THE LOCAL DYNAMICAL MASS DENSITY FROM HIPPARCOS K GIANTS

A Mixture-Model Jeans-Poisson Analysis with Survey-Completeness Marginalisation

[1]Vidika Mirpuri
[1]Student
[1]The Cathedral and John Connon School, Mumbai, India

Abstract : We measure the local dynamical mass density, ρ_0, from the vertical motions of K giants in the Hipparcos catalogue, as distributed in the public HYG database, using a self-consistent plane-parallel Jeans-Poisson model. The giants are split into three absolute-magnitude groups that share a mixture of isothermal kinematic components but have their own mixing weights, which removes most of the bias that age mixing causes in single-temperature fits. The likelihood is a Poisson process in height times a Gaussian in space velocity, evaluated over the latitude-dependent complete volume of the survey, with the missing radial velocities integrated out exactly. On mock catalogues the pipeline recovers an input density of 0.100 $M_\odot pc^{-3}$ (means 0.093-0.099 for two or three components), while a single temperature overestimates it by about a factor of two. Parallax errors of 0.8 mas do not bias the result. For the real data, our pre-specified fit gives $\rho_0 = 0.063 \pm 0.016$ $M_\odot pc^{-3}$ (statistical). The main systematic is the placement of the survey completeness limit: shifting it by ±0.3 mag changes the result by about ±0.017. Marginalising over this gives $\rho_0 = 0.063$ (+0.023, −0.018) $M_\odot pc^{-3}$, which is below the commonly quoted 0.10 by about 1.5 standard deviations. We find the result stable against sky region, latitude, velocity cuts and priors, but an analysis of A-F stars fails because they are not in equilibrium. The Hipparcos K giants cannot separate a total density of 0.06 from 0.10, and a definitive answer needs Gaia.

***IndexTerms* - Galactic disc, local dark matter density, vertical Jeans equation, Hipparcos, K giants, Bayesian inference, mixture model, MCMC.**

I. INTRODUCTION

The amount of matter in the Galactic disc near the Sun is one of the oldest quantities in Galactic dynamics, and it is still not settled. Oort [1] showed that the vertical motions of stars above and below the plane tell us the gravitational force pulling them back, and from that force one can work out how much mass is there. Applied to a well-defined population of tracer stars, the same idea gives the local dynamical mass density, ρ_0, in $M_\odot pc^{-3}$. Subtracting the mass in stars, gas and remnants then leaves whatever is dark. Direct-detection experiments for dark matter need that residual as an input [2], so the number matters well outside Galactic dynamics.

Early determinations from K dwarfs and K giants disagreed over whether the disc hid a large amount of dark matter [3, 4]. Hipparcos [5] changed the picture because it supplied parallaxes and proper motions for stars that had never had them. Crézé et al. [6] used Hipparcos stars to map the potential well and obtained $\rho_0 = 0.076 \pm 0.015$ $M_\odot pc^{-3}$. Holmberg and Flynn [7], working with a volume-complete sample of A and F stars, found 0.102 ± 0.010, and a later analysis of the surface density [8] stayed close to the visible mass. Korchagin et al. [9] used red giants with radial velocities and reached about 0.1 $M_\odot pc^{-3}$ within 50 pc. More recent work in the Gaia era points to somewhat higher values, around 0.12 [10]. The spread in these numbers is not just random error. It comes from the choice of tracer, from how the survey's completeness is treated, and from whether a single isothermal population can stand in for what is in reality a blend of ages.

This paper goes back to the Hipparcos giants and asks how well the vertical Jeans-Poisson method can be done with only the public HYG compilation of the catalogue [11, 12], so that every step can be repeated by anyone. We changed three things compared with a standard single-population analysis. The first is that K giants in three absolute-magnitude windows are fitted together, sharing a common set of velocity components but with their own mixing weights, so that the age mixture is part of the model. The second is that the volume of each group is built from the documented latitude-dependent completeness of the Hipparcos survey [9, 13] rather than from one magnitude cut. The third is that stars with no radial velocity are kept, using the exact marginal over the missing component, and we test on mock catalogues whether that is safe.

We are going to be plain about what the data can and cannot do. The method recovers the input density on mocks, but the real catalogue gives an answer that moves by an amount comparable to its statistical error when the completeness limit is shifted by a few tenths of a magnitude. We therefore report the result that we fixed before running the fit, then show the full sweep of that systematic, a marginalised value that folds it in, and a set of consistency checks that include sky splits, prior changes and convergence tests.

The paper is organised as follows. Section II sets out the equilibrium model. Section III describes the catalogue, the survey volume and the tracer groups. Section IV derives the likelihood, including the treatment of missing radial velocities, and the sampling scheme. Section V tests the full pipeline on mock catalogues, including ones with parallax errors. Section VI presents the results and every robustness test. Section VII discusses limitations and Section VIII concludes.

II. EQUILIBRIUM MODEL

A. Vertical Jeans equation

Let $f(z, v_z)$ be the phase-space density of a tracer population in a static, plane-parallel slab, with z the height above the mid-plane and v_z the vertical velocity. In steady state the collisionless Boltzmann equation reduces to $v_z\, \partial f/\partial z - (d\Phi/dz)\, \partial f/\partial v_z = 0$, where $\Phi(z)$

is the gravitational potential. Taking the first velocity moment, and writing $\nu(z) = \int f \, dv_z$ for the tracer number density and σ_W^2 for the vertical dispersion, one obtains the Jeans equation [14]

$$d(\nu\, \sigma_W^2)/dz + \nu\, d\Phi/dz = 0, \text{ i.e. } d(\nu\, \sigma_W^2)/dz = \nu\, K_z \tag{1}$$

where $K_z = -d\Phi/dz$ is the vertical force. The equation links three things that can all be estimated from tracer stars: the density profile, the dispersion profile, and the force. Once K_z is known, Poisson's equation in the thin-disc limit, $dK_z/dz = -4\pi G\, \rho(z)$, turns it into the total mass density. We neglect the radial and tilt terms, which are small close to the Sun for a disc whose scale length is many times the survey depth [4, 15].

B. Mass model and isothermal tracers

For the total mass density we adopt a sech² profile, $\rho(z) = \rho_0 \operatorname{sech}^2(z/z_\rho)$ (Eq. (2) below). This is the exact solution for an isothermal self-gravitating sheet, and it is smooth and positive, which suits a sampler. The vertical force follows from integrating Poisson's equation, and the potential is

$$\rho(z) = \rho_0 \operatorname{sech}^2(z/z_\rho) \tag{2}$$

$$\Phi(z) = 4\pi G\, \rho_0\, z_\rho^2 \ln \cosh(z/z_\rho) \tag{3}$$

with $G = 4.30091 \times 10^{-3}$ pc (km s^{-1})^2 M$_\odot^{-1}$. A tracer population that is isothermal with dispersion σ_W satisfies Eq. (1) with a constant σ_W, so that its density is

$$\nu_k(z) \propto \exp[\, -\Phi(z) / \sigma_{W,k}^2 \,] \tag{4}$$

up to a normalisation. The surface density within $|z| < h$ and the vertical force are derived quantities,

$$\Sigma(|z|<h) = 2\, \rho_0\, z_\rho \tanh(h/z_\rho) \tag{5}$$

$$K_z(h) = 4\pi G\, \rho_0\, z_\rho \tanh(h/z_\rho) \tag{6}$$

where for z_ρ much larger than h the force is proportional to ρ_0 alone, so it stays well defined even when z_ρ is poorly constrained.

C. Why a single temperature is not enough

A sample of K giants contains stars with a wide range of ages, and vertical dispersion rises with age. A sum of isothermal components with dispersions $\sigma_{W,k}$ and fractions w_k has a density that falls off more slowly than any single isothermal component of matching mean dispersion, because the hot component puts stars at large height that a single colder component would not. If this blend is fitted with a single σ_W, the fit sees a profile that is too extended for its σ_W and compensates by making the potential too deep. The resulting ρ_0 is biased high. The effect is a consequence of the convexity of $\exp(-\Phi/\sigma^2)$ as a function of $1/\sigma^2$, and we confirm its size on mocks in Section V. We therefore model the tracer density as a mixture,

$$\nu_g(z) = \Sigma_k\, w_{gk}\, \nu_k(z)/I_{gk}, \quad \Sigma_k\, w_{gk} = 1 \tag{7}$$

with the w depending on the magnitude group g because brighter giants are on average younger and more massive, while the dispersions and mean velocities are shared. This shared-kinematics assumption is what lets three groups pin down the mixture, and it is also its main simplification.

D. From density to dark matter

The total density includes stars, gas and any dark component. The model of McKee et al. [16] gives a total of 0.097 M$_\odot$ pc^{-3} with 0.013 of it dark, which leaves about 0.084 in stars, gas and remnants. We write the implied dark matter density as $\rho_{DM} = \rho_0 - \rho_b$, with $\rho_b = 0.084$ M$_\odot$ pc^{-3} as a fixed reference value. The baryon inventory is itself uncertain at a level comparable to ρ_{DM}, so we regard this as an illustration of the leverage of the measurement rather than a result.

III. Data and Sample Selection

A. Catalogue

We use HYG v4.1 [11], which carries the Hipparcos new reduction [12] for about 118,000 stars together with Gliese and Yale Bright Star entries. We keep entries that have a Hipparcos identifier, a parallax-based distance and a *B-V* colour. HYG gives distance, *V*-band apparent and absolute magnitude, proper motion and, where available, radial velocity, but no per-star uncertainties. A radial velocity of exactly zero means "not available" in HYG, and the value 9999.99 in the proper-motion columns marks a missing proper motion. Both are treated as missing. Distances are used as given, so selection is on the observed distance and absolute magnitude, exactly as it would be for any user of this table. We return to the effect of parallax errors in Section V-D.

B. Coordinates and velocities

Equatorial positions and proper motions are converted to Galactic coordinates (l, b) with astropy [17]. With d in parsecs and proper motions in mas yr^{-1} the tangential velocities are $v_l = 4.74047\, d\, \mu_{l*}/10^3$ and similarly for b, in km s^{-1}. The observed vector for a star is $y = (v_l, v_b, v_r)$, and it is related to the Galactic heliocentric velocity $v = (U, V, W)$ by $y = \mathbf{R}\, v$, with

$$\mathbf{R} = [\, (-\sin l, \cos l, 0); (-\sin b \cos l, -\sin b \sin l, \cos b); (\cos b \cos l, \cos b \sin l, \sin b) \,] \tag{8}$$

where the rows are the unit vectors along increasing l, increasing b and the line of sight. A star's height is $Z = d \sin b$ relative to the Sun, and its height above the mid-plane is $z = Z + z_0$. The solar motion is absorbed into the fitted component means, with $W_\odot \approx 7.25$ km s^{-1} [18].

C. Survey completeness and the sample volume

The Hipparcos Input Catalogue was built to be complete to a limiting magnitude that depends on spectral type and on Galactic latitude. For stars redder than $B\text{-}V = 0.8$ the documented limit is $V \leq 7.3 + 1.1 \sin|b|$ [9, 13], with the survey about 93 per cent complete

relative to Tycho in the regions where it is tested [9]. For a group with absolute magnitudes up to M_{hi} the completeness condition on distance is

$$d_{max}(b) = \min[\, d_{cap},\, D_0\, 10^{1.1 \sin|b|/5} \,] \tag{9}$$

with $D_0 = 10^x$ pc, $x = (v_0 - M_{hi} + 5)/5$, and $v_0 = 7.3$. We truncate at $d_{cap} = 200$ pc, remove $|b| < 10°$ where extinction and crowding matter most, and keep stars inside the resulting volume. The volume is not a sphere. The cross-sectional area at a height Z, which we need to normalise the vertical profile, is

$$A_g(Z) = \pi\, [\, d_{max}^2 - Z^2 \,] \;\; \text{(limited by } |b| \geq b_{min}) \tag{10}$$

where $d_{max}(b)$ is the limit above. We obtain the largest s at each Z by bisection on $d_{max}(b(s, Z))$, and then apply the latitude cut through $s \leq |Z| \cot b_{min}$ (Fig. 1). Using a sphere instead is not harmless. A constant-limit analysis of the same stars gave 0.12 $M_\odot\,pc^{-3}$ for a limit of 7.0 and 0.08 for 7.3, which is how we found out that the latitude dependence had to be modelled.

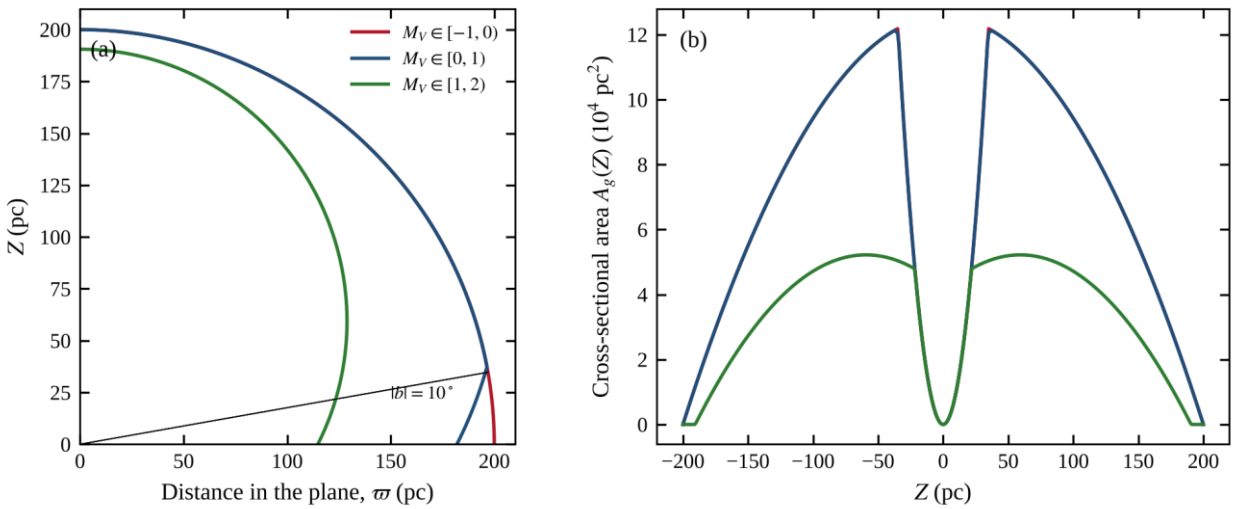

Fig. 1. (a) Boundary of the complete volume in the (ϖ, Z) plane for the three M_V windows at the adopted limit. (b) Cross-sectional area $A_g(Z)$ entering the normalisation of Eq. (10). The jump near $|Z| \approx 30$ pc is where the $|b| \geq 10°$ cut stops binding.

D. Tracer groups

The tracers are K giants with $0.9 < B\text{-}V < 1.6$, in three windows of absolute magnitude, $M_V \in [-1, 0)$, $[0, 1)$ and $[1, 2)$. The brighter windows are younger on average. Together they cover a range of ages while all lying on the red-giant branch and red clump, and all being visible to 200 pc. Stars with a tangential speed above 150 km s^{-1}, or a radial velocity above the same value, are removed as halo and thick-disc contaminants. Table I lists the sample. The decision about the definition of the volume, v_0 and the slope 1.1, was taken before any fit using it was run, following the cited documentation. The group boundaries and the cuts were also chosen before the fit.

Table I. K-giant sample used in the primary fit

Group	M_V window	N	With RV	d_{max} in plane (pc)
1	[−1, 0)	464	458	200
2	[0, 1)	2054	1204	181
3	[1, 2)	1067	454	115

d_{max} is D_0 of Eq. (9), the completeness distance in the Galactic plane. At higher latitude the limiting magnitude is fainter and the volume reaches further, up to the 200 pc cap.

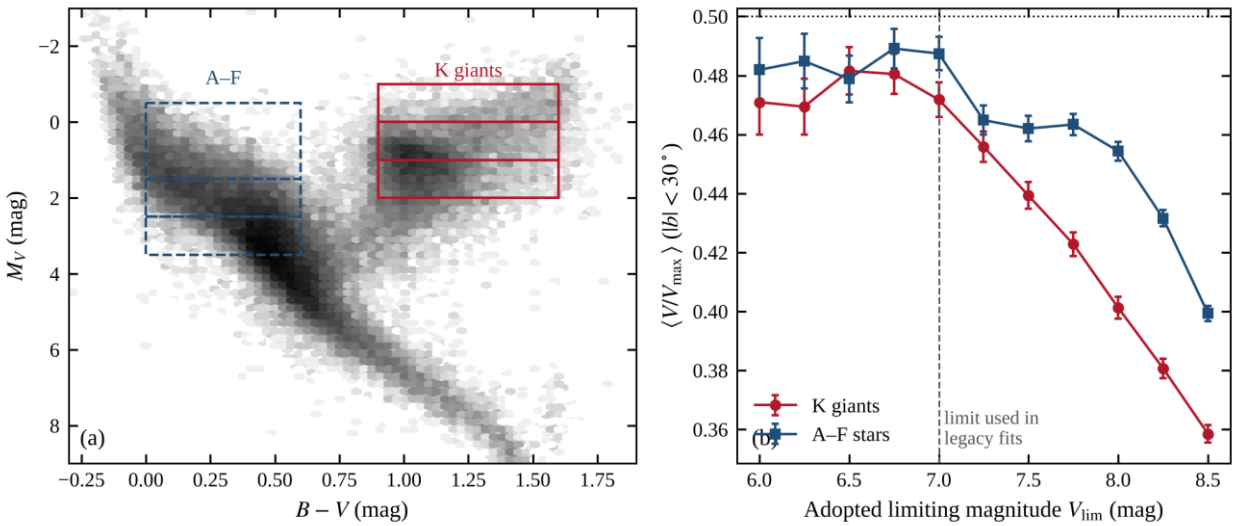

Fig. 2. (a) Colour-magnitude diagram of Hipparcos stars with distances in HYG, with the K-giant boxes (solid) and the A-F boxes used as a diagnostic (dashed). (b) Mean V/V_{max} at low latitude for a spherical volume with a constant limiting magnitude V_{lim}. A complete sample sits near 0.5. The decline beyond $V \approx 7$ for a constant limit motivated modelling the latitude-dependent limit.

IV. Statistical Model

A. Spatial likelihood

Stars in a given group are treated as a realisation of an inhomogeneous Poisson process in position, with intensity proportional to the equilibrium density $\nu(z)$ inside the complete volume V and zero outside. For a single component the likelihood of the observed heights $z_1, \ldots, z_n$ is

$$\ln L_z = \Sigma_i \ln [\, \nu(z_i) / \textstyle\int_V \nu(z)\, dV \,] \tag{11}$$

where the amplitude N_0 has been profiled out, which is the standard result when the total count carries no information on the shape. The normalising integral is taken over the same volume as the selection. With the survey geometry of Section III-C it reduces to a one-dimensional integral over the area function,

$$I_{gk} = \int \exp[\, -\Phi(Z+z_0) / \sigma_{W,k}^2 \,]\, A_g(Z)\, dZ \tag{12}$$

which we evaluate with the trapezoidal rule on a grid of 601 heights. The accuracy of this quadrature was checked against a ten-times finer grid and the integrated volume against Monte Carlo sampling; the two agree to about 10^{-5}, and the integrated volume agrees with a Monte Carlo estimate to 0.1 per cent.

B. Velocity likelihood and missing radial velocities

Given component k with mean m_k and diagonal dispersion tensor $\Sigma_k = \mathrm{diag}(\sigma_U^2, \sigma_V^2, \sigma_W^2)$, the observed vector $y_i = \mathbf{R}_i v$ is Gaussian with mean $\mathbf{R}_i m_k$ and covariance $\mathbf{R}_i \Sigma_k \mathbf{R}_i^T$. We do not rotate stars to Galactic velocities, because that would need the radial velocity, which about half the sample lacks. Instead, for a star without a radial velocity the radial row of $\mathbf{R}$ is dropped, and the likelihood uses the 2×2 block of the covariance. This is the exact marginal over the unobserved radial velocity, since marginalising a Gaussian only deletes the corresponding row and column.

This marginalisation is valid when the chance that a star has a radial velocity depends only on quantities that enter the likelihood, namely its position and proper motion. That condition is called missing at random [19]. It would fail if radial velocities had been preferentially measured for stars with unusually large or small radial speed, which is a real possibility for a patchwork of historical programmes, and we return to it in Section VI-H. The full likelihood for star i in group g is

$$\ell_i = \Sigma_k\, w_{gk} \cdot \exp[\, -\Phi(z_i) / \sigma_{W,k}^2 \,] / I_{gk} \cdot \square(y_i;\, \mathbf{R}_i m_k,\, \mathbf{R}_i \Sigma_k \mathbf{R}_i^T) \tag{13}$$

with w_{gk} the weight of component k in group g and $\square$ the Gaussian density. Eq. (13) is a proper mixture, not a product of separate spatial and velocity terms, because the vertical dispersion that sets the spatial profile is the same $\sigma_{W,k}$ that enters the velocity part. This coupling is what gives the method its leverage on the potential.

C. Priors and parametrisation

The parameter vector contains ρ_0, z_ρ, z_0, six numbers per component (three means and three dispersions) and $G(K - 1)$ mixing parameters. The weights use a stick-breaking construction, $w_{g1} = f_{g1}$, $w_{g2} = f_{g2}(1 - f_{g1})$, and so on, with each f uniform on (0, 1). The components are ordered by $\sigma_{W,1} < \sigma_{W,2} < \ldots$ to remove label switching. All other priors are uniform: ρ_0 between 0.02 and 0.30 $M_\odot$ pc^{-3}, z_ρ between 150 and 700 pc, z_0 between −10 and 60 pc, component means within ±80 km s^{-1} and dispersions between 2 and 80 km s^{-1}. The $K = 2$ model has 18 parameters and the $K = 3$ model 27.

D. Sampling and convergence

We sample with the affine-invariant ensemble sampler of Goodman and Weare [20] as implemented in emcee [21], but with a mixture of differential-evolution [22] and snooker moves (80 and 20 per cent), which handled the strongly correlated mixture parameters better. There are three walkers per parameter, started in a tight ball around a Nelder-Mead optimum. We use the integrated autocorrelation time τ to judge convergence, and require the chain to be longer than about 50 τ. For the adopted $K = 2$ fit we ran two

independent chains with different seeds and compared them with the potential scale reduction factor of Gelman and Rubin [23]. A single evaluation of the full likelihood takes about 2 ms.

E. Model comparison

To choose the number of components we compare the maximum log-likelihood, the Bayesian information criterion BIC = $k \ln n - 2 \ln L_{max}$ [24] and the Akaike criterion AIC = $2k - 2 \ln L_{max}$, with k the number of free parameters and n the number of stars. These are used as guides, not as proof, because the stars are not independent draws from a perfectly specified model.

V. Validation on Mock Catalogues

A. Construction

Before using the pipeline on the catalogue we tested it on synthetic stars. Mock stars are drawn from the equilibrium model with a known $\rho_0 = 0.100$ $M_\odot\,pc^{-3}$, $z_\rho = 350$ pc and $z_0 = 10$ pc. Positions are drawn by rejection from $\exp(-\Phi/\sigma_{W,k}^2)$ inside the volume, velocities from the component Gaussians, and the three groups contain 1,500, 2,000 and 1,000 stars inside spheres of radius 200, 158 and 100 pc, with the same M_V windows in mind. Radial velocities are marked missing with a probability that depends on proper motion, with fewer radial velocities for stars with small proper motion, so that the missingness is a function of an observed quantity (MAR) and not trivially random. We used two truths: two components with dispersions of 10 and 22 km s^{-1}, and four components from 7 to 24 km s^{-1}. Each truth was fitted with $K = 1$, 2 and 3, using six independent mock catalogues for each combination, so that a truth with four components tests the response to a model that is too simple. A further set of six mocks uses the survey-shaped volume of Section III-C and is described in Section V-C.

B. Recovery of the input density

Table II and Fig. 3(a) show the mock fits. With $K = 2$ and $K = 3$ the recovered density averages 0.093-0.099 $M_\odot\,pc^{-3}$ for an input of 0.100, the scatter of the recovered medians is about the same as the quoted uncertainty, and the rms pull lies between 0.88 and 1.15. All four means fall slightly below the input, by 1 to 7 per cent, which is a fraction of one standard error. With six catalogues per row we cannot tell whether that is noise or a small downward bias, and we carry it as a possible systematic of order 5 per cent. A truth with four components that is fitted with two shows no worse behaviour than the matched case. This is what the structure of the model suggests, since the potential is constrained mainly by the vertical profile of the whole sample, and the mixture only needs to reproduce its overall shape.

The single-component fit is a different story. When the truth is a mixture and we fit $K = 1$, the recovered density averages 0.182 for two components and 0.197 for four, close to twice the input. In all twelve catalogues the estimate is too high, by between 1.0 and 3.7 standard errors. This is the bias that motivated the mixture model, and it is large. The agreement of a single-temperature fit to the real data with the literature value should therefore not be read as support for the literature value.

Table II. Recovery of the input density on mock catalogues (true $\rho_0 = 0.100$ $M_\odot\,pc^{-3}$)

True model	Fitted K	Mocks	Mean ± scatter of medians	Mean 1σ	rms pull
Two components	1	6	0.182 ± 0.025	0.047	1.82
Two components	2	6	0.099 ± 0.024	0.027	0.89
Two components	3	6	0.094 ± 0.028	0.025	1.01
Four components	1	6	0.197 ± 0.037	0.041	2.53
Four components	2	6	0.093 ± 0.020	0.024	0.88
Four components	3	6	0.096 ± 0.031	0.024	1.15
Two components, survey volume	2	6	0.105 ± 0.014	0.015	0.95

Pull = (median − 0.100)/σ, with σ half the 68 per cent interval of each mock. A well-calibrated pipeline has rms pull near 1 and a scatter of medians close to the mean 1σ.

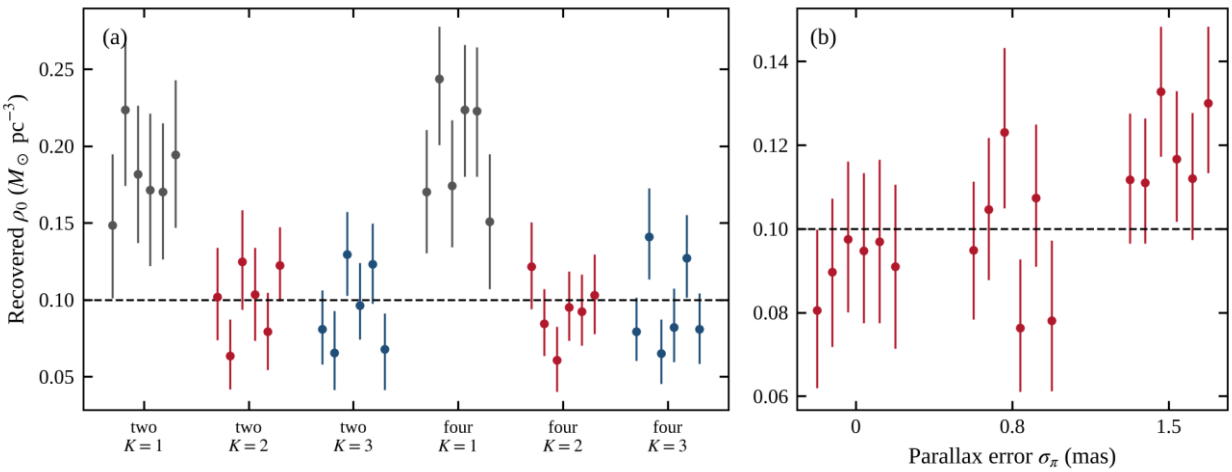

Fig. 3. (a) Recovered ρ_0 (median and 68 per cent interval) for each mock catalogue, by true model (two or four components) and fitted number of components K. The dashed line is the input value 0.100 $M_\odot$ pc^{-3}. (b) The same for mocks with Gaussian parallax errors, with the sample selected on the observed distance.

C. Survey-shaped volume

The mocks above use spherical volumes, whereas the real sample lives in the latitude-dependent volume of Eq. (9). We therefore generated six further catalogues with the exact survey geometry for the three windows of Table I, with about 800, 2,400 and 820 stars per group, a two-component truth with the same parameters as before, and fitted them with $K = 2$. The recovered densities average 0.105 with a scatter of 0.014, for an input of 0.100, and the rms pull is 0.95. Here the mean lies about 5 per cent above the input, and in the spherical $K = 2$ mocks about 1 to 7 per cent below, so there is no consistent sign and we treat the pipeline bias as within about ±5 per cent, small next to the statistical error of 0.016 on the real data.

D. Effect of parallax errors

HYG carries no parallax errors. The Hipparcos catalogue has a median parallax uncertainty of about 1 mas [13], and the new reduction improves the accuracy of stars brighter than $H_p = 8$ by up to a factor of four [12], so for our stars the errors are probably a few tenths of a mas, which is a few per cent of the parallax at 150-200 pc. The selection in our sample is made on the observed distance and absolute magnitude, so errors move stars across the boundaries of the volume and spread them in height. The likelihood does not model this, so we tested its effect directly. We generated mock catalogues with Gaussian parallax errors of 0, 0.8 and 1.5 mas, selected them on the observed distance in the same way, and fitted them with the error-free likelihood. Six catalogues were run for each value. The mean recovered density is 0.092 with no error (a control, with the same spherical geometry as the earlier mocks), 0.097 at 0.8 mas and 0.119 at 1.5 mas, with rms pulls of 0.54, 0.96 and 1.31 (Fig. 3(b)). At 0.8 mas, which is a conservative value for this sample, we see no significant bias. Larger errors push the answer upwards. The sign matters here: errors of this size cannot be what makes the real-data density low, and if anything they would bias a measurement of the real catalogue high. These mocks use a spherical volume, not the survey-shaped one, so we regard them as an indication and not a calibration. We do not apply any correction for the Lutz-Kelker bias [25].

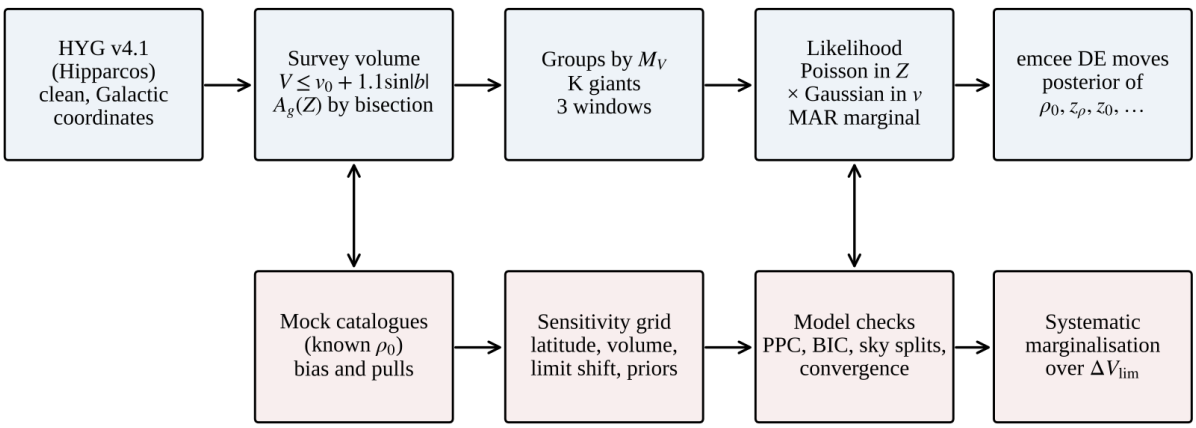

Fig. 4. Outline of the analysis. The upper row is the inference pipeline. The lower row lists the validation and robustness tests described in Sections V and VI, each of which feeds back on the pipeline.

VI. Results

A. Primary fit

Table III. Posterior summary of the adopted $K = 2$ fit

Parameter	Median and 68 per cent interval
ρ_0 ($M_\odot$ pc^{-3})	0.063 ± 0.016
z_ρ (pc)	401 ± 192
z_0 (pc)	8.2 ± 10.2
Cold component, $\langle U \rangle$ (km s^{-1})	−5.9 ± 0.7
Cold component, $\langle V \rangle$ (km s^{-1})	−11.0 ± 0.5
Cold component, $\langle W \rangle$ (km s^{-1})	−6.3 ± 0.4
Cold component, σ_U (km s^{-1})	26.4 ± 0.7
Cold component, σ_V (km s^{-1})	15.1 ± 0.5
Cold component, σ_W (km s^{-1})	11.7 ± 0.4
Hot component, $\langle U \rangle$ (km s^{-1})	−14.5 ± 1.6
Hot component, $\langle V \rangle$ (km s^{-1})	−33.0 ± 1.4
Hot component, $\langle W \rangle$ (km s^{-1})	−9.1 ± 0.9
Hot component, σ_U (km s^{-1})	43.5 ± 1.1
Hot component, σ_V (km s^{-1})	27.4 ± 0.7
Hot component, σ_W (km s^{-1})	25.5 ± 0.8
Cold-component weight, $M_V \in [-1, 0)$	0.67 ± 0.04
Cold-component weight, $M_V \in [0, 1)$	0.66 ± 0.03
Cold-component weight, $M_V \in [1, 2)$	0.56 ± 0.03

The fit uses 3585 stars. Intervals are 16th-84th percentiles of 48600 retained samples. The cold component is the one with the lower σ_W.

The pre-specified $K = 2$ fit gives $\rho_0 = 0.063 \pm 0.016$ $M_\odot$ pc^{-3}, with a 95 per cent interval of 0.033-0.094. All parameters are listed in Table III. The vertical dispersions of the two components are 11.7 ± 0.4 and 25.5 ± 0.8 km s^{-1}, and the cold component carries 0.67 ± 0.04, 0.66 ± 0.03 and 0.56 ± 0.03 of the stars in the three groups from brightest to faintest. The trend with absolute magnitude is in the expected direction, with fainter giants being older and hotter, although the change between the first two groups is small. The Sun's height is $z_0 = 8 \pm 10$ pc, which is consistent with the 20.8 ± 0.3 pc of Bennett and Bovy [26] at about 1.2 standard deviations, and fixing it to that value changes ρ_0 by only a few thousandths (Section VI-G).

The scale height z_ρ is not constrained inside the prior range, and its posterior is nearly flat between 150 and 700 pc (Fig. 6). This is expected. The survey reaches 200 pc, so the potential is probed only over the part of the sech² profile where it is close to parabolic, and over that range the data constrain the combination ρ_0 rather than ρ_0 and z_ρ separately. The posterior of ρ_0 is uncorrelated with z_ρ, as the corner plot shows. For the same reason the extrapolated surface density depends on the prior for z_ρ, and we quote only the well-defined derived quantities, the surface density and force at 100 pc, in Section VI-I.

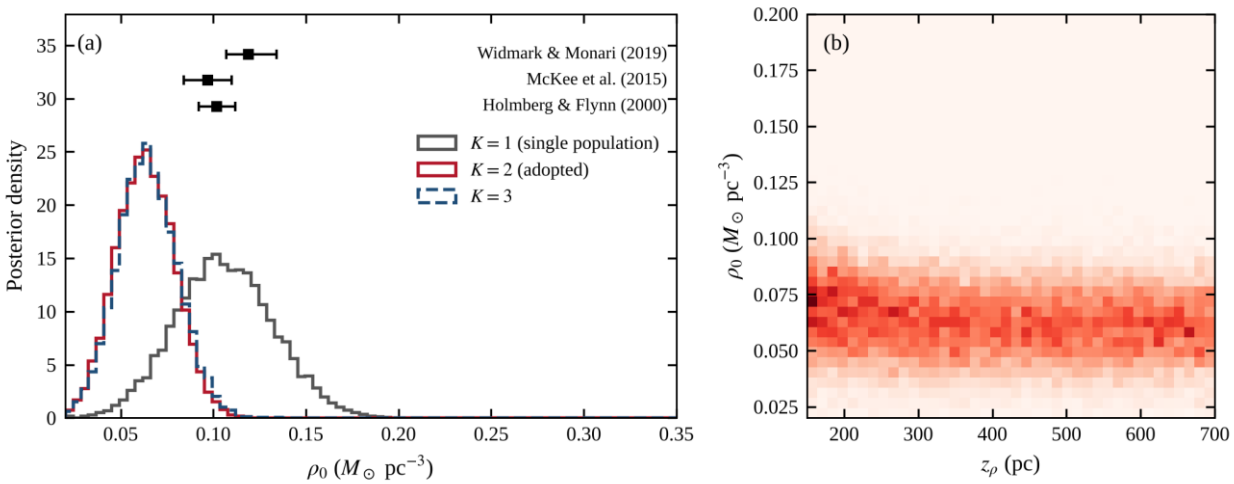

Fig. 5. (a) Posterior of ρ_0 for $K = 1$, 2 and 3 components, with three literature values for comparison. (b) Joint posterior of z_ρ and ρ_0 for $K = 2$.

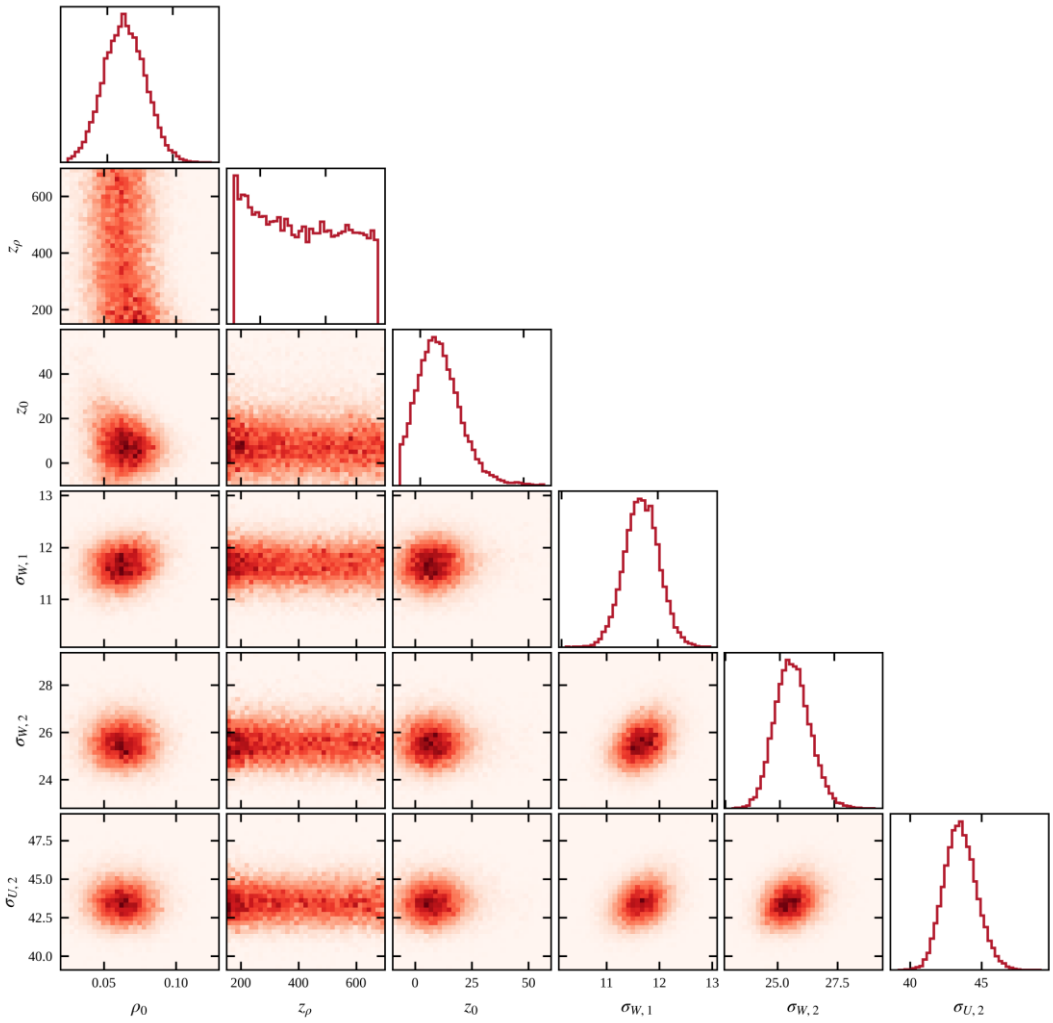

Fig. 6. Marginal and pairwise posteriors for a subset of the $K = 2$ parameters: the density ρ_0, scale height z_ρ, solar height z_0, vertical dispersions of the two components and the radial dispersion of the hot component.

B. Posterior predictive checks and kinematics

Fig. 7 compares the observed and predicted height distributions in each group and the vertical dispersion as a function of height. The model reproduces the vertical distribution of each group, including the dip near the plane produced by the $|b|$ cut, and the flat dispersion profile of the stars that have radial velocities, at about 17-19 km s^{-1} against a model value of about 18 km s^{-1}. The observed profile is flat to within its errors, so the data do not directly show the weak rise with height that a mixture of components predicts, and we draw no conclusions from the small differences.

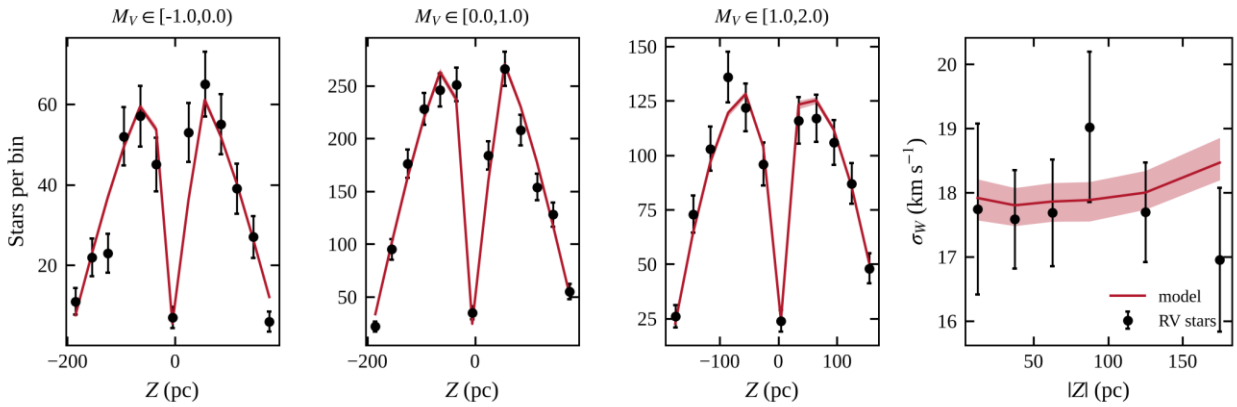

Fig. 7. Observed (points) and posterior predictive (lines and shading, 16-84 per cent) height distributions for the three M_V groups and the vertical dispersion σ_W against $|Z|$ for stars with radial velocities.

The velocity part of the model can be checked against known kinematics. The cold component has a mean W of -6.3 ± 0.4 km s^{-1} and the hot one -9.1 ± 0.9, both consistent with the reflex of the solar vertical motion of about 7 km s^{-1} [18]. The hot component lags the cold one by 22.0 ± 1.4 km s^{-1} in V. The classical asymmetric drift relation [27, 28] predicts a lag of about $(\sigma_{U,\mathrm{hot}}^2 - \sigma_{U,\mathrm{cold}}^2)/80$, which here is 14.9 ± 1.3 km s^{-1}. The sign and size agree, and the observed lag is about 7 km s^{-1} larger than the relation, which we do not consider alarming, because the relation is only approximate and our components are mixtures of ages and not single populations.

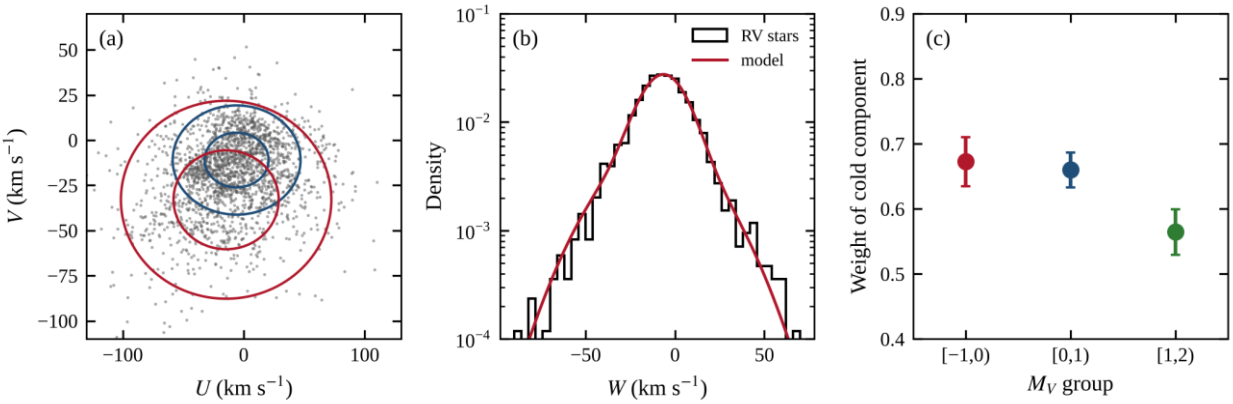

Fig. 8. (a) Heliocentric (U, V) of the stars with radial velocities, with the 1σ and 2σ ellipses of the cold (blue) and hot (red) components. (b) Histogram of W with the fitted mixture. (c) Weight of the cold component in each magnitude group.

C. Number of components

Table IV. Model comparison for the number of components

Model	Parameters	ln L_{max}	BIC	AIC	ρ_0 ($M_\odot pc^{-3}$)
$K = 1$	9	-103548.1	207169.9	207114.3	0.105 ± 0.026
$K = 2$	18	-103115.2	206377.7	206266.3	0.063 ± 0.016
$K = 3$	27	-103064.6	206350.2	206183.2	0.064 ± 0.016

Lower BIC and AIC are preferred. Differences of about 10 or more are considered strong.

Table IV compares one, two and three components. Going from one to two improves ln L by 433, and BIC falls by 792, which is decisive. The step to three components improves ln L by a further 51. BIC falls by 28 and AIC by 83, so both criteria prefer three components, and by the usual scales strongly. The density does not care: $K = 3$ gives 0.064 ± 0.016, the same as $K = 2$. We keep $K = 2$ as the adopted model because it was specified in advance, because it samples cleanly, and because the quantity of interest is unchanged when a third component is added. The $K = 3$ chain is much harder to sample, as we describe next.

D. Convergence

The two independent $K = 2$ chains give $\rho_0 = 0.063 \pm 0.016$ and 0.063 ± 0.016, and the potential scale reduction factor is at most 1.001 over all 18 parameters (it is 1.000 for ρ_0). The longest autocorrelation time is 108 and 110 steps for the two chains, which are 84 and 127 times shorter than the run lengths of 9,000 and 14,000 steps. The $K = 3$ chain is much worse. Its longest autocorrelation time is 1438 steps, against a run length of 12,000, so it is not converged by our own criterion. The marginal posterior of ρ_0 for $K = 3$ agrees with the converged $K = 2$ result, but we use it only as a consistency check.

E. Dependence on the completeness limit

Table V. Dependence of ρ_0 on the shift ΔV of the completeness limit

ΔV (mag)	N per group	ρ_0 ($M_\odot pc^{-3}$)
−0.5	464 / 1778 / 476	0.097 ± 0.019
−0.4	464 / 1859 / 557	0.090 ± 0.018
−0.3	464 / 1942 / 660	0.083 ± 0.018
−0.2	464 / 2003 / 779	0.078 ± 0.017
−0.1	464 / 2040 / 922	0.064 ± 0.016
0.0	464 / 2054 / 1067	0.063 ± 0.016
+0.1	464 / 2054 / 1223	0.058 ± 0.014
+0.2	464 / 2054 / 1406	0.052 ± 0.014
+0.3	464 / 2054 / 1581	0.050 ± 0.013
+0.4	464 / 2054 / 1746	0.050 ± 0.013
+0.5	464 / 2054 / 1892	0.051 ± 0.013

A positive shift moves the limit fainter, so that more stars at large height and distance enter the sample. $\Delta V = 0$ is the pre-specified analysis.

The dominant systematic in the analysis is the placement of the completeness limit. We repeated the primary fit with the limit moved in steps of 0.1 mag between −0.5 and +0.5 (Table V and Fig. 9(a)). The density falls steadily as the limit is moved fainter, from 0.097 ± 0.019 at −0.5 mag, where the sample holds only the nearer stars, to 0.051 ± 0.013 at +0.5 mag, and flattens beyond about +0.2. Over the range ±0.3 mag the density changes by about 0.03 $M_\odot pc^{-3}$, which is twice the statistical error. The sample sizes grow smoothly with the limit, so the counts alone do not show where the survey stops being complete. The direction of the trend is informative. Moving the limit fainter mostly adds stars at high latitude and large height, and these are more numerous than a high-

density isothermal model predicts, so that the profile looks flatter and the fitted potential shallower. Simple incompleteness near the limit would act the other way, because it removes distant stars and steepens the profile, so it is not an obvious explanation. A real flattening of the vertical profile, for instance from a larger thick-disc contribution or a departure from equilibrium at those heights, or a slightly wrong slope in the completeness law, would fit better. We cannot separate these possibilities with this catalogue.

To fold this in, we marginalise over a shift with a Gaussian prior of mean 0 and width 0.3 mag, truncated to the grid, by combining the posteriors in proportion to the prior weights. The width of 0.3 mag reflects how well the limit is known for K giants [9, 13], and it was fixed before the marginalisation. The result is $\rho_0 = 0.063$ (+0.023, −0.018) $M_\odot\,pc^{-3}$, with a 95 per cent interval of 0.031-0.111 (Fig. 9(b)). Widths of 0.2 and 0.5 mag give medians of 0.063 and 0.064, so the answer is not sensitive to this choice. The median does not move because the posterior at $\Delta V = 0$ sits near the middle of the sweep, but the interval widens on the high side. This is the number we regard as the main result of the paper.

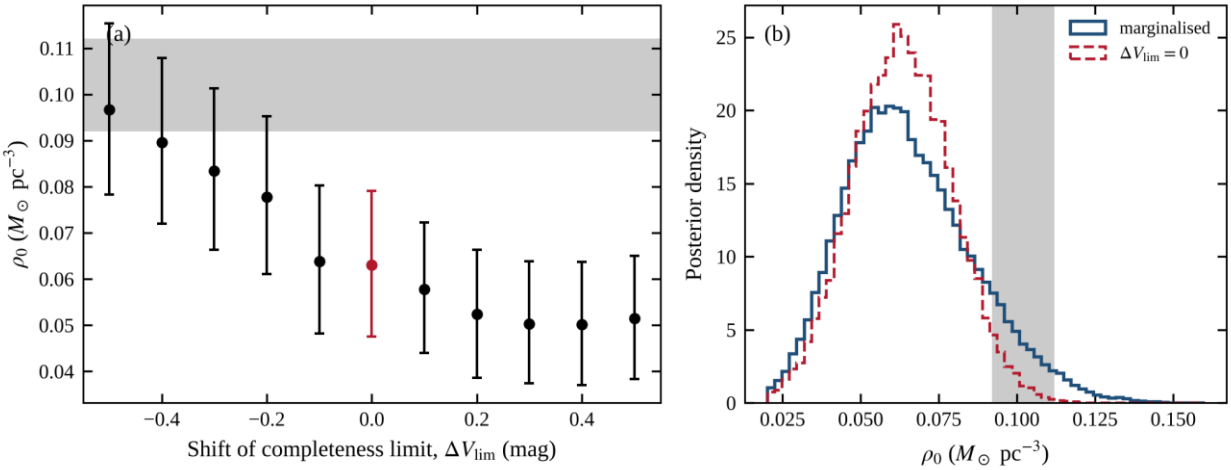

Fig. 9. (a) Posterior median and 68 per cent interval of ρ_0 as a function of the shift of the completeness limit. The red point is the pre-specified analysis, the grey band the Holmberg and Flynn [7] value. (b) Posterior at zero shift (dashed) and after marginalising over a Gaussian prior of width 0.3 mag on the shift.

F. Other sensitivity tests

Table VI. Sensitivity of ρ_0 to selection and sampling choices ($K = 2$)

Variant	N per group	ρ_0 ($M_\odot\,pc^{-3}$)
Adopted	464 / 2054 / 1067	0.063 ± 0.016
Second chain, different seed	464 / 2054 / 1067	0.063 ± 0.016
$\lvert b\rvert \geq 20°$	374 / 1593 / 930	0.055 ± 0.017
$\lvert b\rvert \geq 30°$	276 / 1211 / 778	0.059 ± 0.020
$\lvert b\rvert \geq 40°$	182 / 854 / 600	0.058 (+0.025, −0.021)
$d < 150$ pc	242 / 989 / 813	0.100 ± 0.036
$d < 120$ pc	127 / 538 / 448	0.083 (+0.049, −0.037)
Constant $V_{\rm lim} = 7.3$, spherical	464 / 1647 / 382	0.080 ± 0.022
$\lvert v\rvert < 100$ km s^{-1}	456 / 2019 / 1041	0.062 ± 0.015
$\lvert v\rvert < 250$ km s^{-1}	465 / 2058 / 1068	0.063 ± 0.015
Windows [−0.5, 0.5), [0.5, 1.5), [1.5, 2.5)	912 / 2408 / 342	0.085 ± 0.016
Colour $1.0 < B\text{-}V < 1.5$	322 / 1374 / 803	0.063 ± 0.020

Errors are 68 per cent credible intervals.

The remaining tests are listed in Table VI and shown in Fig. 10. Restricting the sample to higher latitudes, where extinction is smaller and the volume is shaped differently, leaves the result at 0.055-0.059, so dust close to the plane is not what lowers it. Velocity cuts at 100 and 250 km s^{-1} change nothing at the level of 0.001, which shows that the halo and thick-disc contamination at the fastest end is not driving the answer. Truncating the volume at 150 or 120 pc gives higher central values but with much larger errors, 0.100 and 0.083, and these are consistent with the adopted value. Shifting the M_V windows by half a magnitude gives 0.085, and a narrower colour range gives the same as the adopted fit. A spherical volume with a constant limit of 7.3 gives 0.080. None of these differs from the adopted value by more than the uncertainty of the variant itself, and several of the variants that raise the density do so by shrinking the volume.

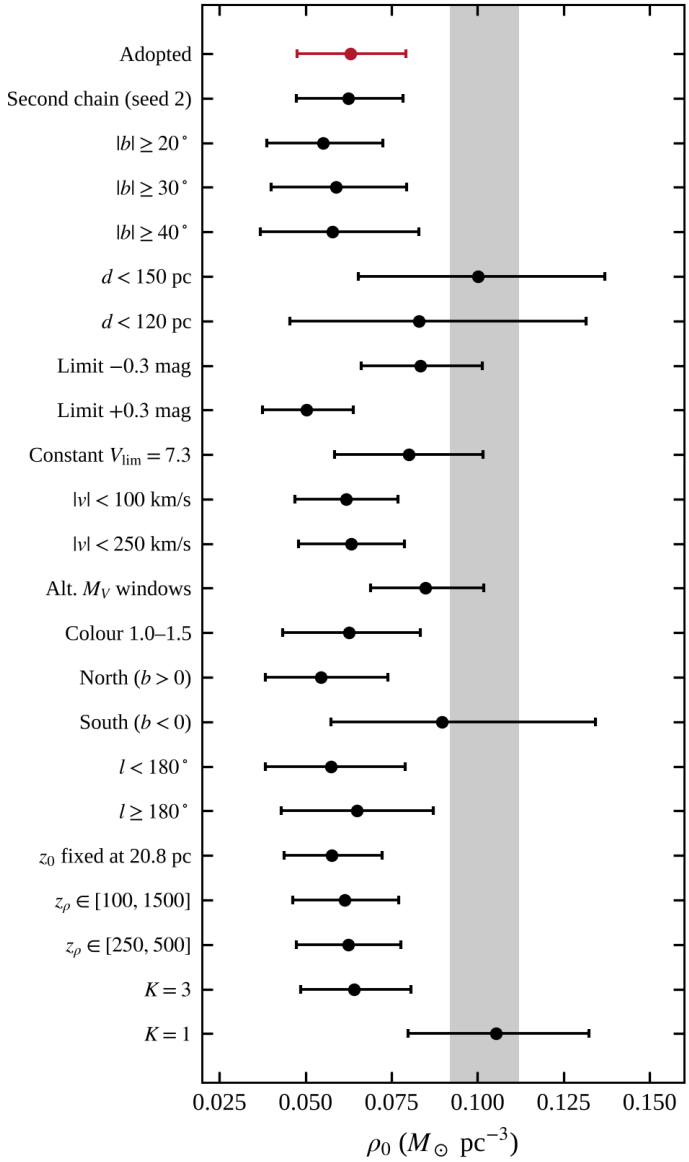

Fig. 10. Posterior median and 68 per cent interval of ρ_0 for every analysis variant. The grey band is the Holmberg and Flynn [7] value.

G. Sky splits and priors

Table VII. Consistency across the sky and sensitivity to priors (K = 2)

Variant	N per group	ρ_0 ($M_\odot$ pc^{-3})	z_0 (pc)
North of the plane (b > 0)	247 / 1004 / 507	0.054 (+0.020, −0.016)	28.5 ± 23.1
South of the plane (b < 0)	217 / 1050 / 560	0.090 (+0.045, −0.032)	29.4 (+20.5, −25.1)
l < 180°	226 / 1037 / 540	0.057 ± 0.020	13.8 (+17.1, −13.1)
l ≥ 180°	238 / 1017 / 527	0.065 ± 0.022	6.7 (+14.5, −10.3)
z_0 fixed to 20.0-21.6 pc	464 / 2054 / 1067	0.058 ± 0.014	20.8 ± 0.5
z_ρ prior 100-1500 pc	464 / 2054 / 1067	0.061 ± 0.015	8.4 ± 9.8
z_ρ prior 250-500 pc	464 / 2054 / 1067	0.062 ± 0.015	7.6 (+10.7, −8.9)

The hemisphere fits use the survey volume restricted to one side of the plane, and the longitude fits to half of the azimuth.

If the model is adequate, independent parts of the sky should give the same density. The two hemispheres give 0.054 (+0.020, −0.016) and 0.090 (+0.045, −0.032), which differ by 0.8 standard deviations, and the two longitude halves give 0.057 ± 0.020 and 0.065 ± 0.022, which differ by 0.3. There is no sign of an inconsistency with the adopted value, though the single-hemisphere fits prefer larger z_0 (about 28 pc each) than the full sample. The vertical offset and the density are partly degenerate when only one side of the plane is observed, and a sample that spans both sides is more reliable for the solar height. Changing the priors makes very little difference (Table VII). Fixing z_0 to the Bennett and Bovy [26] value gives 0.058 ± 0.014, and widening or narrowing the range of z_ρ gives 0.061 ± 0.015 and 0.062 ± 0.015. The posterior is therefore driven by the data, not the priors.

H. Tests that failed

Two further tests did not work, and we report them because they bear on how far the method can be trusted. The first was to apply exactly the same machinery to A and F stars ($0 < B\text{-}V < 0.6$, in three windows of absolute magnitude from −0.5 to 3.5, with the limit for earlier types $v_0 = 7.9$), which are the tracers of Holmberg and Flynn [7]. The fit returns $\rho_0 = 0.143$ (+0.010, −0.008) with $z_0 = -1.3 \pm 1.4$ pc, and a restriction to 100 pc gives 0.138 ± 0.024 with $z_0 = -5.7$ (+3.4, −2.7) pc. These values are well above the published 0.102 and imply a negative solar height, which contradicts the independently known value. The statistical error of the first fit is under 10 per cent, so the answer is confidently wrong. A and F stars are young. They trace the Gould Belt and local spiral structure [29] and are not in vertical equilibrium, so the isothermal assumption fails for them. The fit's confident but wrong answer in this case is a reminder that the posterior width only reflects statistical error within the model.

The second test was a fit using only stars with measured radial velocities, which is the cleanest data but not a complete sample. It returns $\rho_0 = 0.289$ (+0.008, −0.014), at the upper edge of the prior, and a negative z_0. In HYG the availability of a radial velocity depends on a star's brightness and on historic observing programmes, so the stars with a radial velocity are not a volume-complete subset, and our geometry is wrong for them. We think this is the reason, but we did not test it directly. It also means the assumption of missingness at random is doing real work in the adopted analysis, and we have no way of checking it with HYG alone. The mocks in Section V assume it holds.

I. Derived quantities

Within the observed volume the data constrain the integrated mass better than the central density. For the adopted fit the surface density is $\Sigma(|z| < 50\text{ pc}) = 6.3 \pm 1.6$ and $\Sigma(|z| < 100\text{ pc}) = 12.2 \pm 3.0\ M_\odot\,pc^{-2}$, and the vertical force at 100 pc is $K_z = 330 \pm 82$ (km $s^{-1})^2$ kpc^{-1}. Korchagin et al. [9], using Hipparcos red giants, reported $10.5 \pm 0.5\ M_\odot\,pc^{-2}$ within 50 pc, which is higher than our value by about 2.5 standard deviations, with the caveat that their tracer sample and method differ from ours in several ways. Taking the visible matter to be $0.084\ M_\odot\,pc^{-3}$ [16], the implied dark matter density for the adopted fit is -0.021 ± 0.016, which is negative at the 1.3σ level, and for the marginalised result it is -0.021 (+0.023, −0.018). With the limit of 0.111 on the total density at 95 per cent, the data permit at most about $0.03\ M_\odot\,pc^{-3}$ of dark matter, which does not exclude the usual 0.01 and does not confirm it either. We stress that this is not a dark-matter measurement. The uncertainty on ρ_0 is larger than the quantity it is being asked to pin down.

J. Comparison with the literature

Table VIII. Comparison with published local dynamical densities

Study	Tracer and data	ρ_0 ($M_\odot\,pc^{-3}$)
Crézé et al. [6]	Hipparcos stars	0.076 ± 0.015
Holmberg and Flynn [7]	A-F stars, Hipparcos	0.102 ± 0.010
Korchagin et al. [9]	Red giants, Hipparcos + RVs	≈ 0.1 within 50 pc
McKee et al. [16]	Baryon and dark matter inventory	0.097 ± 0.013
Widmark and Monari [10]	Gaia DR1 stars, Jeans equation	0.119 (+0.015, −0.012)
This work, adopted	K giants, Hipparcos (HYG)	0.063 ± 0.016
This work, marginalised over limit	K giants, Hipparcos (HYG)	0.063 (+0.023, −0.018)
This work, single temperature ($K = 1$)	K giants, Hipparcos (HYG)	0.105 ± 0.026

Literature values are quoted as given in the cited papers. They use different tracers and methods.

Table VIII places the result among published values. The adopted value is consistent with Crézé et al., who also used Hipparcos. With statistical errors alone it lies 2.1 standard deviations below Holmberg and Flynn when the two errors are combined in quadrature. After marginalising over the completeness limit the marginalised value is 1.7 standard deviations below Holmberg and Flynn and 2.3 below Widmark and Monari, and about one and a half below 0.10. The tension with the Gaia-based value is the larger one, and we cannot rule out that it is real, although the Gaia analysis relies on different tracers and a much cleaner selection. We therefore do not claim that Hipparcos K giants disagree with a density near $0.1\ M_\odot\,pc^{-3}$, only that they do not prefer it. The most we can say is that with this catalogue and this method a value between 0.03 and 0.11 is allowed at 95 per cent.

VII. Discussion

The result of this paper is more modest than the method deserves. The pipeline passes its tests: it recovers the input density on mocks with different numbers of components, with parallax errors of the size expected for these stars, and with a survey-shaped volume. The chains converge, the posterior predictive checks look reasonable, and the kinematic parameters are physically sensible. The real-data answer, however, depends on a quantity that we cannot pin down with this catalogue, which is where the survey stops being complete.

A. Why the completeness limit matters so much

The vertical density profile is the main handle on the potential, and it is read from the distribution of heights inside the survey volume. Close to the plane the volume is bounded by the $|b|$ cut and by the 200 pc cap. Far from the plane it is bounded by the magnitude limit, which at high latitude lies beyond 200 pc for the brighter groups but cuts the faintest group at 115-190 pc. When the limit is moved by 0.3 mag the number of stars in the faintest group changes by 40 to 50 per cent, and the density changes with it. As Section VI-E explains, the sign of the change is the opposite of what simple incompleteness would give: the stars that are added when the limit moves fainter lower the density. The survey is stated to be about 93 per cent complete relative to Tycho [9], so incompleteness cannot be ruled out, but it does not explain the trend on its own. The distance errors in the catalogue may matter too, since the absolute-magnitude windows are defined from the observed distance, and stars near the window edges are moved in and

out by errors that grow with distance. We did not model this, and it is a plausible contributor. A catalogue with a measured selection function and small distance errors, which Gaia provides, would remove these ambiguities.

B. Limitations

HYG inherits Hipparcos, so it gives no per-star errors, no flags for binaries or variable stars, and no information on the way radial velocities were assembled. Binary motion can inflate proper motions and thereby the dispersions of some stars. The model also assumes a plane-parallel, static disc, which ignores the radial terms in the Jeans equation [15] and the vertical waves of a few km s^{-1} in amplitude found in the solar neighbourhood [26]. An imbalance of a few km s^{-1} in the mean vertical velocity as a function of height would mimic a change of potential, and the model does not allow for it. Each component is isothermal, whereas the true dispersion rises slowly with height, and the shared-kinematics assumption means all three groups are forced to draw on the same set of component velocity distributions. We saw no sign that this causes difficulty in the posterior predictive check, but a group-dependent mean velocity is a natural extension. Lastly, the mocks are generated from the same family of models as the fit, with an isothermal component structure and Gaussian velocities. They test the code and the likelihood, not whether the real disc follows that structure.

The results also depend on our use of observed distances in the selection. This is how the catalogue is used in practice, and our mocks show that with errors of 0.8 mas the bias is small. Errors at the larger end of what Hipparcos has would raise the estimate, as the 1.5 mas mocks show, so a correction for them, if we could make one, would go towards a lower density and not a higher one.

C. What would be needed

The cleanest improvement is to repeat the analysis with Gaia, where parallax errors are an order of magnitude smaller at these distances and the completeness of bright stars is characterised. A second is to put an explicit completeness function into the likelihood, for example a smooth decline in the selection probability near the limit with its shape fitted jointly, instead of a sharp edge with a systematic shift. With Hipparcos alone this extra freedom would be only weakly constrained, as the sweep shows, but it would make the dependence part of the posterior and not a separate step. Finally, a tracer population with a known age structure, for example red-clump stars, would let the mixture be tied to physical ages.

VIII. Conclusions

We have built a Bayesian Jeans-Poisson analysis for Hipparcos K giants that handles the age mixture with a shared-kinematics mixture of isothermal components, the latitude-dependent completeness of the survey with an exact volume integral, and missing radial velocities by exact marginalisation. We tested it on mock catalogues. The conclusions are these.

First, the method works. On mocks it recovers a true density of 0.100 $M_\odot pc^{-3}$ with two or three components (means 0.093-0.099), whereas a single-temperature fit overestimates it by about a factor of two. Parallax errors of 0.8 mas leave the estimate unbiased within the scatter, and errors of 1.5 mas push it up by about 20 per cent.

Second, for the real catalogue, our pre-specified analysis gives $\rho_0 = 0.063 \pm 0.016$ $M_\odot pc^{-3}$ (statistical). The chains are converged, with $\hat{R} \leq 1.001$, and the result is stable against latitude, velocity and prior choices, and consistent between hemispheres and longitudes. It is however sensitive to the completeness limit of the survey. Marginalising over a shift of 0.3 mag in that limit gives 0.063 (+0.023, −0.018) $M_\odot pc^{-3}$, with a 95 per cent interval of 0.031-0.111. This is lower than the commonly quoted 0.10 but consistent with it at about 1.5 standard deviations, and consistent with the earlier Hipparcos value of Crézé et al.

Third, the single-temperature fit gives 0.105 and the A-F stars give 0.14, in line with or above the literature, but the first is known from the mocks to be biased high and the second fails for physical reasons. Agreement with a published number is not evidence that the number is right, and we do not claim it.

Fourth, the data do not allow a statement about dark matter. The uncertainty on the total density is larger than the amount of dark matter it is supposed to reveal. We conclude that with this method the legacy Hipparcos K giants cannot give a total density tighter than about ±0.03 $M_\odot pc^{-3}$, because the answer depends on how the stars far from the plane enter the sample at a level comparable to the statistical error. A definitive measurement needs the better-characterised completeness and precision of Gaia.

References

[1] J. H. Oort, "The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems," *Bull. Astron. Inst. Netherlands*, vol. 6, pp. 249-287, 1932.

[2] J. I. Read, "The local dark matter density," *J. Phys. G: Nucl. Part. Phys.*, vol. 41, 063101, 2014.

[3] J. N. Bahcall, "Self-consistent determinations of the total amount of matter near the Sun," *Astrophys. J.*, vol. 276, pp. 169-181, 1984.

[4] K. Kuijken and G. Gilmore, "The mass distribution in the galactic disc - II. Determination of the surface mass density of the galactic disc near the Sun," *Mon. Not. R. Astron. Soc.*, vol. 239, pp. 605-649, 1989.

[5] M. A. C. Perryman *et al.*, "The HIPPARCOS Catalogue," *Astron. Astrophys.*, vol. 323, pp. L49-L52, 1997.

[6] M. Crézé, E. Chereul, O. Bienaymé and C. Pichon, "The distribution of nearby stars in phase space mapped by Hipparcos: I. The potential well and local dynamical mass," *Astron. Astrophys.*, vol. 329, pp. 920-936, 1998.

[7] J. Holmberg and C. Flynn, "The local density of matter mapped by Hipparcos," *Mon. Not. R. Astron. Soc.*, vol. 313, pp. 209-216, 2000.

[8] J. Holmberg and C. Flynn, "The local surface density of disc matter mapped by Hipparcos," *Mon. Not. R. Astron. Soc.*, vol. 352, pp. 440-446, 2004.

[9] V. I. Korchagin, T. M. Girard, T. V. Borkova, D. I. Dinescu and W. F. van Altena, "Local surface density of the Galactic disk from a 3-D stellar velocity sample," *Astron. J.*, vol. 126, pp. 2896-2909, 2003.

[10] J. Widmark and G. Monari, "Local dark matter density and updated constraints on the Milky Way disc from Gaia DR1 and the vertical Jeans equation," *Mon. Not. R. Astron. Soc.*, vol. 482, pp. 262-277, 2019.

[11] D. Nash, "HYG stellar database, version 4.1," https://github.com/astronexus/hyg-database (CC BY-SA 4.0), accessed 2026.

[12] F. van Leeuwen, "Validation of the new Hipparcos reduction," *Astron. Astrophys.*, vol. 474, pp. 653-664, 2007.

[13] ESA, *The Hipparcos and Tycho Catalogues*, ESA SP-1200. Noordwijk, The Netherlands: ESA Publications Division, 1997.

[14] J. Binney and S. Tremaine, *Galactic Dynamics*, 2nd ed. Princeton, NJ, USA: Princeton Univ. Press, 2008.

[15] J. Bovy and H.-W. Rix, "A direct dynamical measurement of the Milky Way's disk surface density profile, disk scale length, and dark matter profile at 4 kpc $\lesssim$ R $\lesssim$ 9 kpc," *Astrophys. J.*, vol. 779, 115, 2013.

[16] C. F. McKee, A. Parravano and D. J. Hollenbach, "Stars, gas, and dark matter in the solar neighborhood," *Astrophys. J.*, vol. 814, 13, 2015.

[17] Astropy Collaboration, "The Astropy Project: Sustaining and growing a community-oriented open-source project and the latest major release (v5.0) of the core package," *Astrophys. J.*, vol. 935, 167, 2022.

[18] R. Schönrich, J. Binney and W. Dehnen, "Local kinematics and the local standard of rest," *Mon. Not. R. Astron. Soc.*, vol. 403, pp. 1829-1833, 2010.

[19] D. B. Rubin, "Inference and missing data," *Biometrika*, vol. 63, pp. 581-592, 1976.

[20] J. Goodman and J. Weare, "Ensemble samplers with affine invariance," *Commun. Appl. Math. Comput. Sci.*, vol. 5, pp. 65-80, 2010.

[21] D. Foreman-Mackey, D. W. Hogg, D. Lang and J. Goodman, "emcee: The MCMC hammer," *Publ. Astron. Soc. Pac.*, vol. 125, pp. 306-312, 2013.

[22] C. J. F. ter Braak, "A Markov chain Monte Carlo version of the genetic algorithm differential evolution: easy Bayesian computing for real parameter spaces," *Stat. Comput.*, vol. 16, pp. 239-249, 2006.

[23] A. Gelman and D. B. Rubin, "Inference from iterative simulation using multiple sequences," *Stat. Sci.*, vol. 7, pp. 457-472, 1992.

[24] G. Schwarz, "Estimating the dimension of a model," *Ann. Stat.*, vol. 6, pp. 461-464, 1978.

[25] R. C. Lutz and T. E. Kelker, "On the use of trigonometric parallaxes for the calibration of luminosity systems: Theory," *Publ. Astron. Soc. Pac.*, vol. 85, pp. 573-578, 1973.

[26] J. Bennett and J. Bovy, "Vertical waves in the solar neighbourhood in Gaia DR2," *Mon. Not. R. Astron. Soc.*, vol. 482, pp. 1417-1425, 2019.

[27] G. Strömberg, "An analysis of the velocities of Cepheids and long period variables," *Astrophys. J.*, vol. 104, pp. 12-40, 1946.

[28] W. Dehnen and J. J. Binney, "Local stellar kinematics from Hipparcos data," *Mon. Not. R. Astron. Soc.*, vol. 298, pp. 387-394, 1998.

[29] W. G. L. Pöppel, "The Gould Belt system and the local interstellar medium," *Fundam. Cosmic Phys.*, vol. 18, pp. 1-271, 1997.

Copyright & License:
© Authors retain the copyright of this article. This work is published under the Creative Commons Attribution 4.0 International License (CC BY 4.0), permitting unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.