rm LMC}, c, q)$. The overall sign of each summary is conventional, so only relative signs within a row are meaningful. The matrix is diagonally dominant, so each summary tracks its intended parameter, and the largest off diagonal entries reflect the mass–concentration and mass–mass degeneracies.

n_{max} = 30 is the empirical sweet spot: higher orders add unstable derivative estimates without improving constraints.

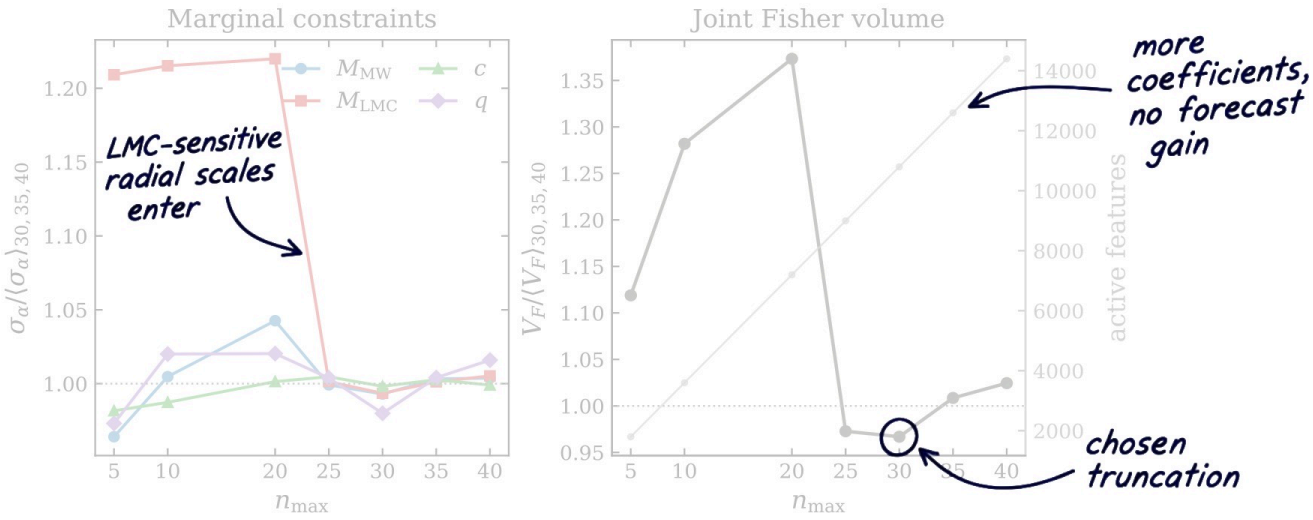


Figure 7. Radial truncation diagnostic for the BFE+MOPED compression. The left panel shows the four marginal constraints, each normalised by the mean over the plateau values n_{max} = 30, 35, and 40, so that M_{MW}, M_{LMC}, c, and q appear together. The right panel shows the joint Fisher volume normalised the same way, with the number of active BFE features on the secondary axis. The joint volume reaches its minimum at n_{max} = 30, after which adding radial orders increases the active feature count without improving the forecast.

cal rather than guaranteed: although the BFE coefficients are particle sums, the radial basis functions contain nodes and steep gradients, so finite-particle coefficients can carry high-leverage contributions and non-Gaussian tails. The Fisher matrix therefore provides only a local forecast, whose likelihood assumptions must be checked with validation diagnostics. Table 2 compares these Fisher constraints directly with the marginal posteriors from the full phase-space CFM and the velocity-moment Fisher forecast already reported in Table 1.

The BFE+MOPED constraints lie between the two benchmarks: they are (3.1, 1.8, 3.3, 1.2) times tighter than the velocity-moment forecast for (M_{MW}, M_{LMC}, c, q), while remaining (3.2, 3.1, 1.3, 2.0) times looser than the full-phase-space CFM benchmark. Thus the compression recovers part of the field-level gain in summaries whose physical content we examine in Section 3.4.

The gain over velocity moments follows from the angular structure of the LMC perturbation: the reflex response and dynamical-friction wake are anisotropic and spatially localised, so the BFE spherical harmonics capture them where all-sky velocity moments in broad radial shells partly smooth them out. To understand why the BFE+MOPED constraints nonetheless fall short of the CFM benchmark, we test two effects internal to the summary pipeline. First, radial truncation does not remove useful modes: the n_{max} sweep of Section 3.2 shows no improvement beyond n_{max} = 30. Second, to test whether parameter-dependent covariance contributes appreciable information, we decompose the Gaussian Fisher matrix into a mean-response term and a covariance-derivative term,

$$F_{\alpha\beta} = \left(\frac{\partial \boldsymbol{\mu}}{\partial \theta_\alpha}\right)^{\top} \mathbf{C}^{-1} \frac{\partial \boldsymbol{\mu}}{\partial \theta_\beta} + \frac{1}{2}\,\mathrm{Tr}\left[\mathbf{C}^{-1}\frac{\partial \mathbf{C}}{\partial \theta_\alpha}\mathbf{C}^{-1}\frac{\partial \mathbf{C}}{\partial \theta_\beta}\right], \tag{7}$$

where $\boldsymbol{\mu}$ and $\mathbf{C}$ are the summary mean and covariance. The covariance-derivative term is [illegible] of the mean-response term, [illegible] cent. The residual gap to CFM is therefore not driven by radial-order truncation or by parameter-dependent covariance within the BFE+MOPED summaries, but by information outside the four linear projections, including the higher-order and non-Gaussian phase-space correlations available to the particle-level neural transport likelihood.

In the exact Fisher limit, extra features cannot hurt; post-30 loosening is a finite-simulation effect.

The remaining CFM gap points beyond truncation to higher-order and non-Gaussian phase-space correlations.

3.4 Physical interpretability of BFE+MOPED channels

The Fisher tests quantify how much information the BFE+MOPED summaries retain. We now identify which phase-space structures supply that information. Although the MOPED projection preserves the Fisher information of the BFE features, the resulting weight vectors are linear combinations of many spatial and kinematic modes. We therefore back-project the MOPED weights to all-sky maps of the per-particle contributions to each summary. This per-particle decomposition is possible because both the BFE projection and the MOPED transformation are linear.

In practice, we draw the particles from the fiducial simulation in the 60–90 kpc Galactocentric shell, calculate the single-particle BFE features, and apply the fiducial MOPED transfer matrix channel by channel to obtain the per-particle contribution $q^{(i)}_{\alpha,c}$ to the α-th compressed summary from each BFE channel c.

For the first-order velocity channels, the three Cartesian momentum contributions (b_x, b_y, b_z) are rotated into the local Galactocentric spherical basis to give the radial, latitudinal, and longitudinal components (v_r, v_b, v_ℓ). The resulting per-particle scores are smoothed onto a Mollweide grid in Galactocentric longitude and latitude for visualisation.

Figure 8 shows how the density and first-moment velocity channels enter the four scalar summaries. Each column is a parameter summary and each row a physical channel, so a bright region marks where particles in that channel and sky direction contribute strongly to the summary, not where the underlying field is largest. Red and blue give the sign of the MOPED weight, which is arbitrary under the QR convention, so the interpretable content is the angular morphology and the physical channel.

The first-order maps show that $s_{M_{LMC}}$ is tied to the LMC-induced asymmetry. Its density and velocity panels select the same large-scale reflex dipole and local wake seen in the BFE reconstruction in Figure 5. The strongest radial-velocity contribution aligns with the wake region, while the tangential-velocity panels capture coherent large-scale reflex motion signatures, including a global bias in latitudinal velocity (v_b). Thus the LMC summary is not just a high-dimensional

Back-projecting the compressed summaries reveals which sky regions and physical channels drive each parameter.

Table 2. Fiducial marginal 1σ parameter constraints from the field-level CFM benchmark, standard velocity moments, and BFE+MOPED summaries, with the joint BFE+MOPED plus velocity moment combination.

Parameter	CFM (phase-space limit)	Velocity moments	BFE+MOPED alone	Joint (19D)
$M_{\rm MW}$	0.066	0.655	0.212	0.188
$M_{\rm LMC}$	0.36	2.02	1.10	1.03
c	0.52	2.26	0.676	0.651
q	0.0084	0.021	0.0173	0.0147

Notes. Mass constraints are in units of $10^{11}\,\mathrm{M}_\odot$ for $M_{\rm MW}$ and $10^{10}\,\mathrm{M}_\odot$ for $M_{\rm LMC}$; c and q are dimensionless. The CFM values are marginal posterior widths from the particle-level likelihood grid. The velocity-moment and BFE+MOPED values are local Fisher forecasts at the same fiducial point. The joint (19D) values are the local Fisher forecast for the concatenated BFE+MOPED plus velocity moment vector measured on common realisations.

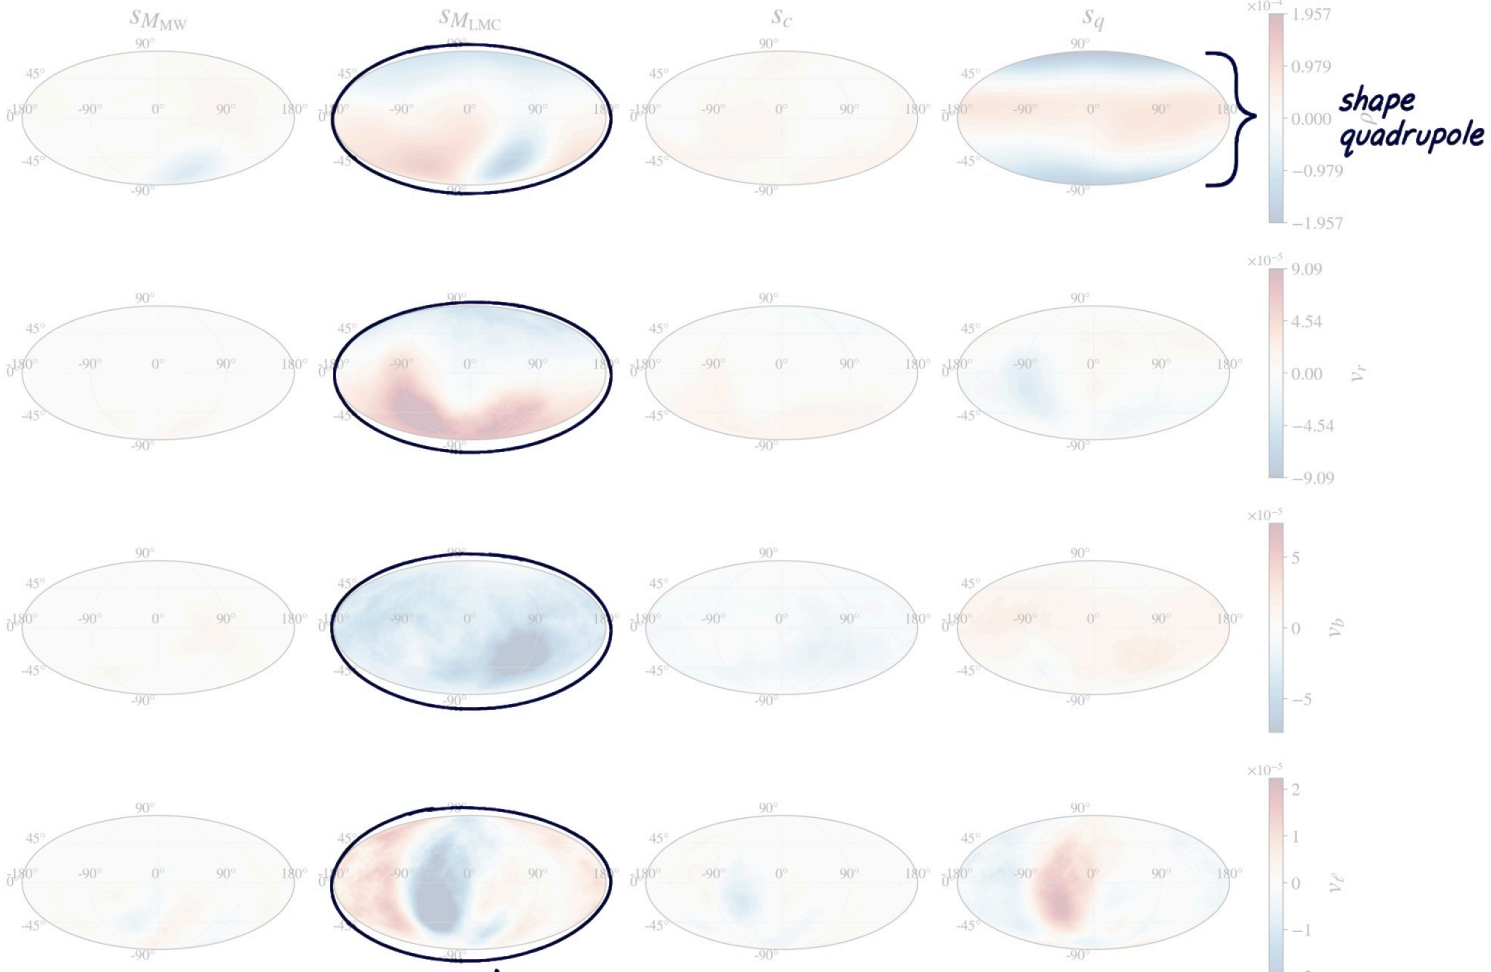


Figure 8. Per-particle MOPED contribution maps from density and first-moment velocity BFE channels in the 60–90 kpc shell. Columns correspond to the four scalar BFE+MOPED summaries, and rows correspond to ρ, v_r, v_b, and v_ℓ. These maps are not maps of the fields themselves; they show where particles in each channel increase or decrease each s[illegible] is conventional, while the angular pattern identifies the structures selected by MOPED.

These are per-particle summary-contribution maps, not maps of density or velocity themselves.

coefficient combination. It can be traced back to the expected reflex and wake response.

The same figure shows a different morphology for s_q. The density panel is dominated by a quadrupolar pattern, with opposite contributions near the equatorial and polar regions. This is the expected angular signature of changing the halo flattening at fixed radius. The first-order velocity panels carry weaker but coherent structures. The $s_{M_{\rm MW}}$ and s_c columns show little structure, for different reasons: the MW-mass response is dominated by the velocity-dispersion channels discussed below, while the largest single contribution to s_c is the angularly uniform $\ell = 0$ density monopole (24%; Table 3), which produces no visible all-sky morphology.

To complete the channel interpretation, we also back-project the diagonal second-moment velocity channels. The six independent second-moment contributions c_{kl} are contracted with the local $(\hat{\mathbf{e}}_r, \hat{\mathbf{e}}_b, \hat{\mathbf{e}}_\ell)$ basis to give the diagonal velocity-dispersion contributions (v_r^2, v_b^2, v_ℓ^2) in the same 60–90 kpc shell.

Figure 9 completes the interpretation by showing the velocity-dispersion channels. The $s_{M_{\rm MW}}$ maps are nearly monopolar in all three v^2 rows, because increasing the MW mass raises the all-sky velocity dispersion in the shell; the s_c maps show a weaker monopolar response of opposite sign, a secondary contribution that complements the density monopole and first-moment terms discussed above. Since increasing either $M_{\rm MW}$ or c raises the enclosed mass in the shell and therefore the dispersion, MOPED assigns opposite weights to this shared response to separate an overall mass increase from a change in radial profile shape. This is the velocity-space form of the NFW M_{200}–c_{200} degeneracy in tracer kinematics (e.g., Cautun et al. 2020). The weaker structures in the $s_{M_{\rm LMC}}$ and s_q columns are consistent with the dipolar and quadrupolar signatures seen in the first-order channels.

We quantify the same back-projection by grouping the per-particle MOPED weights into physical channels. For each MOPED summary s_α and channel group G, we sum the absolute values of the single-particle contributions $q^{(i)}_{\alpha,c}$ over all particles i in the 60–90 kpc shell and all BFE channels c assigned to that group:

$$P_{\alpha,G} = \frac{\sum_{c\in G}\sum_i |q^{(i)}_{\alpha,c}|}{\sum_{G'}\sum_{c\in G'}\sum_i |q^{(i)}_{\alpha,c}|}. \tag{8}$$

The density coefficients are grouped into $\ell = 0$, $\ell = 1$, and $\ell \geq 2$ multipoles. The first-moment velocity coefficients are grouped into (v_r, v_b, v_ℓ), and the diagonal second-moment coefficients are grouped into (v_r^2, v_b^2, v_ℓ^2). The absolute value is used because the MOPED weight signs are conventional and opposite-sign sky regions can otherwise cancel; the percentages therefore measure which physical channels supply the amplitude of each summary in this shell, as a local diagnostic tied to the fiducial transfer matrix.

Table 3 quantifies the same pattern: $s_{M_{\rm LMC}}$ is dominated by the first-moment velocity channels (60%) plus the $\ell = 1$ density dipole (8%), $s_{M_{\rm MW}}$ by the second-moment velocity channels (77%), and s_c splits between the density monopole (24%, its largest single channel) and the dispersion channels (37% in total), while s_q draws non-negligible signal from all three sectors. These fractions explain the physical content of the BFE+MOPED summaries, but they do not [illegible] missing relative to the standard velocity moments; [illegible]

Channel fractions explain the summaries physically, but they do not measure redundancy between summary sets.

BFE+MOPED and velocity moments overlap substantially, yet each retains information the other discards.

4 COMPLEMENTARITY OF BFE AND VELOCITY MOMENTS

BFE+MOPED summaries and velocity moments are two summaries of the same tracer catalogue, and they need not be treated as mutually exclusive alternatives. The BFE+MOPED summaries retain angular structure in the density and velocity fields, including the reflex dipole, wake, and halo-shape response, while the velocity moments provide a compact measurement of the all-sky mean and dispersion profiles in radial shells. Since these summaries compress different projections of the same phase-space perturbation, the relevant question is not only which one performs better in isolation, but whether their combination retains information that either summary loses alone. This section first uses mutual information to test whether the two summaries are complementary, and then reports the corresponding joint Fisher constraints.

4.1 Complementarity from Mutual Information

Fisher constraints rank summaries by the marginal 1σ constraints they provide on each parameter, assuming a Gaussian summary likelihood and a local comparison around the fiducial point. To go beyond these assumptions, we compare the summary representations using their total information content. Let $\mathbf{s}$ denote any summary vector derived from the same tracer catalogue, with $\mathbf{m}$ denoting the four BFE+MOPED summaries and $\mathbf{v}$ denoting the 15 standard velocity moments. The mutual information (MI) between the physical parameters $\boldsymbol{\theta}$ and a summary vector $\mathbf{s}$ is

$$I(\boldsymbol{\theta};\mathbf{s}) = \mathbb{E}_{p(\boldsymbol{\theta},\mathbf{s})}\left[\log \frac{p(\boldsymbol{\theta},\mathbf{s})}{p(\boldsymbol{\theta})\,p(\mathbf{s})}\right]. \tag{9}$$

Mutual information tests complementarity across the prior without assuming a local Gaussian posterior.

MI is parameterisation-invariant and assumes neither a Gaussian likelihood nor a local expansion around a fiducial point. It measures how much a summary contracts the posterior volume on average over the prior, so a higher MI means the summary retains more of the information in the raw 6D phase-space data. The joint mutual information $I(\boldsymbol{\theta};\mathbf{m},\mathbf{v})$ of the combined BFE+MOPED and velocity moment summaries can be decomposed using the chain rule

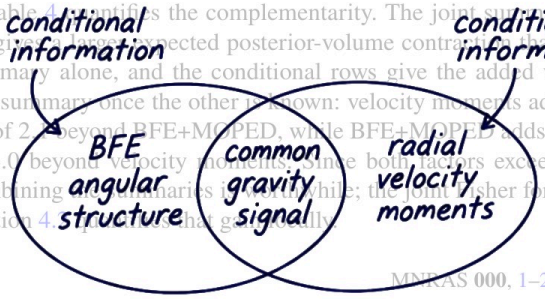
$$I(\boldsymbol{\theta};\mathbf{m},\mathbf{v}) = I(\boldsymbol{\theta};\mathbf{m}) + I(\boldsymbol{\theta};\mathbf{v}\mid\mathbf{m}), \tag{10}$$

Conditional mutual information is the decisive quantity: nonzero means concatenation can genuinely help.

where the conditional mutual information $I(\boldsymbol{\theta};\mathbf{v}\mid\mathbf{m})$ [illegible] $I(\boldsymbol{\theta};\mathbf{m})$ quantifies the additional information contributed by the velocity moments after the BFE+MOPED summaries are already known. A non-zero conditional MI indicates that the two summary sets retain complementary information from the same phase-space dataset.

We estimate the mutual information for the three summary representations using a variational Barber–Agakov lower bound, modeling the variational posterior with conditional flow matching. The detailed neural architecture, flow-matching training procedure, and backward ODE integration settings are described in Appendix C. Table 4 reports the resulting mutual information estimates and the corresponding expected posterior volume contraction factors $\exp(I)$.

Table 4 quantifies the complementarity. The joint summary vector gives the largest expected posterior-volume contraction, [illegible] either summary alone, and the conditional rows give the added value of one summary once the other is known: velocity moments add a factor of 2.[illegible] beyond BFE+MOPED, while BFE+MOPED adds a factor of 4.0 beyond velocity moments. Since both factors exceed unity, combining the summaries is [illegible] while; the joint Fisher forecast in Section 4.[illegible] quantifies that gain locally.



Second-moment channels encode the MW mass-concentration response, while joint summaries retain the most information.

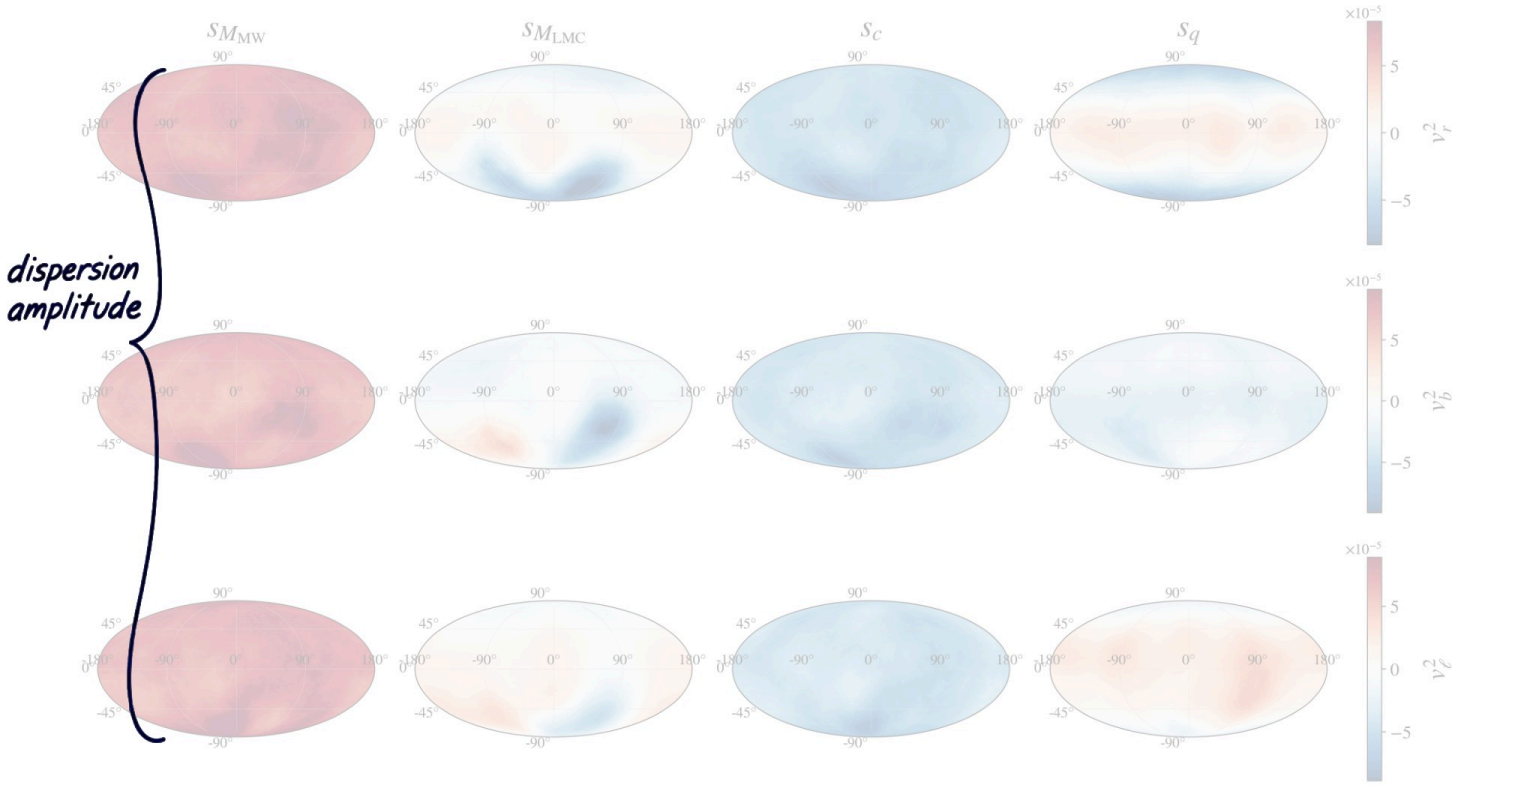


Figure 9. Per-particle MOPED contribution maps from the diagonal second-moment velocity channels in the 60–90 kpc shell. These maps show where v_r^2, v_b^2, and v_ℓ^2 contribute to each scalar summary. The dominant response is in $s_{M_{\rm MW}}$ and s_c, whose nearly monopolar patterns reflect the MW mass–concentration degeneracy in tracer dispersions.

Opposite dispersion weights help MOPED separate an overall mass increase from a change in radial profile shape.

Table 3. Channel-wise signal decomposition of the BFE+MOPED summaries in the 60–90 kpc shell. This table is an interpretability diagnostic for **m**, not the mutual-information complementarity test.

	Density multipoles ρ			First-moment velocities v			Second-moment velocities v^2		
MOPED Summary	$\ell=0$	$\ell=1$	$\ell\geq 2$	v_r	v_b	v_ℓ	v_r^2	v_b^2	v_ℓ^2
$s_{M_{\rm MW}}$	4%	5%	3%	4%	3%	3%	**25%**	**26%**	**26%**
$s_{M_{\rm LMC}}$	7%	8%	5%	**23%**	**27%**	**10%**	8%	8%	4%
s_c	**24%**	**10%**	**11%**	6%	7%	4%	**12%**	**13%**	**12%**
s_q	**14%**	**16%**	7%	**11%**	**13%**	9%	**11%**	**10%**	9%

77% from second moments

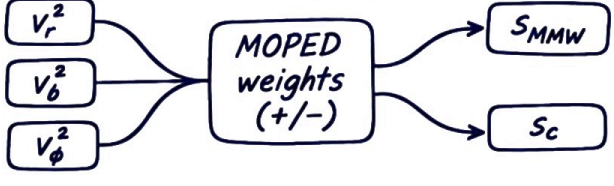
Notes. The percentages are the channel fractions $P_{\alpha,G}$ defined in equation 8. Cell backgrounds highlight dominant channels. High signal shares ($\geq 20\%$) are shaded in dark blue, and moderate shares (10%–19%) are shaded in light blue. The entries describe which BFE channels contribute to each BFE+MOPED summary.

Table 4. Mutual information and expected posterior volume contraction factors for the different summary representations. This table is the diagnostic of complementarity between BFE+MOPED summaries and velocity moments.

Summary representation	Mutual Information $I(\theta;\mathbf{s})$ [nats]	Volume contraction $\exp(I)$
Velocity moments $\mathbf{v}$ (15 dims)	7.07	1,176
BFE+MOPED $\mathbf{m}$ (4 dims)	7.72	2,253
Joint summaries $(\mathbf{m},\mathbf{v})$ (19 dims)	8.46	4,722
Information decomposition:		
Shared information $I(\theta;\mathbf{m}\cap\mathbf{v})$	6.33	561
Conditional moment info $I(\theta;\mathbf{v}\mid\mathbf{m})$	0.74	2.1
Conditional MOPED info $I(\theta;\mathbf{m}\mid\mathbf{v})$	1.39	4.0

complementary gain

Notes. The expected posterior volume contraction factor relative to the prior is given by $\exp(I)$, where I is the mutual information in nats. The shared information is defined as $I(\theta;\mathbf{m})+I(\theta;\mathbf{v})-I(\theta;\mathbf{m},\mathbf{v})$. The conditional rows quantify the information added by one summary set after the other is known.



Concatenating all nineteen summaries gives the best interpretable forecast, though the full 6D likelihood remains tighter.

4.2 Joint Fisher forecast setup

The mutual-information analysis motivates a local Fisher test of the combined summary vector. We form the joint vector by directly concatenating the four BFE+MOPED summaries and the 15 standard velocity moments,

$$\mathbf{y} = (\mathbf{m}, \mathbf{v}) \in \mathbb{R}^{19}. \tag{11}$$

The joint Fisher matrix is then computed from the covariance and parameter-response matrix of this full 19-dimensional vector at the fiducial point. Because $\mathbf{m}$ and $\mathbf{v}$ are measured from the same particle catalogue, the full covariance retains their cross-correlations (median absolute value 0.05, maximum 0.42), rather than treating the two Fisher matrices as additive.

We retain cross-covariances because both summary sets come from the same tracers; adding separate Fisher matrices would assume independence.

The resulting joint constraints, listed in the final column of Table 2, are marginal 1σ values of 0.188×10^{11} $M_\odot$, 1.03×10^{10} $M_\odot$, 0.651, and 0.0147 for $(M_{\rm MW}, M_{\rm LMC}, c, q)$. These tighten the BFE+MOPED constraints by 71.1, 49.1, 71.0, and 30.0 per cent, respectively. Relative to the field-level benchmark, the joint vector closes 16.6, 10.6, 16.1, and 29.5 per cent of the BFE+MOPED-to-CFM gap. Measured from the velocity-moment baseline instead, the joint summaries recover 50–92 per cent of the constraint-width improvement achieved by the field-level likelihood, and the remaining factor of 1.3–2.9 quantifies the information that the current linear compression does not capture.

Figure 10 places all four cases in the same parameter planes and 68 and 95 per cent convention as Figure 3. The Fisher contours nest from the velocity moments through BFE+MOPED to the joint vector, which lies closest to the CFM posterior but still encloses a visibly larger region. Consistent with Table 4, the joint contours sit only slightly inside the BFE+MOPED ones but far inside the velocity-moment ones: the 2.1-fold volume contraction from adding velocity moments is shared across all four parameters, so it appears as the modest 3.7–15.1 per cent marginal tightenings reported above, whereas adding BFE+MOPED to the moments yields the larger 4.0-fold gain.

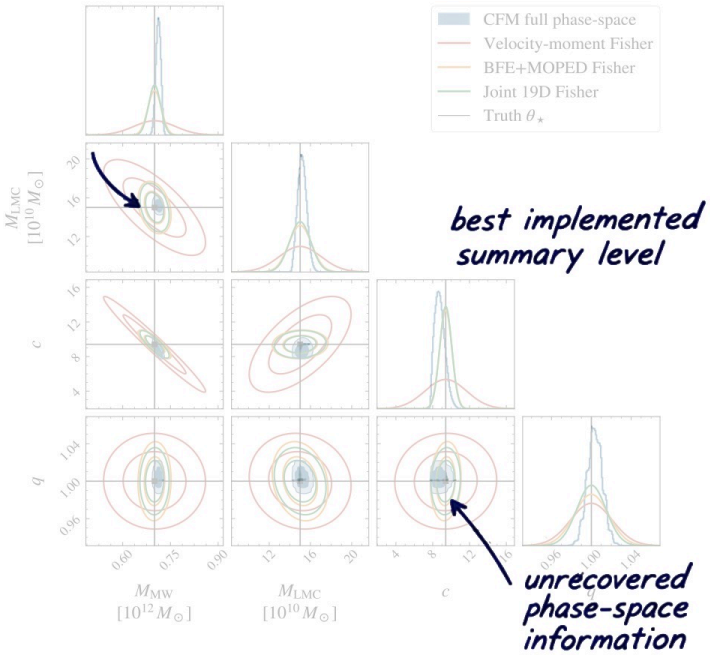


Figure 10. Fiducial marginal constraints from all summary sets against the field-level CFM benchmark, shown as 68 and 95 per cent credible regions in the same convention as Figure 3. The CFM full phase-space posterior (blue, filled) is overlaid with the velocity-moment (red), BFE+MOPED (orange), and joint 19-dimensional (green) Fisher forecasts. The forecasts nest from the loosest velocity moments through BFE+MOPED to the tightest joint vector, which still does not reach the CFM posterior.

5 SIMULATION-BASED INFERENCE ON SUMMARY STATISTICS

The summary-likelihood emulator is used for two closely related inference runs. The primary summary-level constraints use the joint vector $\mathbf{y} = (\mathbf{m}, \mathbf{v}) \in \mathbb{R}^{19}$, where $\mathbf{m}$ contains the four BFE+MOPED summaries and $\mathbf{v}$ contains the 15 velocity moments. The same emulator architecture is also applied to the four-dimensional BFE+MOPED vector $\mathbf{m}$ for the tracer-anisotropy stress test of Section 6.3. In both cases the catalogue is mapped to a summary vector using a fixed MOPED compression operator calibrated at the $\theta_\star$, and the emulator provides the likelihood of that summary as a function of the HaloDance parameters.

We first check that a Gaussian likelihood is adequate for the primary 19-dimensional joint summaries. Pooling 500 realisations at each of 185 held-out parameter points, the whitened Mahalanobis d^2 closely follows the expected χ^2_{19} distribution (mean 18.96 and variance 36.4 against the expected 19 and 38; empirical 68 and 95 per cent coverages of 68.1 and 95.4 per cent), and the marginal skewness and excess kurtosis never exceed 0.026 and 0.040 in absolute value. The four-dimensional BFE+MOPED summaries pass the same diagnostic and are used for the anisotropy test below.

Let $\mathbf{s}$ denote the summary vector used in a given run, with $\mathbf{s} = \mathbf{y}$ for the primary joint analysis and $\mathbf{s} = \mathbf{m}$ for the BFE+MOPED only stress test. We model the summary distribution at fixed parameters as

The Gaussian summary likelihood passed stringent fixed-parameter tests before we used it for simulation-based inference.

$$\mathbf{s} \mid \theta \sim \mathcal{N}\big(\mu(\theta), \Sigma(\theta)\big), \tag{12}$$

and emulate the mean and covariance with a mixture-density-network-style (MDN-style) Gaussian likelihood emulator. In the present implementation the mixture has a single multivariate Gaussian component, so the network predicts the conditional mean and covariance of $p(\mathbf{s} \mid \theta)$ rather than a flow or a particle-level density. For each training parameter point, we draw 500 independent realisations of 5,000 particles from the 30–120 kpc shell and pass each through the chosen summary compression. These realisations define an empirical mean $\hat{\mu}(\theta)$ and covariance $\hat{\Sigma}(\theta)$.

The MDN-style emulator takes the four physical parameters scaled to the unit prior hypercube as input. A shared three-layer trunk with 256 hidden units per layer, SiLU activations, and dropout 0.02 feeds two heads. One predicts the standardised summary mean, and the other predicts the $d(d+1)/2$ entries of a lower-triangular Cholesky factor $\mathbf{L}_{\rm NN}$, where d is the summary dimension ($d = 19$ for the joint vector and $d = 4$ for BFE+MOPED alone). The covariance is reconstructed as $\Sigma_{\rm NN} = \mathbf{L}_{\rm NN}\mathbf{L}_{\rm NN}^{\top}$, with the diagonal entries passed through a softplus activation and a 10^{-4} floor to enforce positive definiteness. The network is trained by minimising

$$\mathcal{L} = \mathrm{KL}\big[\, \mathcal{N}(\hat{\mu}, \hat{\Sigma}) \,\|\, \mathcal{N}(\mu_{\rm NN}, \Sigma_{\rm NN}) \big] + \lambda \left\| \mathrm{logdiag}(\Sigma_{\rm NN}) - \mathrm{logdiag}(\hat{\Sigma}) \right\|^2, \tag{13}$$

with $\lambda = 0.1$. Training uses AdamW with learning rate 10^{-3}, weight decay 10^{-4}, and batch size 256 on 1,663 training parameter points, with early stopping on the held-out validation loss.

To propagate emulator-training uncertainty, we train an ensemble of five networks with independent random seeds and combine

A learned mean and covariance make posterior sampling cheap after the catalogue is compressed once.

The summary emulator enables practical inference, while the particle-level flow sets the field-level information benchmark.

their predictions at inference time. The effective likelihood uses the ensemble-averaged mean and a covariance that includes both the predicted summary covariance and the scatter among network mean predictions,

$$\mu_{\rm eff}(\theta) = \frac{1}{M}\sum_{m=1}^{M}\mu_m(\theta),$$

$$\Sigma_{\rm eff}(\theta) = \frac{1}{M}\sum_{m=1}^{M}\Sigma_m(\theta) + \mathrm{diag}\big[\mathrm{Var}_m\,\mu_m(\theta)\big]. \quad (14)$$

The first covariance term represents the conditional summary covariance at fixed θ, and the second term adds a diagonal emulator-uncertainty floor. We retain only the diagonal of the between-network covariance, because off-diagonal entries estimated from just $M = 5$ members would inject noisy correlations; the restriction conservatively inflates the marginal variances. This floor is small, contributing a fractional posterior broadening of only a few per cent on the validation set and at most about 10 per cent in the worst case, while the deterministic five-member ensemble in production without additional per-call predictive resampling.

Five independently trained networks add a small uncertainty floor instead of pretending emulator training is exact.

Given an observed summary vector $\mathbf{s}_{\rm obs}$, the emulator likelihood is

$$\log\mathcal{L}(\theta) = -\frac{1}{2}\left[\delta^{\top}\Sigma_{\rm eff}^{-1}\delta + \log\det\Sigma_{\rm eff}\right], \quad (15)$$

where $\delta = \mathbf{s}_{\rm obs} - \mu_{\rm eff}(\theta)$. The prior is uniform over the HaloDance parameter box and zero outside it. Posterior sampling uses DYNESTY dynamic nested sampling with multi-ellipsoid bounding and the `rwalk` sampler, initialized with 500 live points and 200 live points per dynamic batch, an evidence stopping threshold $\Delta\ln\mathcal{Z}_{\rm init} = 0.05$, and a target posterior effective sample size of 4,000.

This summary-level MDN-style emulator is distinct from the particle-level CFM in Section 2.2. The particle-level CFM operates on individual 6D phase-space coordinates and defines the field-level information benchmark, whereas the summary-level emulator operates on $\mathbf{s}$ and provides the practical posterior sampler over the HaloDance parameter space. This separation is useful for planned H3 and SDSS-V applications, because the observed catalogue only has to be converted once into the adopted summary vector. Survey masks, tracer selection, and anisotropy priors can then be tested at the summary-likelihood level. The validation results for the implemented pipeline, including P-P ranks, TARP joint coverage, point-estimate residuals, and the residual q structure, are given in Appendix D.

The two neural models solve different problems: one models particles, the other models compressed catalogues.

6 DISCUSSION

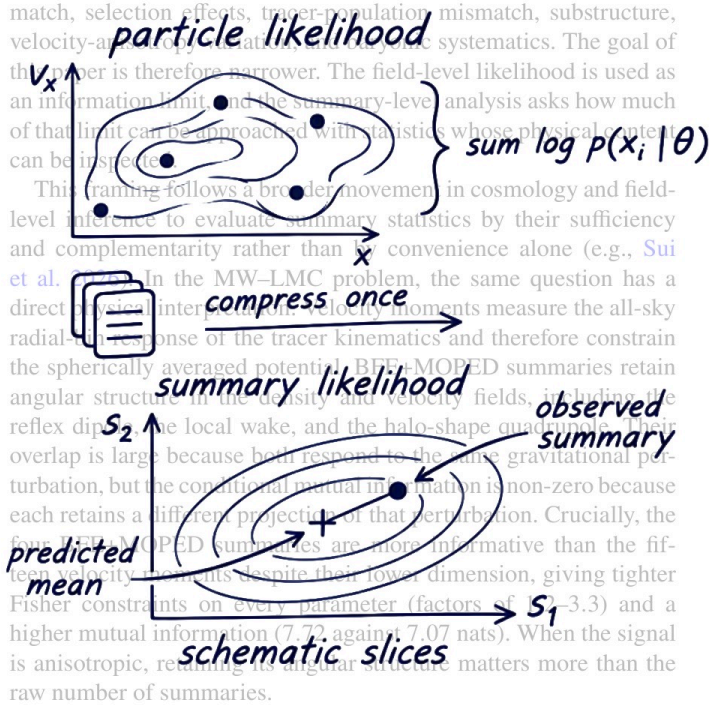
The results define a hierarchy of information content for this simulation setup. The particle-level CFM likelihood gives the tightest local constraints and therefore serves as the field-level reference. Standard velocity moments are more compressed and more interpretable, but lose angular phase-space structure. BFE+MOPED summaries sit between these cases. They recover part of the CFM gain over velocity moments while allowing the contributing density, first-moment velocity, and second-moment velocity channels to be inspected. The mutual-information test then shows that BFE+MOPED and velocity moments are not redundant, motivating the joint summary vector that we adopt for the primary summary-level constraints.

6.1 Field-level limit and interpretable summaries

The field-level CFM likelihood is useful here because modern generative density models make a particle-level likelihood for a high-dimensional tracer catalogue computationally tractable. In this work the CFM is not only a sampler. Through the continuous-normalising-flow form of the probability-flow ODE, it gives an evaluable likelihood for each 6D particle conditioned on the MW–LMC parameters. This technical step turns the raw HaloDance catalogue into a field-level information benchmark.

We use flow matching rather than a conventional normalising flow or a score-based diffusion model for this benchmark. Standard normalising flows provide exact likelihoods, but at fixed architecture they can be less flexible for the complex phase-space structure of the perturbed halo. Score-based diffusion models are expressive samplers, but their likelihoods usually require approximate probability-flow integration or variational bounds. CFM keeps the continuous-flow likelihood machinery while using a flexible generative training objective, making it a suitable compromise for a likelihood benchmark.

This benchmark should not be read as an argument for applying field-level inference uncritically to data. Quantifying emulator systematics remains difficult because the likelihood can be precise for the trained simulation distribution while still inheriting simulation mismatch, selection effects, tracer-population mismatch, substructure, velocity-anisotropy degeneracies, and baryonic systematics. The goal of this paper is therefore narrower. The field-level likelihood is used as an information limit, and the summary-level analysis asks how much of that limit can be approached with statistics whose physical content can be inspected.



This framing follows a broader movement in cosmology and field-level inference to evaluate summary statistics by their sufficiency and complementarity rather than by convenience alone (e.g., Sui et al. 2026). In the MW–LMC problem, the same question has a direct physical interpretation. Velocity moments measure the all-sky radial-velocity response of the tracer kinematics and therefore constrain the spherically averaged potential. BFE+MOPED summaries retain angular structure in the density and velocity fields, including the reflex dipole, the local wake, and the halo-shape quadrupole. Their overlap is large because both respond to the same gravitational perturbation, but the conditional mutual information is non-zero because each retains a different projection of that perturbation. Crucially, the four BFE+MOPED summaries are more informative than the fifteen velocity moments despite their lower dimension, giving tighter Fisher constraints on every parameter (factors of 1.2–3.3) and a higher mutual information (7.72 against 7.07 nats). When the signal is anisotropic, retaining its angular structure matters more than the raw number of summaries.

6.2 Complementarity and missing information

The two measures play different roles in our analysis: the mutual information of Table 4 diagnoses that BFE+MOPED summaries and velocity moments are complementary, while the joint 19-dimensional Fisher forecast and summary-likelihood emulator (Section 4.2) quantify the resulting constraint improvement.

The joint summaries nevertheless do not exhaust the field-level information: adding the velocity moments closes only 11–30 per cent of the BFE+MOPED-to-CFM gap in marginal constraint width, leaving the joint constraints a factor of 1.3–2.9 looser than the particle-level likelihood. This residual gap is not a failure of combining summaries, but identifies phase-space structure that the current linear compression does not capture, and it sets the concrete target for the improved summaries discussed below.

Comparisons with existing MW–LMC mass estimates must account for the fact that different observables probe different aspects of the same system. Stream deflections, timing arguments, LMC inter-

Interpretability helps when real surveys differ from the trained simulations in unmodeled ways.

nal kinematics, and halo-wake measurements all depend on $M_{\rm LMC}$, $M_{\rm MW}$, and the MW potential, but with different parameter combinations and modelling assumptions (e.g., Erkal et al. 2019; Cautun et al. 2020; Vasiliev et al. 2021; Shipp et al. 2021; Fushimi et al. 2024). Our comparison therefore focuses not on matching a single mass estimate, but on how much phase-space information is available, how much is retained by interpretable summaries, and which density and velocity channels carry that information.

Future improvements should preserve this interpretability while adapting the summaries to observational catalogues. Within the BFE framework, this means targeted additions such as selection-function-aware moments, per-tracer weights in place of the equal-mass particle sums used here, anisotropy and tracer-population parameters, or flexible density estimators on physically grouped coefficient blocks.

6.3 Summary-level validation and systematics

The primary summary-level SBI pipeline is trained and validated on HaloDance halos with an isotropic velocity distribution ($\beta = 0$), using the joint 19-dimensional summary vector $\mathbf{y} = (\mathbf{m}, \mathbf{v})$. Figure 11 shows that the posterior medians broadly follow the one-to-one relation on the held-out validation set, while Appendix D gives the marginal P-P ranks, joint TARP coverage, and point-estimate residuals. Because the velocity moments intentionally include dispersion information that responds directly to tracer anisotropy, the robustness experiment below uses a separately trained four-dimensional BFE+MOPED emulator to isolate the response of the BFE density and velocity channels.

The MDN-style likelihood remains an approximation because it uses a single multivariate Gaussian component. Section 5 and Appendix D show that this approximation is adequate for the primary joint summaries in the isotropic validation set, and the four-dimensional BFE+MOPED summaries pass the same fixed-parameter Gaussianity check for the stress test. For survey applications, masks, selection functions, and tracer-population mixtures may make the whitened summaries less Gaussian. In that case, the same simulation-based workflow can be retained while replacing the single-Gaussian likelihood with a more flexible density estimator for the summary-level CFM model for $p(\mathbf{s} \mid \theta)$.

Observed outer-halo tracers are radially biased (e.g., Chandra et al. 2023; Li et al. 2026), so tracer anisotropy remains a leading systematic for applying these summaries to data. To test this mismatch, we apply the isotropic-trained BFE+MOPED emulator to the separate radially varying HaloDance set of Section 2.1, with $\beta(r) = -0.15 - 0.2\,\alpha(r)$, corresponding to $\beta(r) \simeq 0.3$–0.45 across 30–120 kpc for the halos used here. This is a stress test of model misspecification, not a retrained anisotropic inference model. We select 148 non-training parameter points, draw 10 independent 5,000-particle realisations per point, and process each realisation with the same fiducial BFE+MOPED transfer matrix, MDN-style Gaussian emulator, and DYNESTY sampler used for the isotropic BFE+MOPED validation.

Figure 12 shows that the posterior widths change little in this anisotropic stress test, but the posterior medians shift. The largest shifts are a high bias in concentration, $\Delta c \simeq +1.24$, and a low bias in MW mass, $\Delta M_{\rm MW} \simeq -0.59 \times 10^{11}\,{\rm M}_\odot$. The shifts in $M_{\rm LMC}$ and q are smaller, with $\Delta M_{\rm LMC} \simeq -0.82 \times 10^{10}\,{\rm M}_\odot$ and $\Delta q \simeq +0.024$. The per-grid 68% intervals contain the truth in only 6.8%, 80.4%, 19.6%, and 40.5% of the anisotropic grids for $(M_{\rm MW}, M_{\rm LMC}, c, q)$, compared with 99.5%, 88.7%, 97.8%, and 84.9% in the isotropic case. The channel decomposition in Table 3 explains the direction of the failure. The $M_{\rm LMC}$ summary draws most of its signal from first-moment velocities and the $\ell = 1$ density dipole, which are less sensitive to velocity anisotropy. In contrast, $s_{M_{\rm MW}}$ receives much of its signal (77%) from the second-moment velocity channels, where radial anisotropy directly changes σ_{v_r} at fixed potential, while s_c retains a substantial second-moment component (37%) alongside its leading $\ell = 0$ density monopole. An isotropic-trained emulator can therefore absorb anisotropy mismatch as a change in the inferred MW mass and concentration.

This quantitative-statistic approach also gives a practical way to diagnose such systematics. Because the summaries can be decomposed by density and first-moment and second-moment velocity channels, one can test whether a posterior shift is driven by a physically suspect part of the data vector. In a survey application, the same framework can be rerun after censoring selected channels or radial shells, or after adding nuisance parameters for tracer anisotropy, selection, distance errors, and contamination. This does not remove the need for simulation validation, but it makes the failure mode more local than in a less transparent field-level likelihood.

The anisotropy stress test is a first example of this audit. It identifies the second-moment channels as the place where the isotropic training assumption enters most strongly, and it points to the nuisance parameters and channel tests needed before applying the method to H3 or SDSS-V data.

7 CONCLUSIONS

The recent infall of the LMC has driven the MW outer halo into dynamical disequilibrium that encodes the masses and structural parameters of both galaxies in the 6D phase-space distribution of halo tracers. We compare a field-level CFM benchmark, physically interpretable BFE+MOPED summaries, standard velocity moments, and their joint summary vector for the parameter vector $\theta = (M_{\rm MW}, M_{\rm LMC}, c, q)$. Our main conclusions are:

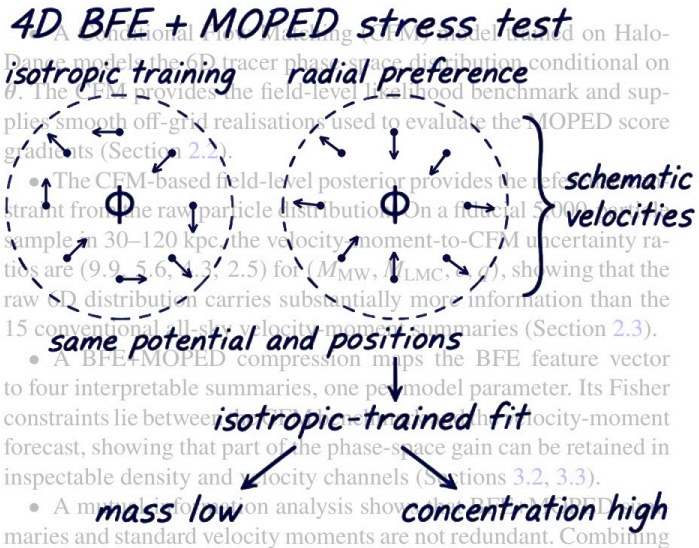
- A conditional flow matching CFM model trained on HaloDance models the 6D tracer phase-space distribution conditional on θ. The CFM provides the field-level likelihood benchmark and supplies smooth off-grid realisations used to evaluate the MOPED score gradients (Section 2.2).
- The CFM-based field-level posterior provides the reference constraint from the raw particle distribution. On a fiducial 5000-particle sample in 30–120 kpc, the velocity-moment-to-CFM uncertainty ratios are (9.9, 5.6, 4.3, 2.5) for $(M_{\rm MW}, M_{\rm LMC}, c, q)$, showing that the raw 6D distribution carries substantially more information than the 15 conventional all-sky velocity-moment summaries (Section 2.3).
- A BFE+MOPED compression maps the BFE feature vector to four interpretable summaries, one per model parameter. Its Fisher constraints lie between the BFE-only and full velocity-moment forecast, showing that part of the phase-space gain can be retained in inspectable density and velocity channels (Sections 3.2, 3.3).
- A mutual information analysis shows that BFE+MOPED summaries and standard velocity moments are not redundant. Combining them into the joint 19-dimensional vector tightens the marginal constraints by up to 15 per cent relative to BFE+MOPED alone and by 30–71 per cent relative to velocity moments alone, coming within a factor of 1.3–2.9 of the field-level benchmark; the remaining gap quantifies the margin left for future summary designs (Section 4).
- Back-projecting the MOPED weights shows what each BFE+MOPED summary measures. The $s_{M_{\rm LMC}}$ summary follows the reflex dipole and wake, s_q follows the quadrupolar halo-shape

Tracer anisotropy can masquerade mainly as shifts in inferred MW mass and concentration.

This intentionally tests misspecification; the inference model was not retrained on anisotropic halos.

The channel decomposition localized the strongest mismatch to second-moment velocity channels.

Posterior widths changed little while medians shifted, a useful warning that precision does not guarantee validity.



Held-out isotropic validation performs well, while a separate BFE+MOPED anisotropy stress test reveals parameter-dependent biases.

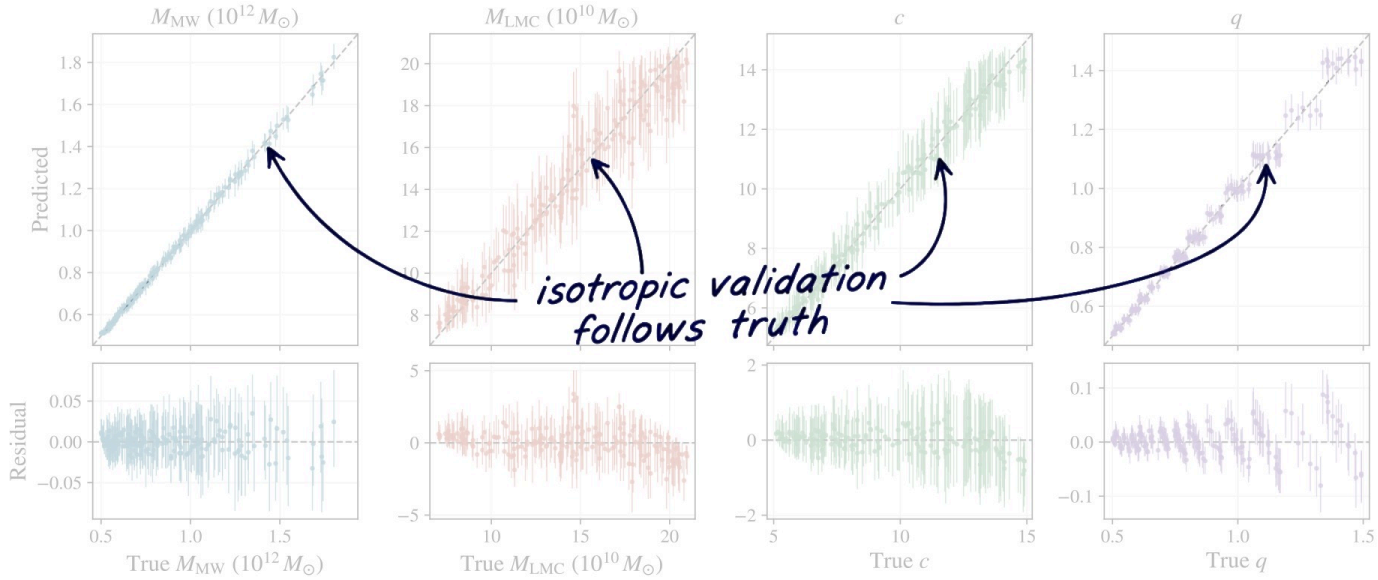


Figure 11. True-vs-predicted accuracy of the MDN-style Gaussian summary-likelihood emulator for the primary joint-summary pipeline on the isotropic 185-grid validation set. The top row shows posterior medians against the true parameters, and the bottom row shows median-minus-true residuals.

Figure 11 validates joint summaries; Figure 12 stress-tests a separate BFE+MOPED pipeline under anisotropy mismatch.

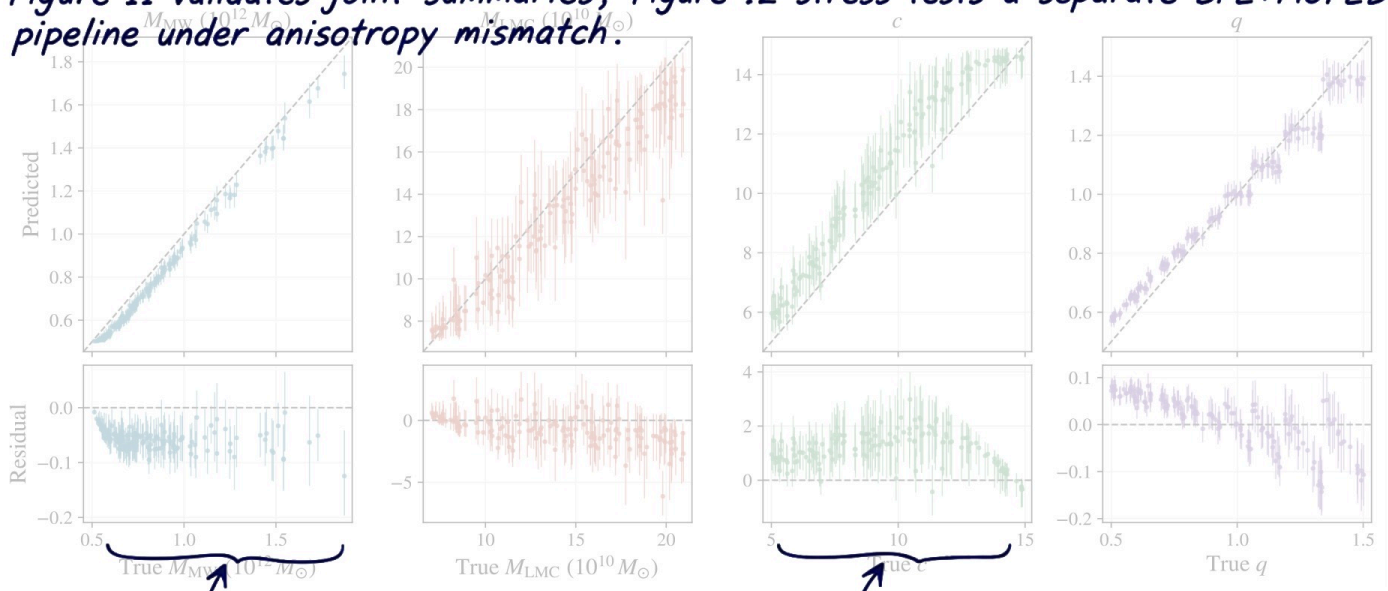


Figure 12. True-vs-predicted accuracy for the anisotropic stress test of the isotropic-trained BFE+MOPED-only pipeline, using the radially varying $\beta(r)$ HaloDance set. Relative to the corresponding isotropic BFE+MOPED validation, the medians shift toward higher c and lower $M_{\rm MW}$, while the shifts in $M_{\rm LMC}$ and q are smaller.

The LMC mass is comparatively robust because its summary relies more on first moments and the density dipole.

The appendix makes the field-level likelihood explicit: normalize coordinates, learn an ODE velocity field, and integrate divergence backward.

response, and $s_{M_{\rm MW}}$ and s_c are dominated by velocity-dispersion monopoles (Section 3.4).

- The primary joint-summary emulator passes the isotropic validation tests on the held-out grids, apart from a small marginal coverage deviation in q. A separate BFE+MOPED-only stress test with radially biased tracers identifies anisotropy as a leading systematic. The $s_{M_{\rm LMC}}$ is least affected, while the inferred $M_{\rm MW}$ and c shift through the second-moment velocity channels (Section 6.3; Appendix D).

The main conclusion is that the outer-halo phase space contains more information than standard velocity moments retain, and that BFE+MOPED summaries recover part of this information in physically interpretable channels. Combining the BFE+MOPED summaries with velocity moments gives the strongest implemented summary-level constraints, but it still falls short of the field-level CFM benchmark. Applying this framework to H3, SDSS-V, or similar catalogues will therefore require systematic tests, including selection functions and tracer-anisotropy nuisance parameters.

ACKNOWLEDGEMENTS

We are grateful to Nico Garavito-Camargo for helpful discussions. YST acknowledges support from NSF Grant AST2406729 and a Humboldt Research Award from the Alexander von Humboldt Foundation. X.-X.X. acknowledges the support from the National Key Research and Development Program of China No. 2024YFA1611902, National Natural Science Foundation of China (NSFC) No. 12588202, CAS Project for Young Scientists in Basic Research grant No. YSBR-062, the Strategic Priority Research Program of Chinese Academy of Sciences grant No. XDB1160102 and grant No. CMS-CSST-2025-A11. We further acknowledge the high-performance computing resources provided by the Australian National Computational Infrastructure (grant y89) through the National and ANU Computational Merit Allocation Schemes.

Large-language-model coding assistants (Anthropic Claude Opus 4.7, OpenAI GPT-5.5) were used during this work to refactor analysis scripts, assist with neural-network fine-tuning, and polish the manuscript draft. The authors reviewed and edited all generated content and take full responsibility for the final manuscript.

DATA AVAILABILITY

The N-body simulation suite developed in this work, HaloDance, will be made publicly available via GitHub at https://github.com/Yanjun-Sheng/HaloDance. The release will include 101 snapshots spanning the past 2 Gyr, which can be used for stellar-stream modelling, halo-kinematic studies, and forward modelling of Gaia observations.

APPENDIX A: CONDITIONAL FLOW MATCHING MODEL DETAILS

The Conditional Flow Matching (CFM) model uses optimal-transport conditional flow matching (OT-CFM) to map a standard Gaussian prior distribution at $t = 0$ to the target parameter-dependent phase-space distribution of the Milky Way (MW) halo dark matter tracer particles at $t = 1$.

For each particle, the input feature vector is $\mathbf{x} = (x, y, z, v_x, v_y, v_z)$, representing Galactocentric positions and velocities, augmented with the Galactocentric radius $r = \sqrt{x^2 + y^2 + z^2}$ and the Galactocentric speed $s = \sqrt{v_x^2 + v_y^2 + v_z^2}$ to incorporate rotation-invariant spatial and velocity information. Coordinates are normalized using standardizing scaling factors derived from the 95th-percentile of the training set:

$$\tilde{\mathbf{x}} = \left(\frac{(x, y, z)}{s_{\rm pos}}, \frac{(v_x, v_y, v_z)}{s_{\rm vel}} \right), \tag{A1}$$

Radius and speed provide rotation-invariant context without replacing the six Cartesian phase-space coordinates.

where $s_{\rm pos}$ and $s_{\rm vel}$ are the 95th-percentile radius and speed respectively.

The vector field $\mathbf{u}_t(\mathbf{x} \mid \theta)$ is parameterised by a neural network. The time variable $t \in [0, 1]$ is represented via sinusoidal Fourier embeddings with 16 log-spaced frequencies spanning $[1, 1000]$, producing a 32-dimensional embedding vector $\gamma(t) = [\sin(2\pi f_k t), \cos(2\pi f_k t)]_{k=1}^{16}$. The physical parameters $\theta = (M_{\rm MW}, M_{\rm LMC}, c, q)$ are first standardized, then passed through a two-layer Multilayer Perceptron (MLP) with layer widths $4 \to 128 \to 256$ and GELU activation to output a 256-dimensional conditioning vector. The core network is a seven-layer MLP with 1,024 hidden units per layer and GELU activations. It takes the concatenated normalized coordinates, time embeddings, and parameter embeddings as input, and outputs a 6D velocity vector $\mathbf{u}_t(\mathbf{x} \mid \theta)$. The total number of parameters is approximately 8.5×10^6.

The OT-CFM straight-line probability path is

$$p_t(\mathbf{x} \mid \mathbf{x}_1) = \mathcal{N}\big(\mathbf{x} \mid t\,\mathbf{x}_1,\ (1-t)^2\,\mathbf{I}\big), \tag{A2}$$

with constant target vector field

$$\mathbf{u}_t(\mathbf{x} \mid \mathbf{x}_1) = \frac{\mathbf{x}_1 - \mathbf{x}}{1 - t}. \tag{A3}$$

Training minimises the conditional flow-matching loss over parameter-space coordinates,

$$\mathcal{L}_{\rm CFM} = \frac{\sum_i w_i \left\| \mathbf{u}_t(\mathbf{x}_t^{(i)} \mid \theta^{(i)}) - (\mathbf{x}_1^{(i)} - \mathbf{x}_0^{(i)}) \right\|^2}{\sum_i w_i}, \tag{A4}$$

where $\mathbf{x}_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, $\mathbf{x}_1$ is a particle coordinate vector from the training set, $t \sim \mathcal{U}(\varepsilon, 1-\varepsilon)$ (with $\varepsilon = 10^{-3}$), and $\mathbf{x}_t = (1-t)\mathbf{x}_0 + t\mathbf{x}_1$. The sample-specific weights w_i implement a curriculum where the loss initially concentrates on parameter points close to the fiducial cosmology, then smoothly broadens to include the whole parameter space. Training uses AdamW with learning rate 10^{-4}, weight decay 10^{-4}, and batch size 512 for 1.5×10^6 steps. The learning rate follows a linear warmup over 5×10^4 steps, remains constant until step 5×10^5, and cosine decays to zero. Gradients are clipped at 1.0.

To compute the log-likelihood of a single particle at a parameter point θ, we integrate the probability flow Ordinary Differential Equation (ODE) backward from $t = 1$ to $t = 0$,

$$\frac{\mathrm{d}}{\mathrm{d}t} \begin{bmatrix} \mathbf{x}_t \\ \log p_t \end{bmatrix} = \begin{bmatrix} \mathbf{u}_t(\mathbf{x}_t \mid \theta) \\ -\operatorname{div}(\mathbf{u}_t)(\mathbf{x}_t \mid \theta) \end{bmatrix}, \qquad \begin{bmatrix} \mathbf{x}_1 \\ \log p_1 \end{bmatrix} = \begin{bmatrix} \tilde{\mathbf{x}} \\ 0 \end{bmatrix}. \tag{A5}$$

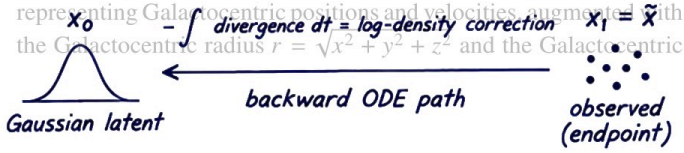
Backward integration accumulates the Jacobian correction that turns a generator into a density estimator.

The divergence $\operatorname{div}(\mathbf{u}_t)$ is approximated using the Hutchinson trace estimator with $n_{\rm H} = 8$ Gaussian probe vectors. The physical coordinate log-likelihood is then

$$\log p(\mathbf{x} \mid \theta) = \log p_0(\mathbf{x}_0) - \int_0^1 \operatorname{div}(\mathbf{u}_t)\,\mathrm{d}t - 3\log s_{\rm pos} - 3\log s_{\rm vel}. \tag{A6}$$

Eight Hutchinson probes trade exact divergence for tractable likelihood evaluation in six dimensions.

Integration uses a fixed-step RK4 ODE solver with 128 steps.



Standardized score gradients and QR define four MOPED summaries; a conditional flow estimates their mutual-information lower bound.

APPENDIX B: BFE AND MOPED COMPRESSION DETAILS

This appendix gives the construction of the linear MOPED projection matrix $\mathbf{Q}_{\rm score}$ mapping high-dimensional Basis Function Expansion (BFE) coefficient vectors $\mathbf{f} \in \mathbb{R}^D$ to the four compressed BFE+MOPED summaries $\mathbf{m} \in \mathbb{R}^P$, with $P = 4$.

Before compression, features are standardized using a diagonal covariance approximation. Using $N_{\rm cov} = 2{,}000$ independent realisations of the MW halo at the fiducial parameter point $\boldsymbol{\theta}_\star = (0.7 \times 10^{12}\,{\rm M}_\odot, 1.5 \times 10^{11}\,{\rm M}_\odot, 9.415, 1.0)$, we compute the per-feature mean vector $\bar{\mathbf{f}}_\star$ and standard deviation vector $\boldsymbol{\sigma}_\star$ for each of the ten channels. The standardized features are

$$\tilde{\mathbf{f}} = \frac{\mathbf{f} - \bar{\mathbf{f}}_\star}{\boldsymbol{\sigma}_\star}. \tag{B1}$$

This per-feature scaling sets the diagonal entries of the covariance matrix to unity. The MOPED construction below then treats the remaining off-diagonal terms as negligible, so that $\mathbf{C} \approx \mathbf{I}$ and the MOPED weight vectors reduce to the scores of the standardized means. We adopt this diagonal approximation for numerical stability: the full $D \times D$ covariance estimated from $N_{\rm cov}$ realisations is rank-deficient and cannot be reliably inverted. The choice is conservative rather than lossless. The neglected off-diagonal correlations are individually small (median absolute value 0.015, 95th percentile 0.046, 99th percentile 0.066, with a maximum of 0.99 from a few strongly coupled modes), but in invertible lower-dimensional subsets including them changes the four-parameter Fisher volume by a factor of 0.15–0.51, so some information is left unrealised.

The full covariance cannot be reliably inverted because features outnumber fiducial realizations by several times.

To test whether this information can be recovered in the full feature space, we regularise the covariance as $\mathbf{C}_\alpha = \alpha\mathbf{I} + (1-\alpha)\mathbf{S}$, where $\mathbf{S}$ is the sample covariance of the standardized features over the $N_{\rm cov}$ fiducial realisations. We choose α by two-fold cross-validation, recomputing $\mathbf{S}$ and the MOPED weights on one half of the realisations and evaluating the Gaussian log-likelihood on the other half, so that the information gain is never scored on the data used to build the compression. The best cross-validated point, $\alpha = 0.9$, tightens the four marginal constraints by factors of 0.97, 0.97, 0.99, and 0.93 and reduces the Fisher volume by 13 per cent, while the median ratio to the CFM benchmark remains 2.44; more aggressive off-diagonal weighting appears tighter in-sample but degrades under cross-validation, indicating covariance overfitting. We therefore retain the diagonal covariance for the production summaries.

Cross-validation rejected aggressive covariance weighting, so we kept the more conservative diagonal approximation.

The score matrix $\mathbf{D} = [\tilde{\boldsymbol{\mu}}_{,1} \mid \cdots \mid \tilde{\boldsymbol{\mu}}_{,P}] \in \mathbb{R}^{D\times P}$ contains the gradients of the mean standardized feature vector with respect to each parameter θ_α. The derivatives are evaluated by central finite differences,

$$\tilde{\boldsymbol{\mu}}_{,\alpha} \approx \frac{\langle\tilde{\mathbf{f}}\rangle_{\boldsymbol{\theta}_\star + \Delta\theta_\alpha\,\hat{\mathbf{e}}_\alpha} - \langle\tilde{\mathbf{f}}\rangle_{\boldsymbol{\theta}_\star - \Delta\theta_\alpha\,\hat{\mathbf{e}}_\alpha}}{2\Delta\theta_\alpha}, \tag{B2}$$

where $\hat{\mathbf{e}}_\alpha$ is the unit vector in the direction of parameter θ_α and $\Delta\theta_\alpha$ is the step size. Rather than running new N-body simulations at these perturbed coordinates, we evaluate the mean feature vectors by averaging over $N_{\rm deriv} = 100$ independent realizations drawn from the trained CFM emulator, which interpolates smoothly across the discrete simulation grid.

With this diagonal covariance approximation ($\mathbf{C} \approx \mathbf{I}$), the Gram–Schmidt orthogonalisation of the score directions is computed directly via QR-decomposition of the score matrix,

$$\mathbf{D} = \mathbf{Q}_{\rm score}\mathbf{R}, \tag{B3}$$

where $\mathbf{Q}_{\rm score} \in \mathbb{R}^{D\times P}$ is an orthonormal matrix whose columns are orthogonal under the identity metric, and $\mathbf{R} \in \mathbb{R}^{P\times P}$ is upper-triangular. The compressed summaries are then

$$\mathbf{m} = \mathbf{Q}_{\rm score}^T(\tilde{\mathbf{f}} - \tilde{\boldsymbol{\mu}}_\star), \tag{B4}$$

where $\tilde{\boldsymbol{\mu}}_\star \equiv \langle\tilde{\mathbf{f}}\rangle_{\boldsymbol{\theta}_\star}$ is the mean standardized feature vector at the fiducial point. A first-order Taylor expansion at parameters near the fiducial point gives

$$\langle\mathbf{m}\rangle \approx \mathbf{Q}_{\rm score}^T\mathbf{D}\,\delta\boldsymbol{\theta} = \mathbf{R}\,\delta\boldsymbol{\theta}, \tag{B5}$$

which links the summaries directly to the parameter deviations $\delta\boldsymbol{\theta} = \boldsymbol{\theta} - \boldsymbol{\theta}_\star$.

APPENDIX C: VARIATIONAL MUTUAL INFORMATION ESTIMATOR DETAILS

This appendix gives the technical details of the variational Mutual Information (MI) estimator used to compare different summary representations.

The mutual-information estimate is a variational lower bound, not an exact equality.

The mutual information $I(\boldsymbol{\theta};\mathbf{s})$ is estimated via the Barber–Agakov variational lower bound (Barber & Agakov 2003; Poole et al. 2019),

$$I(\boldsymbol{\theta},\mathbf{s}) \geq \mathbb{E}_{p(\boldsymbol{\theta},\mathbf{s})}\left[\log q_\phi(\boldsymbol{\theta}\mid\mathbf{s})\right] + h(\boldsymbol{\theta}), \tag{C1}$$

where $h(\boldsymbol{\theta}) = -\mathbb{E}_{p(\boldsymbol{\theta})}[\log p(\boldsymbol{\theta})]$ is the analytic prior entropy. For the uniform prior adopted, the entropy is $h(\boldsymbol{\theta}) = \sum_\alpha \log(\theta_{\alpha,\max} - \theta_{\alpha,\min})$. The bound is tight in the limit where the variational posterior $q_\phi(\boldsymbol{\theta}\mid\mathbf{s})$ approaches the true posterior $p(\boldsymbol{\theta}\mid\mathbf{s})$.

The variational posterior $q_\phi(\boldsymbol{\theta}\mid\mathbf{s})$ is modelled as a conditional flow-matching posterior over the 4D parameter space. Three independent posteriors are trained: (i) $q_\phi(\boldsymbol{\theta}\mid\mathbf{m})$ conditioned on the 4 BFE+MOPED summaries $\mathbf{m}$, (ii) $q_\phi(\boldsymbol{\theta}\mid\mathbf{v})$ conditioned on the 15 velocity moments $\mathbf{v}$, and (iii) $q_\phi(\boldsymbol{\theta}\mid\mathbf{m},\mathbf{v})$ conditioned on the 19 joint summaries. The training dataset uses the same Latin-hypercube parameter grid split as the MDN-style Gaussian likelihood emulator, consisting of 1,663 training points and 185 held-out validation points. For each parameter point, 500 independent 5,000-particle realisations are used to calculate the respective summaries. The vector field $u_\phi(\mathbf{x}_t, t, \mathbf{c}_n)$ parameterising the variational posterior flow is a four-layer MLP with a hidden layer width of 512 and SiLU activations, where $\mathbf{c}_n$ is the conditioning summary vector for the n-th training example. The scalar time t is encoded using a sinusoidal Fourier embedding of 16 frequencies. The training loss minimises:

$$\mathcal{L}_{\rm CFM} = \mathbb{E}\left[\left\|u_\phi(\mathbf{x}_t,t,\mathbf{c}_n) - (\boldsymbol{\theta}_n - \mathbf{z})\right\|_2^2\right], \tag{C2}$$

with $\mathbf{z} \sim \mathcal{N}(0,\mathbf{I})$, $t \sim \mathcal{U}(\varepsilon, 1-\varepsilon)$ (with $\varepsilon = 10^{-3}$), and $\mathbf{x}_t = (1-t)\mathbf{z} + t\boldsymbol{\theta}_n$. Standardized parameter coordinates are used for target flows. We optimize the network using AdamW with a learning rate of 10^{-4} and weight decay of 10^{-4}, with a batch size of 512, for up to 7.5×10^5 steps. The model checkpoint with the minimum validation loss is selected.

To evaluate the lower bound, we compute the log-posterior density of the validation samples under the trained q_ϕ using backward probability flow integration. Because the parameter space is low-dimensional ($D = 4$), the vector-field divergence is computed exactly by summing the diagonal entries of the Jacobian matrix, evaluated using reverse-mode automatic differentiation,

$$\mathrm{div}(\mathbf{u}_t) = \sum_{k=1}^{4}\frac{\partial u_{t,k}}{\partial x_k}. \tag{C3}$$

This is not a Hutchinson trace estimator, which is stochastically noisy Hutchinson trace estimator. The backward integration is solved using a standard Euler solver with 100 steps.

$D = Q_{score} R$

$D = Q_{score} \mid R$ — four orthonormal columns

$Q_{score}^T \tilde{f} \rightarrow m \rightarrow I(\theta;s) \geq E[\log q_\phi(\theta \mid s)] + h(\theta)$

project centered standardized features

MI lower bound

Calibration is tested both parameter by parameter with ranks and jointly with TARP coverage.

APPENDIX D: VALIDATION AND CALIBRATION OF THE SUMMARY SBI PIPELINE

This appendix describes the calibration and accuracy tests for the summary-level Simulation-Based Inference (SBI) pipeline.

We validate the posterior calibration of the individual parameters by evaluating their marginal percentile–percentile (P-P) ranks. For each parameter θ_α and a validation dataset of repeated realisations at known parameter values, we compute the posterior rank of the true value,

$$r_\alpha = \frac{1}{N_{\rm samp}} \sum_{j=1}^{N_{\rm samp}} \mathbb{I}\left[\theta_\alpha^{(j)} < \theta_{\alpha,\rm true}\right], \tag{D1}$$

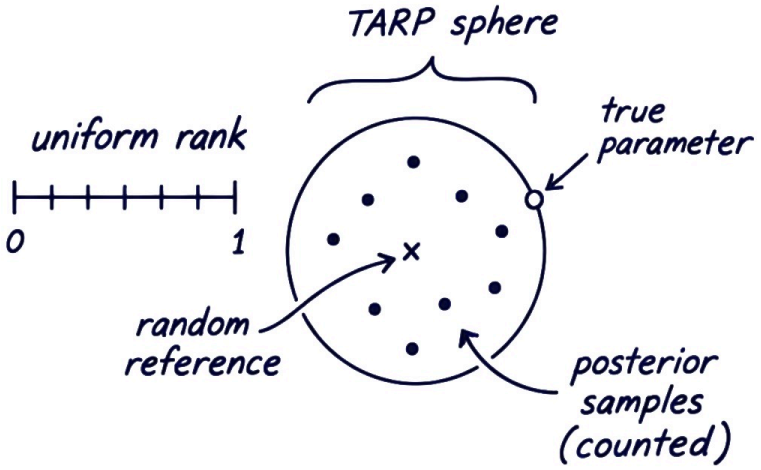
where $\theta_\alpha^{(j)}$ is the value of the parameter in the j-th posterior sample, $\theta_{\alpha,\rm true}$ is the true parameter value, and $\mathbb{I}$ is the indicator function. A calibrated marginal posterior yields ranks that are uniformly distributed between 0 and 1. The P-P curve is the empirical cumulative distribution function (CDF) of these ranks. The top-left panel of Figure D1 shows the marginal P-P plots for $(M_{\rm MW}, M_{\rm LMC}, c, q)$. The Kolmogorov–Smirnov (KS) test values are $p_{\rm KS} = 0.065$, 0.44, and 0.42 for $M_{\rm MW}$, $M_{\rm LMC}$, and c, indicating statistical consistency with perfect calibration. The q parameter shows a slight departure from uniformity ($p_{\rm KS} = 3.1 \times 10^{-9}$), reflecting marginal coverage errors, though the absolute deviation remains small.

Marginal calibration alone can miss incorrect parameter correlations, which is why we also use the joint TARP diagnostic.

To assess whether the posteriors capture the joint parameter correlations, we apply the TARP diagnostic (Lemos et al. 2023). TARP measures the joint coverage by calculating the posterior mass inside a Euclidean sphere centred on an independent random point $\boldsymbol{\theta}_{\rm ref}$, with a radius equal to the distance from $\boldsymbol{\theta}_{\rm ref}$ to the true parameter $\boldsymbol{\theta}_{\rm true}$,

$$\alpha_{\rm TARP} = \frac{1}{N_{\rm samp}} \sum_{j=1}^{N_{\rm samp}} \mathbb{I}\left[\left\|\boldsymbol{\theta}^{(j)} - \boldsymbol{\theta}_{\rm ref}\right\| < \left\|\boldsymbol{\theta}_{\rm true} - \boldsymbol{\theta}_{\rm ref}\right\|\right]. \tag{D2}$$

For a calibrated joint posterior, the distribution of $\alpha_{\rm TARP}$ values across the validation set is uniform. The top-right panel of Figure D1 shows that the empirical TARP curve matches the diagonal within the 95% bootstrap uncertainty, indicating that the joint four-dimensional posterior passes this coverage test.



The flattening ranks depart significantly from uniformity, though the absolute deviation remains small.

We evaluate the point-estimate accuracy using True-vs-Predicted validation grids. Figure 11 shows the posterior medians and 16th-to-84th percentile uncertainties against the true parameters. The median absolute residuals are $0.15 \times 10^{11}\,{\rm M}_\odot$ for $M_{\rm MW}$, $0.79 \times 10^{10}\,{\rm M}_\odot$ for $M_{\rm LMC}$, 0.38 for c, and 0.015 for q. These residuals are smaller than the median 68% posterior widths (0.23, 1.21, 0.63, and 0.020, respectively), indicating that the point estimates recover the true values without systematic offsets.

The marginal posterior medians of the flattening parameter q show a small over-clustering at recurring locations in the P-P ranks and point estimates. To determine whether this crowding is caused by discontinuities in the emulator likelihood, we evaluate the MDN-style Gaussian likelihood ensemble for the joint 19-dimensional summaries on a dense grid scan of q at the fiducial point, varying $q_{\rm true}$ from 0.5 to 1.4 in steps of 10^{-3}. The bottom panels of Figure D1 show that both the likelihood maximum $q_{\rm MAP}$ and the posterior median follow the diagonal smoothly. This rules out discontinuities in the emulator likelihood. The same clustering appears when the validation is repeated with the BFE+MOPED summaries or the velocity moments alone, so it is not introduced by the joint-summary construction. Instead, the clustering reflects parts of the prior where the summaries have weak leverage on q, so the marginal posterior is shaped mainly by the prior and by degeneracies with the other parameters.

A dense scan ruled out likelihood discontinuities as the source of the recurring flattening estimates.

The joint posterior passes the TARP test, while flattening shows a small but real marginal calibration imperfection.

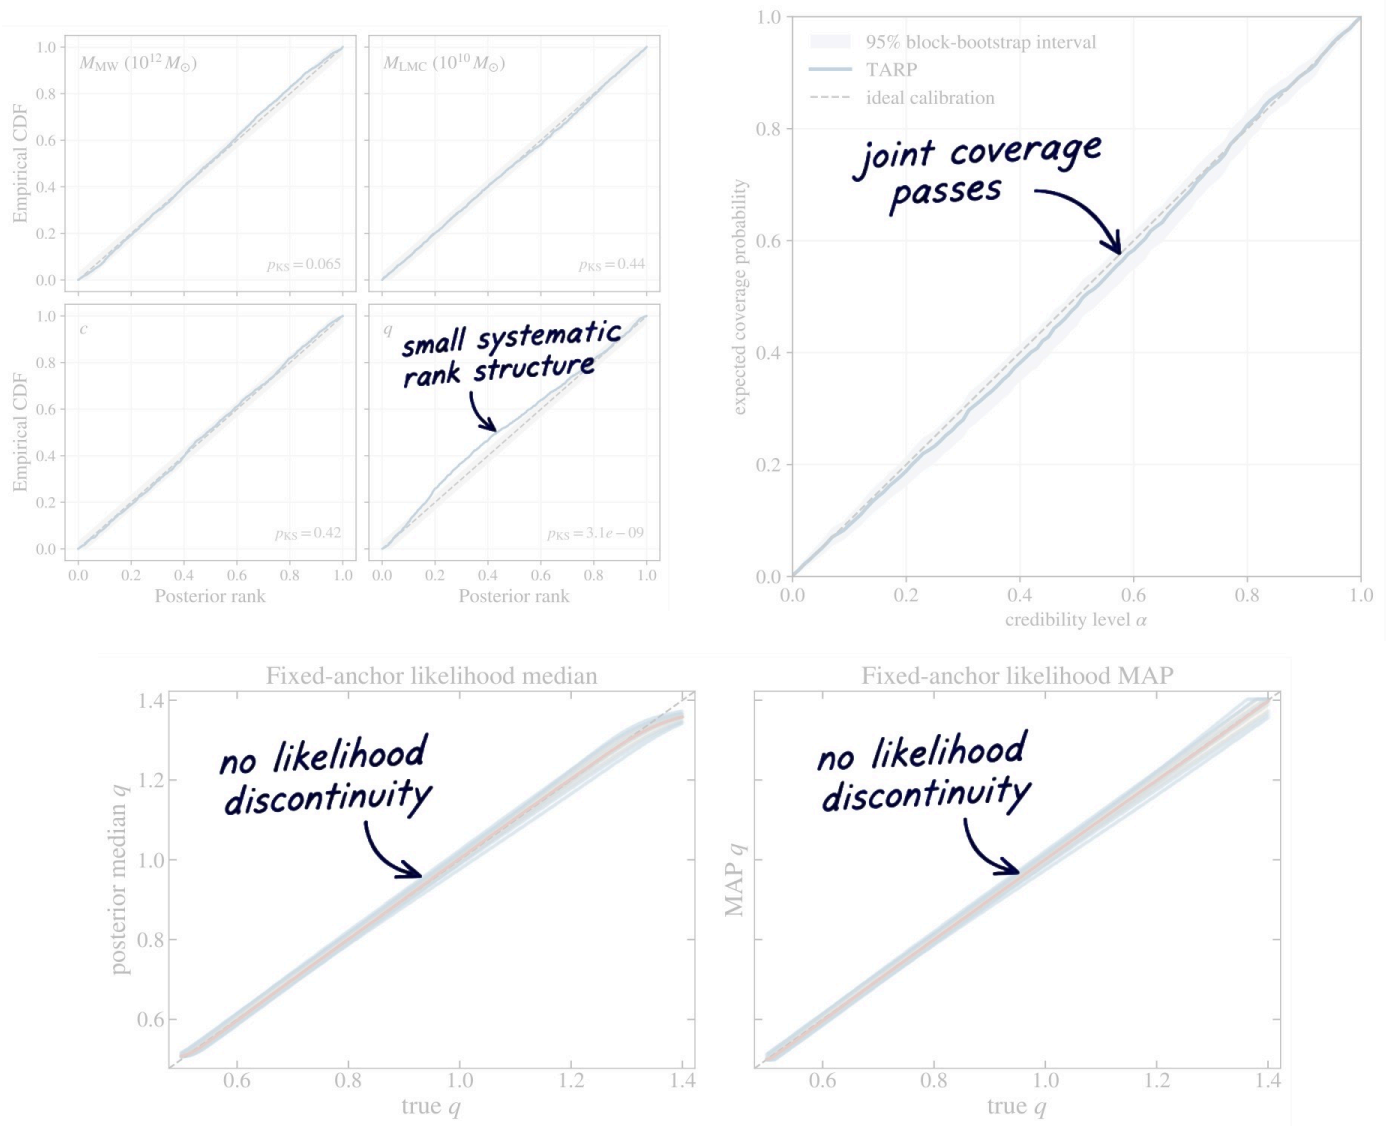


Figure D1. Validation diagnostics for the MDN-style Gaussian summary-likelihood emulator. The top-left panel shows marginal posterior P-P ranks on the validation set, with dashed diagonals for calibrated marginals and grey bands for 95% acceptance regions. The top-right panel shows joint TARP coverage, with the light blue band marking block-bootstrap uncertainty. The bottom panels show the dense q scan at the fiducial anchor, where both the posterior median and likelihood maximum track the one-to-one relation smoothly. The faint lines are the individual repeated realisations at each scanned q, and the solid line is their mean.

Smooth scans rule out likelihood discontinuities; weak leverage, prior influence, and degeneracies drive the flattening pattern.

REFERENCES

Barber D., Agakov F., 2003, in Thrun S., Saul L., Schölkopf B., eds, Advances in Neural Information Processing Systems Vol. 16, Advances in Neural Information Processing Systems. MIT Press, https://proceedings.neurips.cc/paper_files/paper/2003/file/a6ea8471c120fe8cc35a2954c9b9c595-Paper.pdf

Belokurov V., Deason A. J., Erkal D., Koposov S. E., Carballo-Bello J. A., Smith M. C., Jethwa P., Navarrete C., 2019, MNRAS, 488, L47

Besla G., Kallivayalil N., Hernquist L., Robertson B., Cox T. J., van der Marel R. P., Alcock C., 2007, ApJ, 668, 949

Besla G., et al., 2018, MNRAS, 480, 3376

Bland-Hawthorn J., Gerhard O., 2016, ARA&A, 54, 529

Brooks R. A. N., Sanders J. L., Chandra V., Garavito-Camargo N., Dillamore A. M., Price-Whelan A. M., Ting Y.-S., 2025, arXiv e-prints, p. arXiv:2510.04735

Brooks R. A. N., Sanders J. L., Dillamore A. M., Garavito-Camargo N., Chandra V., Price-Whelan A. M., Cargile P., 2026a, MNRAS,

Brooks R. A. N., Sanders J. L., Dillamore A. M., Garavito-Camargo N., Price-Whelan A. M., 2026b, MNRAS, 545, staf2111

Byström A., et al., 2025, MNRAS, 542, 560

Cautun M., et al., 2020, MNRAS, 494, 4291

Cavieres M., et al., 2025, ApJ, 983, 83

Chamberlain K., Price-Whelan A. M., Besla G., Cunningham E. C., Garavito-Camargo N., Peñarrubia J., Petersen M. S., 2023, ApJ, 942, 18

Chandra V., et al., 2025, ApJ, 988, 156

Chandrasekhar S., 1943, ApJ, 97, 263

Conroy C., Naidu R. P., Garavito-Camargo N., Besla G., Zaritsky D., Bonaca A., Johnson B. D., 2021, Nature, 592, 534

Cunningham E. C., et al., 2020, ApJ, 898, 4

Darragh-Ford E., et al., 2025, arXiv e-prints, p. arXiv:2511.02031

Durkan C., Bekasov A., Murray I., Papamakarios G., 2019, arXiv e-prints, p. arXiv:1906.04032

Erkal D., et al., 2019, MNRAS, 487, 2685

Erkal D., et al., 2021, MNRAS, 506, 2677

Foote H. R., et al., 2023, ApJ, 954, 163

Foote H. R., et al., 2026, ApJ, 1001, 37

Fushimi K. J., Mosquera M. E., Dominguez M., 2024, A&A, 688, A147

Garavito-Camargo N., Besla G., Laporte C. F. P., Johnston K. V., Gómez F. A., Watkins L. L., 2019, ApJ, 884, 51

Garavito-Camargo N., Besla G., Laporte C. F. P., Price-Whelan A. M., Cunningham E. C., Johnston K. V., Weinberg M., Gómez F. A., 2021, ApJ, 919, 109

Garver B., Nidever D., Debattista V., Deg N., 2026, arXiv e-prints, p. arXiv:2602.05021

Gómez F. A., Besla G., Carpintero D. D., Villalobos Á., O'Shea B. W., Bell E. F., 2015, ApJ, 802, 128

Hansen S. H., Moore B., 2006, New Astron., 11, 333

Heavens A. F., Jimenez R., Lahav O., 2000, MNRAS, 317, 965

Hernquist L., 1990, ApJ, 356, 359

Ho J., Jain A., Abbeel P., 2020, arXiv e-prints, p. arXiv:2006.11239

Jimenez Rezende D., Mohamed S., 2015, arXiv e-prints, p. arXiv:1505.05770

Kallivayalil N., van der Marel R. P., Alcock C., Axelrod T., Cook K. H., Drake A. J., Geha M., 2006, ApJ, 638, 772

Kallivayalil N., van der Marel R. P., Besla G., Anderson J., Alcock C., 2013, ApJ, 764, 161

Koposov S. E., et al., 2023, MNRAS, 521, 4936

Kravtsov A., Winney S., 2024, The Open Journal of Astrophysics, 7, 50

Laporte C. F. P., Gómez F. A., Besla G., Johnston K. V., Garavito-Camargo N., 2018a, MNRAS, 473, 1218

Laporte C. F. P., Johnston K. V., Gómez F. A., Garavito-Camargo N., Besla G., 2018b, MNRAS, 481, 286

Lemos P., Coogan A., Hezaveh Y., Perreault-Levasseur L., 2023, 40th International Conference on Machine Learning, 202, 19256

Li S., et al., 2026, arXiv e-prints, p. arXiv:2604.01628

Lilleengen S., et al., 2023, MNRAS, 518, 774

Lin D. N. C., Lynden-Bell D., 1982, MNRAS, 198, 707

Lucchini S., Han J. J., Mishra S., Fox A. J., 2025, arXiv e-prints, p. arXiv:2510.03395

Miyamoto M., Nagai R., 1975, PASJ, 27, 533

Murai T., Fujimoto M., 1980, PASJ, 32, 581

Navarro J. F., Frenk C. S., White S. D. M., 1997, ApJ, 490, 493

Peñarrubia J., Gómez F. A., Besla G., Erkal D., Ma Y.-Z., 2016, MNRAS, 456, L54

Petersen M. S., Peñarrubia J., 2020, MNRAS, 494, L11

Petersen M. S., Peñarrubia J., 2021, Nature Astronomy, 5, 251

Petersen M., Weinberg M., 2025, The Journal of Open Source Software, 10, 7302

Poole B., Ozair S., van den Oord A., Alemi A. A., Tucker G., 2019, arXiv e-prints, p. arXiv:1905.06922

Rozier S., Famaey B., Siebert A., Monari G., Pichon C., Ibata R., 2022, ApJ, 933, 113

Sheng Y., Ting Y.-S., Xue X.-X., Chang J., Tian H., 2024, MNRAS, 534, 2694

Sheng Y., Ting Y.-S., Xue X.-X., 2025, MNRAS, 544, 2434

Shipp N., et al., 2021, ApJ, 923, 149

Sohl-Dickstein J., Weiss E. A., Maheswaranathan N., Ganguli S., 2015, arXiv e-prints, p. arXiv:1503.03585

Song Y., Sohl-Dickstein J., Kingma D. P., Kumar A., Ermon S., Poole B., 2020, arXiv e-prints, p. arXiv:2011.13456

Springel V., Pakmor R., Zier O., Reinecke M., 2021, MNRAS, 506, 2871

Sui C., Mao Y., Zhao X., Jing T., Wandelt B. D., 2026, ApJ, 998, 329

Suzuki Y., Chiba M., Wyse R. F. G., 2026, arXiv e-prints, p. arXiv:2604.10281

Tamfal T., Mayer L., Quinn T. R., Capelo P. R., Kazantzidis S., Babul A., Potter D., 2021, ApJ, 916, 55

Tremaine S. D., 1976, ApJ, 203, 72

Vasiliev E., 2023, Galaxies, 11, 59

Vasiliev E., 2024, MNRAS, 527, 437

Vasiliev E., Belokurov V., Erkal D., 2021, MNRAS, 501, 2279

Viterbo G., Buck T., 2026, A&A, 707, A363

Watkins L. L., van der Marel R. P., Bennet P., 2024, ApJ, 963, 84

Yunus M., Wu J. F., Holwerda B. W., 2025, arXiv e-prints, p. arXiv:2512.05078

Yurin D., Springel V., 2014, MNRAS, 444, 62

van der Marel R. P., Kallivayalil N., 2014, ApJ, 781, 121

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