An overlooked source of uncertainty in the mass of the Milky Way

Kyle A. Oman [1,2,3]★ and Alexander H. Riley [1,3]

[1] *Institute for Computational Cosmology, Durham University, South Road, Durham DH1 3LE, UK*
[2] *Centre for Extragalactic Astronomy, Durham University, South Road, Durham DH1 3LE, UK*
[3] *Department of Physics, Durham University, South Road, Durham DH1 3LE, UK*

ABSTRACT

In the conventional approach to decomposing a rotation curve into a set of contributions from mass model components, the measurements of the rotation curve at different radii are taken to be independent. It is clear, however, that radial correlations are present in such data, for instance (but not only) because the orbital speed depends on the mass distribution at all (or, minimally, inner) radii. We adopt a very simple parametric form for a covariance matrix and constrain its parameters using Gaussian process regression. Applied to the rotation curve of the Milky Way, this suggests the presence of correlations between neighbouring rotation curve points with amplitudes of $< 10\,\mathrm{km\,s^{-1}}$ over length scales of 1.5–2.5 kpc regardless of the assumed dark halo component. We show that accounting for such covariance can result in a ~ 50 per cent lower total mass estimate for the Milky Way than when it is neglected, and that the uncertainty in model parameters increases such that it seems more representative of the uncertainty in the rotation curve measurement. The statistical uncertainty associated with the covariance is comparable to or exceeds the total systematic uncertainty budget. Our findings motivate including more detailed treatment of rotation curve covariance in future analyses.

Key words: methods: statistical – Galaxy: kinematics and dynamics – Galaxy: structure.

1 CORRELATIONS ACROSS RADII IN ROTATION CURVE MEASUREMENTS

Rotation curves are a widely used dynamical tracer of the mass content of and distribution within galaxies including the Milky Way (e.g. Carignan et al. 2006; de Blok et al. 2008; Kuzio de Naray, McGaugh & de Blok 2008; Swaters et al. 2009; McMillan 2011; Bovy et al. 2012; Sofue 2012; Adams et al. 2014; Pato, Iocco & Bertone 2015; Lelli, McGaugh & Schombert 2016; Eilers et al. 2019; Cautun et al. 2020; Ou et al. 2024, amongst many others). The atomic measurements comprising a rotation curve – the orbital speed at fixed radius – are usually taken to be independent from and uncorrelated with their counterparts at other radii (including in all the references above; see also Põder et al. 2023, for some discussion of concerns around such correlations). However, it is easy to see that correlations across radii must exist. One unambiguous source is the connection between gravitationally-driven kinematics and integrals of the mass distribution – for example, the mass within some central aperture contributes to the kinematics at all larger radii, introducing a correlation. Correlations may also arise due to instrumental effects such as beam smearing (Swaters et al. 2009, and references therein), modelling effects such as a geometrically thick disc being imperfectly separated into rings, or physical effects such as a spiral arm coherently perturbing the kinematics over a range in radius. The presence of correlations can further be inferred by noticing that the scatter implied by the statistical uncertainties

on rotation curve measurements often exceeds the point-to-point scatter measured (but misestimates of the uncertainties could also contribute, including incorrectly assuming that uncertainties are Gaussian distributed).

The heterogenous origins of radial correlations in rotation curve measurements makes them challenging to model explicitly. Posti (2022) proposed the pragmatic approach of assuming a parametric form for the covariance matrix:

$$\mathsf{K}_{ij} = a_k^2 \exp\left[-\frac{1}{2}\left(\frac{|R_i - R_j|}{s_k}\right)^2\right] + \sigma_i^2 \delta_{ij}. \tag{1}$$

This describes a correlation with amplitude[1] a_k between points of a given radial separation $R_i - R_j$ that decays exponentially with a length scale s_k; the uncertainties on the individual measurements are σ_i. There is freedom in the choice of 'kernel function' (the first term in equation 1) but Posti (2022) reports that reasonable variations in the choice do not lead to large differences in results, which is enough for our illustrative purposes in this work. Posti (2022) argued that marginalizing over the possibility of such correlations across radii, even in such a simplistic manner, leads to more realistic and less biased confidence intervals on model parameters of interest, such as those describing the dark halo component of a galaxy.

Analysis of the rotation curve of the Milky Way is distinct from other galaxies in two important ways. First, the systematic uncertainties affecting the measurement are distinct from those

---

[1] Our a_k is equal to $\sqrt{A_k}$ in the notation of Posti (2022).

relevant in other galaxies due to our unique vantage point, especially for recent measurements incorporating high-precision proper motion measurements of stars. Second, we have more constraints on the visible matter content and distribution of our Galaxy than for external galaxies (e.g. Binney & Tremaine 2008, section 2.7; McMillan 2017; Sofue 2020), making mass models of the Milky Way more tightly constrained. These considerations motivate us to apply the approach of Posti (2022), who illustrated it using measurements of NGC 2403, to the Milky Way in order to assess whether accounting for correlations in the measurements make up a significant portion of the uncertainty budget in mass models of the Milky Way. We also explore whether failing to account for such correlations is likely to lead to biases in rotation curve-based measurements of the mass of the Milky Way.

2 MASS MODEL COMPONENTS AND FITTING METHODOLOGY

We use the rotation curve reported by Ou et al. (2024, see their table 1) as a representative example of recent measurements (see also Jiao et al. 2023; Wang et al. 2023; Zhou et al. 2023) incorporating data from the *Gaia* mission (Gaia Collaboration 2023). We also adopt the structural parameters for the baryonic components of the mass model of Ou et al. (2024, see their table 2). The circular velocity curve specified by our model given its parameters is:

$$v^2_{\text{model}}(R) = v^2_{\text{stars}}(R) + v^2_{\text{gas}}(R) + v^2_{\text{dark matter}}(R). \tag{2}$$

We have grouped together the contribution of the stellar bulge and disc (both held fixed in model optimization) in v_{stars}, and likewise the contributions of H I gas, H$_2$ gas, warm dust, and cold dust in v_{gas} (also all held fixed). Whereas Ou et al. (2024) derived the circular velocity of each component from the enclosed mass at each radius, for disc-like components we instead use the expression in terms of the gradient of the potential Φ in the disc mid-plane:

$$v^2 = R \frac{\text{d}\Phi}{\text{d}R}. \tag{3}$$

A derivation of the potential of a thick exponential disc can be found in Binney & Tremaine (2008, section 2.6.1c); we evaluate the relevant integrals numerically and use a cubic spline approximation to measure its radial derivative. For spherically symmetric components we use the usual expression $v_{\text{circ}} = \sqrt{GM_{\text{enclosed}}/r}$.

We use two models for $v_{\text{dark matter}}$: the Navarro, Frenk & White (1996, NFW) profile and pseudo-isothermal (Gunn & Gott 1972) profile. Both can be expressed as:

$$v^2_{\text{dark matter}} = v^2_{200\text{c}} \frac{R_{200\text{c}}}{R} \frac{f_{\text{c}}\left(\frac{cR}{R_{200\text{c}}}\right)}{f_{\text{c}}(c)}, \tag{4}$$

where $v^2_{200\text{c}} = \frac{GM_{200\text{c}}}{R_{200\text{c}}}$, $R_{200\text{c}} = (2GM_{200\text{c}}/\Delta_{\text{crit}}H_0^2)^{\frac{1}{3}}$, and $\Delta_{\text{crit}} = 200$. The mass enclosed within a sphere within which the average density is Δ_{crit} times the critical density for closure, $M_{200\text{c}}$, is a free parameter. The models differ in the definition of the second free parameter – the 'concentration paramter' c. For the NFW profile we adopt the usual definition $c_{\text{NFW}} = R_{200c}/R_{\text{s}}$, where R_{s} is the 'scale radius' in the density profile $\rho(R) \propto [R(1 + R/R_s)]^{-2}$. For the pseudo-isothermal profile we define $c_{\text{PI}} = R_{200c}/R_{\text{c}}$ where R_{c} is the 'core radius' in the density profile $\rho(R) \propto [1 + (R/R_c)^2]^{-1}$. Finally, the

function f_{c} is defined for the two models respectively as[2]:

$$f_{\text{c,NFW}}(x) = \frac{\log(1 + x) - x}{1 + x} \tag{5}$$

and

$$f_{\text{c,PI}}(x) = 1 - \frac{\arctan(x)}{x}. \tag{6}$$

We stress that the choice of these two models for $v_{\text{dark matter}}$ is not motivated by their suitability to describe the Milky Way rotation curve data. Instead, they are chosen for their similarity in terms of simplicity of mathematical form (e.g. two free parameters each, analytic expressions for all needed quantities) but stark dissimilarity in structure (central $\rho \propto r^{-1}$ and outer $\rho \propto r^{-3}$ density profile for the NFW model versus central $\rho \propto r^0$ and outer $\rho \propto r^{-2}$ density profile for the pseudo-isothermal model). We will show below that some outcomes of allowing for correlations in the rotation curve data are common to both models, suggesting that the lessons learned are quite general.

We optimize the two free parameters of the model following exactly the same methodology as Posti (2022) – in fact we use their software implementation, with some modification to include the stellar bulge component and the pseudo-isothermal model in addition to the NFW model, in both the case assuming independent measurements of the rotation curve at each radius and that where correlations across radii are modelled with Gaussian process (GP) regression.

3 IMPACT OF CORRELATION MODELLING ON INFERRED MILKY WAY MASS PROFILES

We show the mass models resulting from our modelling for the two halo models in the cases without (upper panels) and with (lower panels) the GP regression model for radial correlations in the Milky Way rotation curve in Fig. 1. Corresponding best-fitting parameter values and uncertainties are reported in Table 1, and one- and two-dimensional marginalized posterior probability distributions are shown in the Appendix. In the cases without GP regression the confidence intervals on the models' total circular velocity curves are unrealistically small, and the models' inability to adequately capture the data is reflected by a reduced chi-squared (χ^2_ν) of about 10 (NFW halo) or 5 (pseudo-isothermal halo). Qualitatively, these fits are being driven by the measurements near $10 < R/\text{kpc} < 15$ that have very small uncertainties and therefore miss the measurements near $20 < R/\text{kpc} < 25$ where the uncertainties are larger.

In the cases with GP regression the models achieve a more balanced 'compromise' fit because a moderate deviation from the points with small uncertainties ($10 < R/\text{kpc} < 15$) is mitigated by the assumption that these measurements are correlated. The correlation amplitudes and scales preferred by the models seem intuitively plausible for a galaxy like the Milky Way: $(a_k, s_k) \sim (5\,\text{km s}^{-1}, 2.4\,\text{kpc})$ for an NFW halo, or $(3\,\text{km s}^{-1}, 2.0\,\text{kpc})$ for a pseudo-isothermal halo. The confidence intervals on the total circular velocity curves and dark halo components are larger and, we feel, more representative of the statistical uncertainty in the measurements. The fits are formally much better, with $\chi^2_\nu \sim 1$ for both halo models, justifying the addition of the two additional free parameters.

In addition to wider confidence intervals, the models with GP regression are also biased towards having somewhat more centrally

---

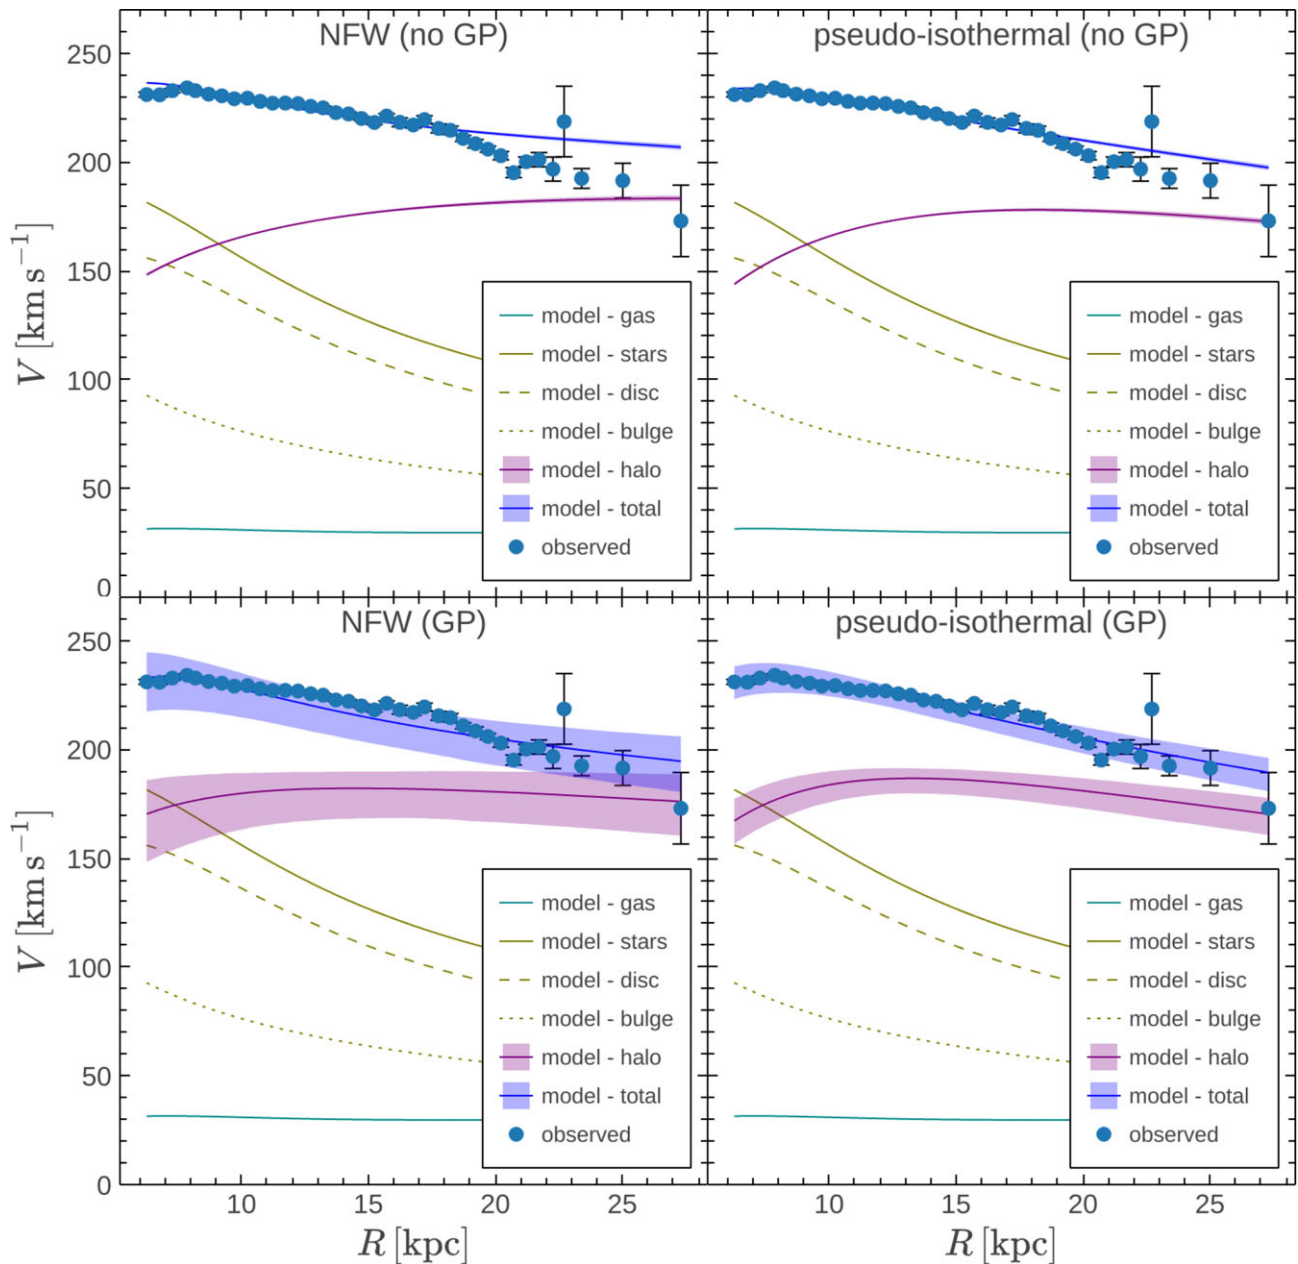[2]We use log and log$_{10}$ to denote the natural and base-10 logarithms, respectively.

Figure 1. Illustrative mass models for the Milky Way, with and without modelling radial correlations. The measurements (points with 1σ error bars) are as presented in Ou et al. (2024). In all cases the stellar disc (dashed yellow) and bulge (dotted yellow; combined stellar components shown with solid yellow), and 'gas' (green; including H I gas, H_2 gas, warm dust, and cold dust) are kept fixed to the model proposed by Ou et al. (2024); see Section 2 for details. In the left panels the halo component (solid purple) is an NFW model, while in the right panels it is a pseudo-isothermal sphere. In the upper panels no correlation ('no GP') between the observed rotation speeds is assumed, while in the lower panels the covariance matrix for the observations is estimated using GP regression as described by Posti (2022). Shaded bands mark the regions enclosing 95 per cent of model curves at each radius for the halo component and model total (in the upper panels these bands are very thin).

concentrated dark halo components, visible in Fig. 1 as an elevated circular velocity for the halo component near the centre ($R \lesssim 10$ kpc). This is compensated by lower halo masses such that the total dark matter mass within 16 and 25 kpc is the same as in the models without GP regression for the NFW and pseudo-isothermal halo cases, respectively (the dark matter mass within 20 kpc for each model is tabulated in Table 1). The marginalized posterior probability distributions for the halo mass for the four models plotted in Fig. 1 are shown in Fig. 2 (see also the Appendix). The variants with GP regression have systematically lower M_{200c}, by $\sim$0.2 dex (58 per cent)

for the NFW halo or $\sim$0.1 dex (26 per cent) for the pseudo-isothermal halo.

4 DISCUSSION

4.1 Appropriateness of the halo models

We briefly discuss whether the models shown in Fig. 1 are realistic. The NFW halo model, in particular, comes with a strong implied prior on the concentration c_{NFW} given the mass M_{200c} (e.g. Ludlow

Table 1. The first four columns show the best-fitting and marginalized 16th–84th percentile confidence intervals for the free parameters of our mass models for the two halo models – NFW and pseudo-isothermal (P–I) – and the case where correlations between rotation curve points are ignored (no GP) or accounted for through GP regression. The last column shows the inference on the mass of the dark matter halo component of the model within a 20 kpc spherical aperture. Values on a linear scale are given here for ease of reference, but we note that the probability distributions are either close to log-normal (M_{200c}, c, $M_{DM}(r < 20 \, kpc)$) or asymmetric (a_k, s_k) – see figures in the Appendix.

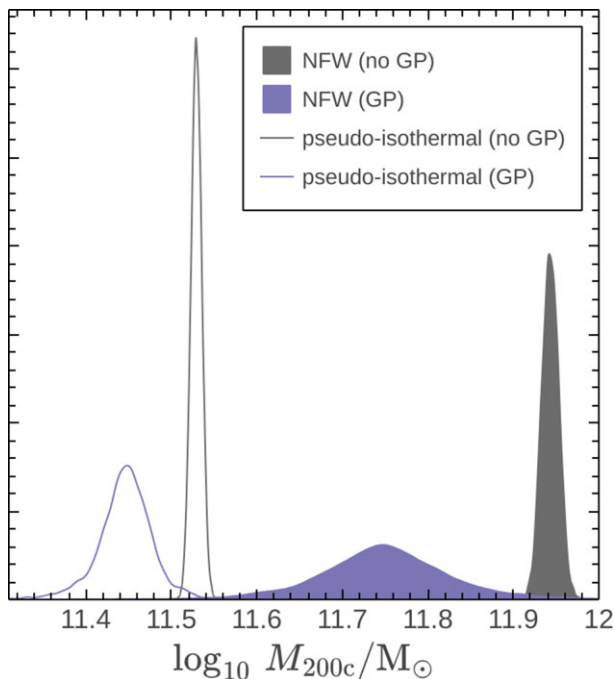| Model	M_{200c} $(10^{11} \, M_\odot)$	c	a_k $(km \, s^{-1})$	s_k (kpc)	$M_{DM}(r < 20 \, kpc)$ $(10^{10} \, M_\odot)$
NFW (no GP)	$8.7^{+0.2}_{-0.2}$	$15.5^{+0.2}_{-0.2}$	–	–	$15.3^{+0.1}_{-0.1}$
P–I (no GP)	$3.4^{+0.1}_{-0.1}$	$12.2^{+0.1}_{-0.1}$	–	–	$14.7^{+0.1}_{-0.1}$
NFW (GP)	$5.6^{+1.2}_{-1.0}$	25^{+6}_{-5}	6^{+3}_{-2}	$2.3^{+0.5}_{-0.4}$	$15.2^{+0.7}_{-0.8}$
P–I (GP)	$2.8^{+0.2}_{-0.2}$	16^{+1}_{-1}	3^{+2}_{-1}	$1.9^{+0.6}_{-0.4}$	$15.2^{+0.4}_{-0.5}$

Figure 2. Marginalized posterior probability distributions for the model halo mass. The four cases shown are for an NFW halo (filled histograms) and a pseudo-isothermal halo (open histograms) for the no-correlations ('no GP'; grey) and covariance estimated with GP regression (purple) models. For both halo models, including an estimate of the covariance results in a significant bias to lower halo mass; the posterior probability distribution also becomes wider. One- and two-dimensional marginalized posterior probability distributions for all model parameters ('corner plots') can be found in Figs A1 & A2.

et al. 2014). We have chosen not to impose this prior in our modelling, primarily to enable a fair comparison with the pseudo-isothermal halo model where no similarly strong prior exists (but see Kormendy & Freeman 2004; de Blok et al. 2008). Given the existence of a mass–concentration relation for the NFW model, it makes sense to ask whether our models are consistent with it. We have deliberately not shown the relation in the c versus M_{200c} panel of Fig. A1 – in fact it lies largely outside of the axes. Both of our models with an NFW halo (with and without GP regression) prefer a region of the parameter space many standard deviations above the locus of the mass-concentration relation. This can reasonably be attributed

to neglecting any response of the halo to the assembly of the Galaxy (see e.g. Cautun et al. 2020, and references therein), possible mismodelling of its baryonic components, or, likely, a combination of these.

We stress, however, that our objective in this work is not to create a realistic mass model for the Milky Way, but to highlight the kinds of systematic biases that can arise when correlations between points in the rotation curve measured at different radii are ignored. With this end in mind, the choice of halo models and how realistic they are (within reason) is irrelevant: our modelling clearly shows that biases of the same sign and similar amplitude arise in both models that we explore, despite their dissimilarity. We therefore expect that more realistic models, which likely have a central density profile somewhere between a steep NFW cusp and the flat core of the pseudo-isothermal model, suffer from the same sorts of systematic biases (e.g. in halo mass; Fig. 2) as our models.

4.2 Importance of statistical uncertainty in the uncertainty budget

Another useful question to consider is whether the statistical uncertainty associated with correlations between circular velocity curve measurements at different radii is comparable to other leading sources of systematic uncertainty in the determination of the circular velocity curve. Ou et al. (2024) provide estimates for the systematic uncertainties associated with effects such as varying the assumed density profile of the kinematic tracer population, the choice to neglect a certain higher order term in the Jeans' equations, the uncertainty in the galactocentric Solar radius, and others in their fig. 6. We reproduce the curve showing the sum in quadrature of all of the systematic uncertainties considered in that work as a function of radius with the black line in Fig. 3. The systematic uncertainty budget is about 2–3 per cent, with a gradual increase between about 5 and 20 kpc, and a sharp increase to >10 per cent at larger radii.

On the same figure we show the statistical uncertainty in the total model circular velocity curve as a function of radius for our models with GP regression. Comparing to the models without GP regression (e.g. shaded bands in Fig. 1), it is clear that the statistical uncertainty is severely underestimated if correlations between rotation curve measurements at different radii are neglected. When correlations are accounted for, the statistical uncertainty is about 2 to 5 per cent (purple solid and dashed lines for the NFW and pseudo-isothermal halo models, respectively). Whereas in the case without GP regression we would conclude that the total uncertainty budget for the model circular velocity curve is strongly systematics dominated, in the case
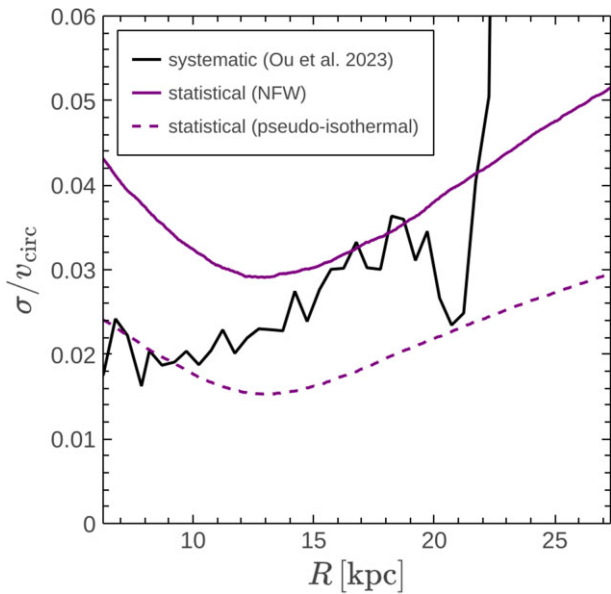
Figure 3. Comparison of the statistical uncertainty as a function of radius for our mass models including GP regression estimates of the covariance matrix. Measurements in the case with the NFW halo model (solid purple) and pseudo-isothermal halo model (dashed purple) are shown. The curves represent the width of the interval enclosing 68 per cent of model curves at each radius, normalized by the model circular velocity at that radius. These confidence intervals are comparable in width to the estimate of the total (1σ) systematic uncertainty estimate of Ou et al. (2024, black, reproduced from their fig. 6).

with GP regression it is clear that the statistical uncertainty cannot be neglected.

4.3 Possible pitfalls in interpretation

The approach of modelling correlations in rotation curve measurements across radii by adding degrees of freedom (a_k and s_k) to the model and marginalizing over them requires care in interpretation. We have experimented with applying the same methodology to external galaxies and have encountered cases where the proposed mass model is unable to describe the data (an NFW dark halo being fit to a linearly rising rotation curve, for example). In such cases the parameter search responds by moving to very large values of a_k (e.g. $> 1000\,\mathrm{km\,s^{-1}}$) and s_k (e.g. $> 100\,\mathrm{kpc}$), bounded above only by the boundary imposed on the (flat) prior. Such cases are clear model failures, but serve to illustrate that the parameters a_k and s_k will respond to essentially any discrepancy between the model and data by proposing a stronger correlation. In our exploration of the Milky Way, it is clear that something is missing from the models because $\chi_\nu^2 \gg 1$ (see Figs A1 & A2). Allowing for correlations in the measurements provides a plausible extension to the models: the goodness of fit improves and the additional parameters converge to plausible values corresponding to a fraction of the rotation speed and linear size of the disc. The question of whether this provides a more compelling explanation than other ways of accounting for the discrepancy between the data and models when correlations are neglected remains.

It is very challenging to account for all possible sources of correlation. Attempting to write down the full covariance matrix for a rotation curve measurement is currently infeasible – indeed this provides the motivation for the approach of Posti (2022). However,

taking some burden of capturing correlations off of the model by describing them in the data and its associated uncertainties would undoubtedly help to mitigate the possible pitfalls described above.

5 SUMMARY

We have shown that accounting for the possibility that rotation curve measurements of the Milky Way at different radii are statistically correlated has significant implications for inference of the Milky Way's total mass and the structure of its dark halo. In particular:

(i) accounting for such correlations can lead to a difference in the inferred mass of the Milky Way (lower by about 50 per cent);

(ii) it also results in larger uncertainty estimates for mass model parameters that we feel are more representative of the constraining power of the data;

(iii) the uncertainty budget in mass models of the Milky Way is likely not dominated by systematic uncertainties once the statistical uncertainty associated to correlations in the measurements is accounted for.

The above conclusions hold whether we assume an NFW or a pseudo-isothermal form for the dark halo component in our mass models, suggesting that they are probably generic for any broadly plausible choice of dark halo model. This strongly motivates including an allowance for correlations in the rotation curve measurements in future efforts to decompose the rotation of the Milky Way into components, but we caution that our assumed form for the covariance matrix (equation 1) is likely too simple to fully capture the correlations likely to be present in the measurements.

DATA AVAILABILITY

The rotation curve and structural parameters of the Milky Way used as inputs in our modelling are tabulated in Ou et al. (2024), tables 1 & 2, respectively. A software implementation of the fitting method that we have slightly adapted for this work is available from https://lposti.github.io/MLPages/gaussian_processes/2022/11/02/gp_rotcurves.html.

[3]bokeh.pydata.org/en/latest/
[4]pypi.org/project/contourpy/
[5]github.com/google/jax

REFERENCES

Adams J. J. et al., 2014, ApJ, 789, 63

Astropy Collaboration, 2022, ApJ, 935, 167

Bingham E. et al., 2019, Journal of Machine Learning Research, 20, 1

Binney J., Tremaine S., 2008, Galactic Dynamics, 2nd edn. Princeton University Press, Princeton, NJ

Bovy J. et al., 2012, ApJ, 759, 131

Carignan C., Chemin L., Huchtmeier W. K., Lockman F. J., 2006, ApJ, 641, L109

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

de Blok W. J. G., Walter F., Brinks E., Trachternach C., Oh S. H., Kennicutt R. C. J., 2008, AJ, 136, 2648

Eilers A.-C., Hogg D. W., Rix H.-W., Ness M. K., 2019, ApJ, 871, 120

Foreman-Mackey D., Yadav S., Theorashid, Fowlie A., Tronsgaard R., Schmerler S., Killestein T., 2022, dfm/tinygp: v0.2.3

Gaia Collaboration, 2023, A&A, 674, A1

Gunn J. E., Gott J. Richard I., 1972, ApJ, 176, 1

Harris C. R. et al., 2020, Nature, 585, 357

Jiao Y., Hammer F., Wang H., Wang J., Amram P., Chemin L., Yang Y., 2023, A&A, 678, A208

Kormendy J., Freeman K. C., 2004, in Ryder S., Pisano D., Walker M., Freeman K., eds, Proc. IAU Symp. 220, Dark Matter in Galaxies. Astron. Soc. Pac., San Francisco, p. 377

Kumar R., Carroll C., Hartikainen A., Martin O., 2019, J. Open Source Softw., 4, 1143

Kuzio de Naray R., McGaugh S. S., de Blok W. J. G., 2008, ApJ, 676, 920

Lelli F., McGaugh S. S., Schombert J. M., 2016, AJ, 152, 157

Ludlow A. D., Navarro J. F., Angulo R. E., Boylan-Kolchin M., Springel V., Frenk C., White S. D. M., 2014, MNRAS, 441, 378

McMillan P. J., 2011, MNRAS, 414, 2446

McMillan P. J., 2017, MNRAS, 465, 76

Navarro J. F., Frenk C. S., White S. D. M., 1996, ApJ, 462, 563

Ou X., Eilers A.-C., Necib L., Frebel A., 2024, MNRAS, 528, 693

Põder S., Benito M., Pata J., Kipper R., Ramler H., Hütsi G., Kolka I., Thomas G. F., 2023, A&A, 676, A134

Pato M., Iocco F., Bertone G., 2015, J. Cosmol. Astropart. Phys., 2015, 001

Phan D., Pradhan N., Jankowiak M., 2019, in NeurIPS 2019 Workshop Program Transformations for Machine Learning, preprint (arXiv:1912.11554)

Posti L., 2022, Res. Notes Am. Astron. Soc., 6, 233

Posti L., Fraternali F., Marasco A., 2019, A&A, 626, A56

Sofue Y., 2012, PASJ, 64, 75

Sofue Y., 2020, Galaxies, 8, 37

Swaters R. A., Sancisi R., van Albada T. S., van der Hulst J. M., 2009, A&A, 493, 871

Wang H.-F., Chrobáková Ž., López-Corredoira M., Sylos Labini F., 2023, ApJ, 942, 12

Zhou Y., Li X., Huang Y., Zhang H., 2023, ApJ, 946, 73

APPENDIX A: MARGINALIZED POSTERIOR PROBABILITY DISTRIBUTIONS

We show the one- and two-dimensional marginalized posterior probability distributions ('corner plots') for the parameters (M_{200c}, c, a_k, and s_k) of our mass models for the case including an NFW halo model in Fig. A1, and for the case including a pseudo-isothermal halo model in Fig. A2. We also include a panel in each figure showing the distribution of χ^2_ν values used in the likelihood function (for a definition see Posti, Fraternali & Marasco 2019, equation 2).
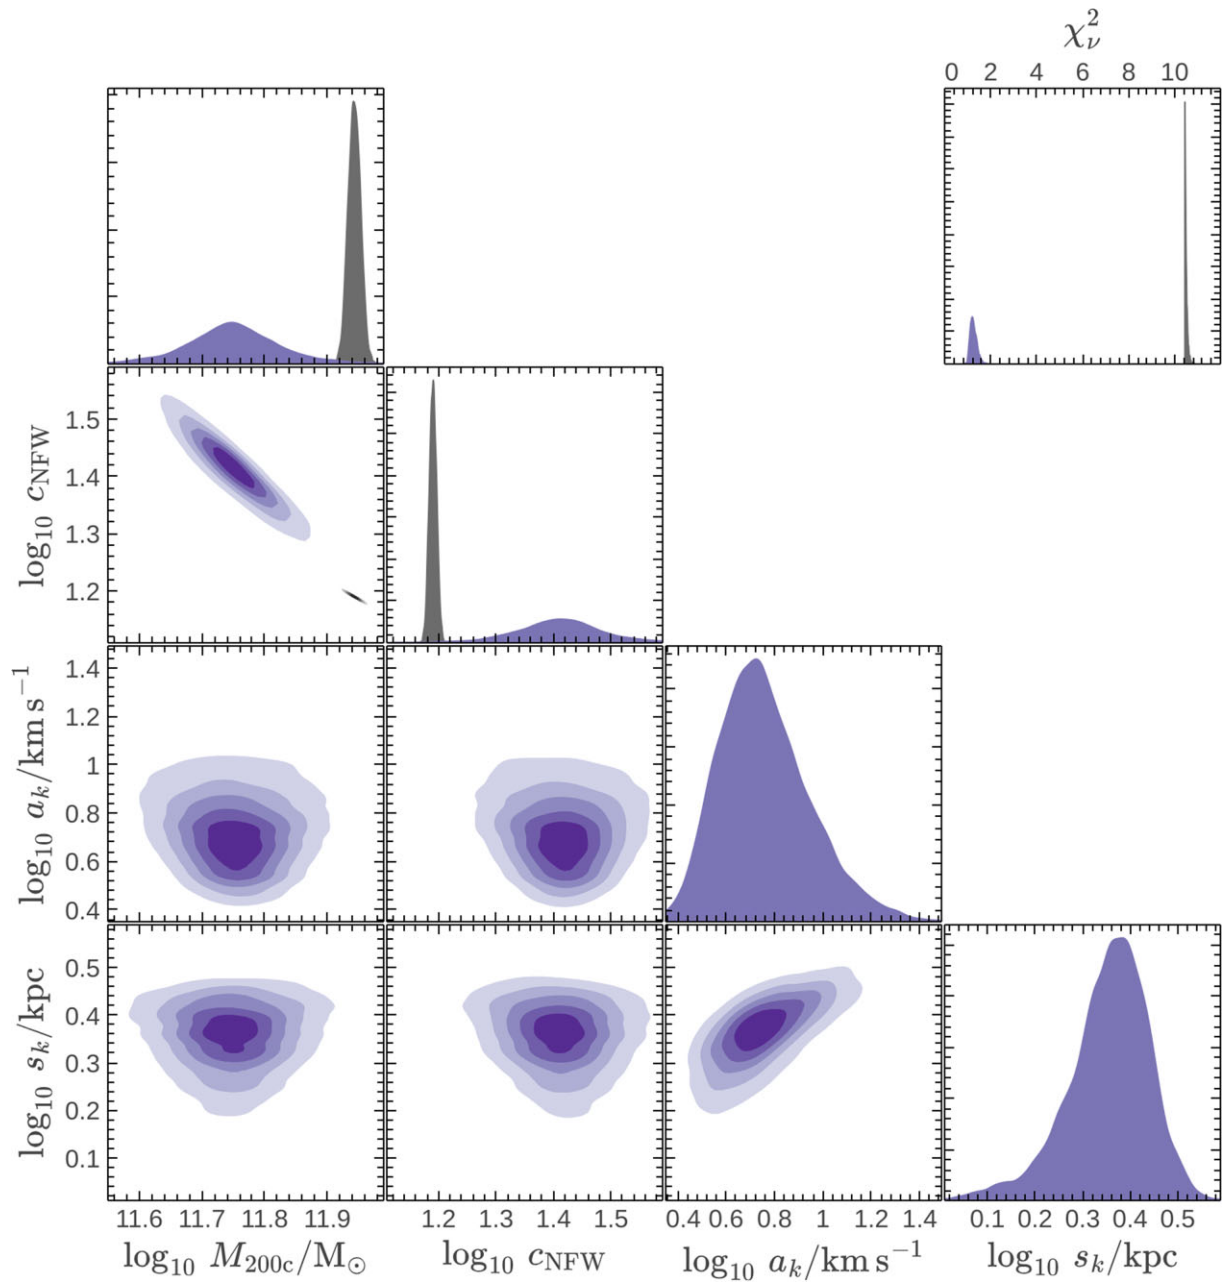
Figure A1. One- and two-dimensional marginalized posterior probability distributions for the parameters of our mass models with an NFW halo model in the 'no GP' (grey) and 'GP' (purple) cases. Model parameters in the 'no GP' case are the halo mass M_{200c} and concentration $c_{NFW} \equiv R_{200c}/R_s$. The GP case supplements these with a correlation amplitude a_k and length scale s_k. The upper right panel shows the distribution of χ^2_ν values for samples in the Markov chains.
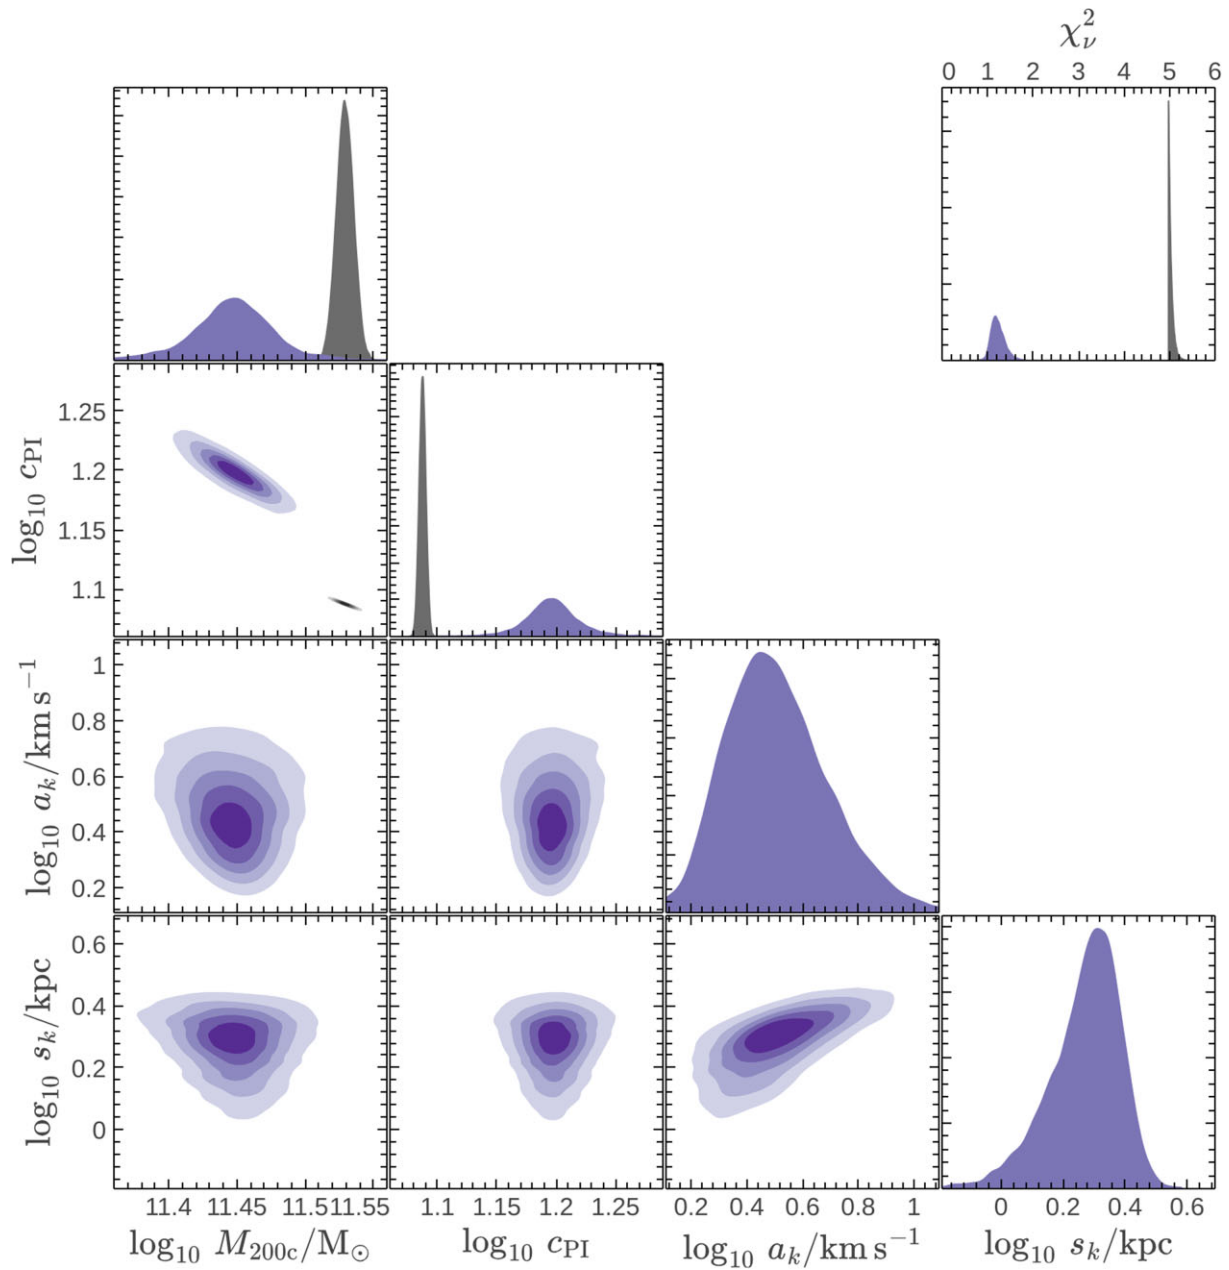
Figure A2. As Fig. A1, but for our mass models with a pseudo-isothermal halo model. For these models the 'concentration' parameter is defined $c_{\mathrm{PI}} \equiv R_{\mathrm{200c}}/R_{\mathrm{c}}$.

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