




Consequences of radially correlated rotation curves for galaxy mass models

Helena Chase ^{1,2}, Diego Dado ^{1,2}, Katherine E. Harborne ^{1,2} and Kyle A. Oman ^{1,2}★

¹*Institute for Computational Cosmology, Physics Department, Durham University, South Road, Durham DH1 3LE, UK*

²*Centre for Extragalactic Astronomy, Physics Department, Durham University, South Road, Durham DH1 3LE, UK*

Accepted 2026 May 20. Received 2026 May 8; in original form 2025 November 28

ABSTRACT

Consecutive points in rotation curve measurements are correlated with each other, but this is usually ignored when constructing galaxy mass models. We apply a recently proposed data-driven approach to include the characteristic amplitude and scale length of such correlations as ‘nuisance parameters’. We construct mass models for 134 galaxies from the SPARC (*Spitzer* Photometry and Accurate Rotation Curves) rotation curve compilation with Navarro–Frenk–White (NFW) and pseudo-isothermal sphere (pISO) models for the dark halo. Allowing for correlations in the rotation curves generally improves the goodness of fit for both halo models, often yielding a formally good fit ($\chi_r^2 \approx 1$) and model uncertainties that seem more representative of the constraining power of the data. For both halo models, the inference on the typical correlation amplitude and scale length are very similar and physically plausible, $\sim 20 \text{ km s}^{-1}$ and $\sim 5 \text{ kpc}$, respectively. The parametric form that we use to describe the correlations is intentionally simple, and our fitting approach makes the parameters describing possible correlations prone to ‘absorbing’ other systematic errors, so we regard these estimates as upper limits. Without allowing for correlations, we find a statistical preference for the pISO over the NFW model for 88/134 galaxies; this preference essentially disappears when correlations are allowed for. Accounting for correlations in rotation curves when constructing mass models fundamentally affects how they are interpreted, highlighting an important systematic uncertainty that affects evidence for cusps or cores in dark matter haloes.

Key words: methods: statistical – galaxies: kinematics and dynamics – dark matter.

1 INTRODUCTION

Rotation curves can provide strong constraints on the mass and structure of galaxies. Since early observations these have showed a discrepancy between the visible matter and the inferred dynamical mass (see G. Bertone & D. Hooper 2018 for a historical perspective). Mass modelling can reveal the distribution of the dark matter within galaxies. Several dark halo models with differing density profiles have been proposed (e.g. J. E. Gunn & J. R. Gott 1972; L. Hernquist 1990; A. Burkert 1995; J. F. Navarro, C. S. Frenk & S. D. M. White 1996; A. Di Cintio et al. 2014; J. I. Read, O. Agertz & M. L. M. Collins 2016a) with the aim of providing the best fit to the observed kinematics within prior constraints.

Several authors (e.g. W. J. G. De Blok, S. S. McGaugh & V. C. Rubin 2001; R. Kuzio de Naray et al. 2006; W. J. G. de Blok 2008; L. M. Fernández-Hernández, A. Montiel & M. A. Rodríguez-Meza 2019; P. Li et al. 2020) have argued that allowing for centrally constant dark matter density ‘cores’ (R. A. Flores & J. R. Primack 1994; B. Moore 1994, and see W. J. G. de Blok 2010; L. V. Sales, A. Wetzel & A. Fattahi 2022 for reviews) can lead to better-fitting mass models than centrally denser dark matter ‘cusps’, especially for low-mass galaxies. However, central density cusps

are a generic prediction of N-body simulations of cosmological structure formation in the lambda cold dark matter (Λ CDM) cosmogony (J. F. Navarro, C. S. Frenk & S. D. M. White 1997).

Mass model fitting is subject to many potential systematic errors arising following uncertain assumptions (e.g. G. Rhee et al. 2004; O. Valenzuela et al. 2007; J. I. Read et al. 2016b; J. C. B. Pineda et al. 2017). This work focuses on one potential source of error: In nearly all prior work, the velocity measurements making up a rotation curve are treated as statistically independent, when in reality they are correlated. Correlation arises in several ways. First and least controversially, the circular velocity is an integral quantity that depends at least on the inner matter distribution (in a spherically symmetric system), or more generally on the entire matter distribution. All points in a rotation curve are therefore at minimum correlated with all points interior to themselves, if the kinematics are gravitationally driven. Instrumental effects, such as beam smearing (e.g. R. A. Swaters et al. 2009), can also introduce local correlations. Physical effects also play a role, for example, a spiral arm can induce a locally correlated perturbation to the rotation pattern – indeed local morphological features often correspond to local features in rotation curves (R. Sancisi 2004, but see also E. Ko et al. 2025).

Given the diversity of origins of correlations in rotation curve data, it is prohibitively difficult to explicitly account for all of them in mass models. Instead, L. Posti (2022) proposed a data-

* E-mail: kyle.a.oman@durham.ac.uk

driven approach. A covariance matrix is included in the mass modelling pipeline and Gaussian process (GP) regression is used to constrain the correlation amplitudes. As a starting point, the correlations are assumed to obey a simple analytic form described by a single amplitude and length scale (per galaxy). His work presented NGC 2403 as a case study and showed that the data were consistent with the presence of correlations of 3 km s^{-1} on scales of 1 kpc. He further showed that failing to account for the possible presence of these correlations can lead to underestimated uncertainties and biased estimates for structural parameters, such as the dark halo mass. K. A. Oman & A. H. Riley (2024) applied the same methodology to the case of the Milky Way and showed that low-amplitude correlations ($\lesssim 10 \text{ km s}^{-1}$) on $\sim 2 \text{ kpc}$ scales are likely to be present and accounting for them leads to a 50 per cent lower total mass for the Galaxy.

In this work, we fit mass models to the rotation curves of 134 galaxies from the SPARC.¹ compilation (Section 2) assuming either a density cusp or core in the dark halo (Section 3). In each case, we also fit models both with and without accounting for possible correlations in the rotation curve data. We investigate the differences between these models with a particular focus on differences in goodness of fit (Section 4), and discuss changes in the maximum-likelihood values and uncertainty in the model parameters. We comment on how our findings affect the interpretation of rotation curve data and where this differs from the conclusions drawn in prior work (Section 5).

We assume a cosmology with a Hubble–Lemaître parameter $H_0 = 73 \text{ km s}^{-1} \text{ Mpc}^{-1}$ for consistency with F. Lelli, S. S. McGaugh & J. M. Schombert (2016). Virial quantities are defined by an overdensity of 200 times the critical density for closure $\rho_{\text{crit}} = 3H_0^2/8\pi G$, where G is Newton’s constant, i.e. M_{200c} is the mass enclosed within a sphere of radius R_{200c} within which the mean density is $200\rho_{\text{crit}}$.

2 ROTATION CURVE SAMPLE

Our sample of rotation curves is drawn from the SPARC (F. Lelli et al. 2016) compilation of 175 nearby galaxies with $3.6\mu\text{m}$ surface photometry and HI and/or H α rotation curves. It spans a wide range of late-type galaxy properties including morphology, luminosity, and surface brightness. The HI/H α measurements aim to trace the circular velocity curves of the galaxies (e.g. corrected for inclination), while the $3.6\mu\text{m}$ band minimizes the uncertainty in the stellar mass-to-light ratio (S. E. Meidt et al. 2014). In many galaxies the rotation curves extend to large radii and can constrain the structural parameters of the dark halo.

We adopt bulk galaxy properties from the SPARC catalogue including the distance and inclination of each galaxy. To isolate and focus our attention on the leading-order influence of allowing for radial correlations in the rotation curve measurements we hold these parameters fixed at their fiducial values (F. Lelli et al. 2016, table 1) throughout this work, but note that in full detail the mass models parameters will covary with these if more freedom is allowed. We take the v_{rot} values and uncertainties tabulated in F. Lelli et al. (2016, table 2) as the rotation curves of the galaxies and fix the shapes of the stellar disc & bulge and cold gas contributions to the mass model to v_{bulge} , v_{disc} and v_{gas} from

the same table.² We fix the normalization of the gas component to the tabulated value, but allow freedom in the stellar mass-to-light ratios as detailed below (Section 3.1).

3 MASS MODEL DEFINITIONS AND FITTING METHOD

3.1 Mass model

We closely follow the methodology of L. Posti (2022). We therefore summarize the fitting procedure emphasizing where we have modified it. The rotation curve is decomposed into components as

$$v_{\text{rot}}^2(R) = v_{\text{gas}}^2(R) + \left(\Upsilon_{\text{disc}}v_{\text{disc}}(R)\right)^2 + \left(\Upsilon_{\text{bulge}}v_{\text{bulge}}(R)\right)^2 + v_{\text{DM}}^2(R), \quad (1)$$

where $v_{\text{gas}}(R)$, $v_{\text{disc}}(R)$, $v_{\text{bulge}}(R)$, and $v_{\text{DM}}(R)$ are the velocity contributions as a function of radius, R , from the gas, stellar disc, stellar bulge,³ and dark matter halo, respectively, and Υ_{disc} and Υ_{bulge} are the stellar mass-to-light ratios of the disc and bulge. We adopt the $v_{\text{gas}}(R)$, $v_{\text{disc}}(R)$, and $v_{\text{bulge}}(R)$ profiles tabulated in the SPARC compilation.

We consider two models for the dark halo: the J. F. Navarro, C. S. Frenk & S. D. M. White (NFW; 1996) profile and the pseudo-isothermal (pISO; J. E. Gunn & J. R. Gott 1972) profile (only the former was used by L. Posti 2022). Although other models are known to be able to provide formally better statistical fits to simulated dark matter haloes and/or galaxy rotation curves (e.g. A. Burkert 1995; D. Merritt et al. 2006; J. F. Navarro et al. 2010; A. Klypin et al. 2016; J. I. Read et al. 2016a; F. Allaert, G. Gentile & M. Baes 2017; H. Katz et al. 2017; A. Lazar et al. 2020; P. Li et al. 2020), we choose these two because: (i) they are simple; (ii) they have been widely used in previous work, facilitating comparison; and most importantly (iii) they maximize the difference between the two models with very different density profile slopes that bracket the range of plausible halo models. This choice helps to emphasize differences in model fits that could be more subtle had we chosen more similar or flexible models. Our aim is not to find the best mass models but instead to explore the influence of allowing for (or neglecting) correlations in the rotation curve data.

The NFW profile is characterized by a central $\rho \propto R^{-1}$ cusp and an outer slope of $\rho \propto R^{-3}$ and has two degrees of freedom. One possible parametrization is the halo mass, M_{200c} , and the concentration, $c_{\text{NFW}} \equiv R_{200c}/R_s$, where R_s is the (unique) radius where⁴ $d \log_{10} \rho / d \log_{10} R = -2$. N-body simulations predict a tight correlation between these two parameters (e.g. A. D. Ludlow et al. 2016). In contrast, the pISO profile is characterized by a central $\rho \propto R^0$ core and an outer density slope of $\rho \propto R^{-2}$. Similarly to the NFW profile, the pISO model can also be characterized by two parameters. We choose an unconventional parametrization in analogy with the NFW parameters. The halo mass M_{200c} is defined in the same way as for the NFW profile, and we define a ‘concentration’ as $c_{\text{pISO}} = R_{200c}/R_c$, where R_c is the core radius

²In that table, these are denoted V_{obs} , V_{bul} , V_{disk} [sic], and V_{gas} , respectively.

³L. Posti (2022) did not include a bulge component since NGC 2403 does not have one.

⁴We use \log_{10} to denote the base-10 logarithm and \log for the natural logarithm.

¹Spitzer Photometry and Accurate Rotation Curves.

where the density has fallen to half of the central asymptotic value. Unlike for c_{NFW} , there is no strong prior for the relationship between M_{200c} and c_{pISO} (but see J. Kormendy & K. C. Freeman 2004).

With these definitions, the two halo models have very similar mathematical forms. v_{DM} can be expressed as

$$v_{\text{DM}}^2(R) = v_{200c}^2 \frac{R_{200c}}{R} \frac{f_c\left(\frac{cR}{R_{200c}}\right)}{f_c(c)}, \quad (2)$$

where

$$v_{200c}^2 = \frac{GM_{200c}}{R_{200c}}, \quad (3)$$

$$R_{200c} = \left(\frac{2GM_{200c}}{200H_0^2}\right)^{\frac{1}{3}}, \quad (4)$$

and c is either c_{NFW} or c_{pISO} depending on the model. The f_c function is defined for the NFW and pISO profiles as

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

and

$$f_{c,\text{pISO}}(x) = 1 - \frac{\arctan(x)}{x}, \quad (6)$$

respectively. This leaves c_{NFW} or c_{pISO} , M_{200c} , Υ_{disc} , and Υ_{bulge} as the free parameters of our mass model that does not account for possible correlations in the rotation curve data.

To introduce the possibility of radial correlations in the rotation curves into our model, we use the parametric covariance matrix of L. Posti (2022). The elements K_{ij} of the matrix are given by the kernel function

$$K_{ij} = k(R_i, R_j) + \sigma_{\text{rot},i}^2 \delta_{ij}, \quad (7)$$

where

$$k(R_i, R_j) = a_k^2 \exp\left[-\frac{1}{2} \left(\frac{|R_i - R_j|}{s_k}\right)^2\right], \quad (8)$$

and δ_{ij} is the Kronecker delta function. a_k and s_k are the characteristic amplitude⁵ (e.g. in km s^{-1}) and scale (e.g. in kpc) of the correlation. a_k is the maximum correlation amplitude, for points separated by distances much smaller than s_k . The strength of the correlation smoothly tends to zero when the separation is much greater than s_k . The correlations, encoded by $k(R_i, R_j)$, are treated as separable from the uncertainties on the individual measurements, $\sigma_{\text{rot},i}$.

This parametrization of the covariance in the data is empirical and can (and will) capture any of the sources of covariance outlined in Section 1. We caution that the true covariance is unlikely to have the exact form given by the kernel in equation (8), which can lead to biases in model fitting. We explore the influence of the choice of kernel in Appendix C1.

3.2 Fitting method

We estimate the posterior probability distributions for our model parameters using Markov Chain Monte Carlo (MCMC) sampling. For our models that neglect the possibility of correlations in the rotation curves, we sample only the parameters $\theta_v = (\log_{10} M_{200c}, \log_{10} c, \log_{10} \Upsilon_{\text{disc}}, \log_{10} \Upsilon_{\text{bulge}})$ – in

galaxies without a bulge Υ_{bulge} is omitted. We adopt log-uniform priors: $8 \leq \log_{10}(M_{200c}/M_{\odot}) \leq 14$ and $0 \leq \log_{10} c_{\text{NFW}} \leq 3$ or $0 \leq \log_{10} c_{\text{pISO}} \leq 3$. For the mass-to-light ratios, we adopt log-normal priors peaking at $\log_{10}(\Upsilon_{\text{disc}}/M_{\odot} L_{\odot}^{-1}) = -0.3$ and $\log_{10}(\Upsilon_{\text{bulge}}/M_{\odot} L_{\odot}^{-1}) = -0.15$ with a standard deviation of 0.2 dex (following F. Lelli et al. 2016). To focus on the influence of the treatment of radial correlations in rotation curves on the halo parameters, we fix the distances and inclinations of galaxies to the values tabulated in the SPARC compilation. The logarithmic likelihood for these models is given by (L. Posti, F. Fraternali & A. Marasco 2019)

$$\log_{10} \mathcal{L}(\theta_v) = -\frac{1}{2} \sum_i \frac{[v_{\text{rot},i} - v_{\text{model}}(R_i|\theta_v)]^2}{\sigma_{\text{rot},i}^2}. \quad (9)$$

To fit the characteristic length and amplitude, we repeat the MCMC sampling with all six parameters: $\theta_k = (\log_{10} M_{200c}, \log_{10} c, \log_{10} \Upsilon_{\text{disc}}, \log_{10} \Upsilon_{\text{bulge}}, \log_{10} a_k, \log_{10} s_k)$.

We keep the same priors as above for the parameters in common and impose log-uniform priors on $-2 \leq \log_{10}(a_k/\text{km s}^{-1}) \leq 2.5$ and $-2 \leq \log_{10}(s_k/\text{kpc}) \leq 3$. Our choice of priors is discussed further in Appendix C2. The log-likelihood is then

$$\log_{10} \mathcal{L}(\theta_k) = -\frac{1}{2} \mathbf{V}_{\text{res}}^T \mathbf{K}^{-1} \mathbf{V}_{\text{res}} - \frac{1}{2} \log_{10} |\mathbf{K}|, \quad (10)$$

where \mathbf{V}_{res} is a vector with components $V_{\text{res},i} = v_{\text{rot},i} - v_{\text{model}}(R_i|\theta_k)$. This is a GP regression problem.

3.3 Configuration and convergence of the MCMC sampler

We draw MCMC samples using NUMPYRO's (D. Phan, N. Pradhan & M. Jankowiak 2019) NUTS⁶ sampler with `dense_mass = True` and an acceptance probability of 0.9; other configurable parameters are set to their default values. All MCMC sampling chains (3 in parallel for each model) were run for 4000 steps with the first 1000 discarded as burn-in. We visually inspected the MCMC chains to confirm that they converge to the likelihood peak before the end of the burn-in phase. We also tested running longer chains with longer burn-in phases, which led to no significant differences. We checked that our results are insensitive to reasonable variations in the number of radial sampling points used in constructing the GP model (1000 by default). For all galaxies in our sample, the sampling produces an average $\hat{R} = 1.03$, with 92.5 per cent satisfying $\hat{R} \leq 1.05$ and 66.9 per cent achieving $\hat{R} = 1.00$. The average effective sample size across all parameters is $N_{\text{eff}} = 1831.66$.

3.4 Example: DDO 154

Fig. 1 shows the result of applying this fitting process to the rotation curve of DDO 154 – one of the earliest instances of a ‘cored’ dwarf galaxy (R. A. Flores & J. R. Primack 1994; B. Moore 1994) – as an illustrative example, for both the NFW and pISO halo model, with (GP) and without (nGP) the GP regression that allows for correlations in the rotation curve. In the nGP case, the likely parameter values are strongly constrained by the few rotation curve points with the smallest uncertainties, resulting in an unrealistically small uncertainty on the best-fitting rotation curve

⁵The A_k of L. Posti (2022) is equal to our a_k^2 .

⁶‘No U-turn sampler’, an adaptive variant of the Hamiltonian Monte Carlo sampler.




Figure 1. Mass models of DDO 154 for four cases. *Upper left:* NFW halo without GP model for correlations in rotation curve. *Upper right:* pISO halo without GP model. *Lower left:* NFW halo with GP model. *Lower right:* pISO halo with GP model. In all panels, rotation curve measurements from the SPARC data base are shown as points with error bars. The mass models show the gas (green line), stars (yellow line with shaded band), and halo (pink line with shaded band) components, along with the total model rotation curve (blue line with shaded band). The shaded bands show the 16th–84th percentile range of MCMC model evaluations at each radius. For both models, our view is that the GP case provides a more realistic description of models compatible with the data. Similar figures showing mass models for all 134 galaxies in our sample are available as Supplementary Material.

and on the model parameters. Loosely speaking, the GP model responds by assuming a correlation between points such that a modest difference between the model and a few adjacent points with small uncertainties is penalized once instead of once per data point. This leads to a larger likely volume in the parameter

space and correspondingly larger uncertainties on the best-fitting rotation curves.

The one- and two-dimensional marginalized posterior probability distributions for the model parameters for DDO 154 for both (GP and nGP) cases can be found in Appendix B. The median val-

Table 1. Parameter values for DDO 154 for the four mass models shown in Fig. 1. Column descriptions: (1) the halo model; (2) whether the GP model for correlations in the rotation curve data is used or not (nGP); (3) halo mass; (4) c , the halo concentration (NFW model) or ratio of R_{200c} and core radius (pISO model); (5) stellar disc-mass-to-light ratio; (6) amplitude of correlations in rotation curve (GP models only); (7) scale length for correlations in rotation curve (GP models only); (8) goodness of fit. All values are medians, with 16th–84th percentile intervals as uncertainties where applicable.

Halo model	GP?	$\log_{10}(M_{200c}/M_{\odot})$	$\log_{10} c$	$\log_{10}(\Upsilon_{\text{disc}}/M_{\odot} L_{\odot}^{-1})$	$\log_{10}(a_k/\text{km s}^{-1})$	$\log_{10}(s_k/\text{kpc})$	χ_r^2
NFW	nGP	$10.80^{+0.06}_{-0.06}$	$0.74^{+0.02}_{-0.03}$	$-0.83^{+0.13}_{-0.14}$	–	–	10.1
NFW	GP	$10.87^{+1.09}_{-0.77}$	$0.65^{+0.34}_{-0.42}$	$-0.38^{+0.18}_{-0.19}$	$1.17^{+0.50}_{-0.41}$	$0.59^{+0.26}_{-0.26}$	1.2
pISO	nGP	$10.10^{+0.04}_{-0.03}$	$0.87^{+0.01}_{-0.01}$	$-0.82^{+0.13}_{-0.14}$	–	–	6.7
pISO	GP	$10.26^{+1.82}_{-0.52}$	$0.84^{+0.18}_{-0.50}$	$-0.36^{+0.17}_{-0.20}$	$1.14^{+0.53}_{-0.41}$	$0.63^{+0.38}_{-0.27}$	1.3




Figure 2. Goodness of fit and correlation amplitudes and scale lengths for DDO 154. *Left panel:* Reduced chi-squared χ_r^2 for each combination of halo model (NFW, solid lines; pISO, dashed lines) and Gaussian process model for radial correlations in the rotation curve (GP, red; nGP, black). In the nGP case, there is a clear preference for the pISO halo model (lower typical χ_r^2), although neither model is a good fit ($\chi_r^2 \gg 1$). This preference vanishes when the GP model is used. *Middle panel:* Marginalized posterior probability distribution for the amplitude of radial correlations, a_k , in the rotation curve. Lines are as in left panel. *Right panel:* Marginalized posterior probability distribution for the scale length, s_k , of radial correlations in the rotation curve. Lines as in the left panel. The a_k and s_k distributions for the two models are nearly identical. All distributions are smoothed using Gaussian kernel density estimation (KDE).

ues and 16th–84th percentile intervals are summarized in Table 1. Although the most likely parameter values differ for parameters in common between the GP and nGP cases, they are consistent within the uncertainties (mostly due to the larger uncertainties on the GP parameter values).

An important difference between the nGP and GP models is the goodness of fit, which we quantify using the distribution of the reduced chi-squared statistic, $\chi_r^2 = \frac{1}{N_{\text{dof}}} (\mathbf{V}_{\text{res}}^T \mathbf{K}^{-1} \mathbf{V}_{\text{res}})$, which we evaluate for each MCMC sample (see Appendix A for a discussion of this choice of statistic for goodness of fit). N_{dof} is the number of degrees of freedom in the model. The left panel of Fig. 2 shows this distribution for the four models. Considering first the nGP models, the pISO halo provides a significantly better fit ($\chi_r^2 \sim 7$) than the NFW halo ($\chi_r^2 \sim 10$). This type of evidence has been used (e.g. W. J. G. de Blok et al. 2001; W. J. G. de Blok et al. 2008; H. Katz et al. 2017; R. Kuzio de Naray et al. 2006; P. Li et al. 2020; S.-H. Oh et al. 2015; M. Spano et al. 2008; see also M. Manju & S. Majumdar 2023 for a different approach to model comparison) to argue that cored halo models are statistically preferred over halo models with cusps, although in this example neither model achieves a good fit (i.e. $\chi_r^2 \sim 1$).

Turning next to the GP models, we find that both halo models result in a statistically good fit with a peak in the χ_r^2 distribution near 1. This is not surprising since these models have consider-

ably more freedom through the correlation parameters a_k and s_k . However, the way that this arises is important. For both halo models, the a_k and s_k parameters converge around maximum-likelihood values (see Appendix B), and interestingly there is no strong covariance between these and M_{200c} , c , or the mass-to-light ratios (although a_k and s_k are strongly covariant). Most interestingly, nearly exactly the same correlation strength and scale are preferred for both halo models, as shown in the centre and right panels of Fig. 2. A slightly (0.1 dex, i.e. ~ 25 per cent) weaker correlation amplitude is preferred for the pISO than for the NFW halo model, but this is much less than the width of the distribution, and only about 1.5 km s^{-1} in absolute terms. This slightly larger correlation amplitude could be responsible for making up the relative difference in goodness of fit in the nGP and GP cases between the two halo models, but we will see below that this is not typical across a larger sample of galaxies.

For the case of DDO 154, the conclusion is straightforward. If we accept that the rotation curve points could be correlated at around the level implied by the most likely a_k and s_k parameters, then there is no statistical preference for one halo model over the other. This possibility is made more compelling by the fact that both models achieve a good fit with nearly the same correlation amplitude and length scale, and our expectation that such correlations must be present (Section 1), albeit with their strength

unknown *a priori*. We next turn to exploring the implications of allowing for such correlations by applying our modelling procedure to a larger sample drawn from the SPARC compilation.

3.5 Sample selection

After obtaining best-fitting parameters for the models defined in Section 3.1, we pruned the initial SPARC sample of 175 to those useful for further discussion of the influence of radial correlations in rotation curve data. We noticed that when the shape of the circular velocity curves allowed by a halo model (NFW or pISO) is incompatible with the rotation curve to be fit, the models including correlations (GP) respond by sampling unrealistically large values of a_k and s_k . Qualitatively speaking, it is trying to achieve a formally good fit by penalizing the shape mismatch only ‘once’ (because all data are assumed to be extremely strongly correlated). In such cases a different, perhaps more flexible halo model might be needed to obtain a good fit with physically plausible values for a_k and s_k . (Or, the rotation curve may not be a good tracer of the mass distribution, for example because the galaxy is out of equilibrium.) We therefore restrict our sample of galaxies to those where a plausible fit of at least one of the NFW and pISO halo models is possible.

The goodness of fit (parametrized by χ_r^2) for our models without allowing for radial correlations (nGP) are in principle a good metric to identify galaxies where a good fit cannot be achieved by either/both models. However, many galaxies whose rotation curve shapes could be described well by one or both halo models have small uncertainties in v_{rot} that result in much larger χ_r^2 values than for galaxies with similar rotation curve shapes but larger uncertainties. DDO 154 is one example of such a galaxy. Therefore, for the purpose of selecting our galaxy sample only, we adopt a constant absolute uncertainty of 5 per cent of the galaxy’s rotation curve maximum for each v_{rot} measurement. These are typically somewhat larger than the nominal uncertainties recorded in the SPARC compilation. For each galaxy, the median value of the χ_r^2 posterior distributions of the nGP models (with these constant uncertainties) for both halo models are plotted against each other in Fig. 3. The 16th and 84th percentiles of the distributions are shown with error bars. Any galaxy with a median χ_r^2 value less than 2 – indicating that the relevant halo model can provide a reasonable match for the shape of the rotation curve – for either halo model is accepted into our sample.

After this initial selection, we also remove any remaining galaxies with less than 10 velocity measurements. As the GP models have up to six fitting parameters, this would leave a maximum of three degrees of freedom in the model. We found that keeping the relative change in the number of degrees of freedom between the nGP and GP models small – by keeping the total number of degrees of freedom relatively large – simplifies interpreting the resulting best-fitting models.

This selection process removes 41 galaxies, 28 of which are removed for having < 10 velocity measurements after the initial selection. The SPARC compilation includes a ‘quality flag’ Q for each galaxy, providing a heuristic measure of the quality of the HI or H α data, and the probability that a rotation curve will be a good tracer of the circular velocity. For example, very asymmetric galaxies, galaxies with strong non-circular motions and galaxies with gas spatially offset from their stars are penalized in Q . Our cuts remove 75 per cent of the galaxies flagged $Q = 3$ by F. Lelli et al. (2016) from our sample (we do not explicitly use this quality flag in our selection). 45 per cent of $Q = 2$ galaxies




Figure 3. Comparison of the χ_r^2 values for the pISO (abscissa) and NFW (ordinate) halo models when constant uncertainties of $0.05 \max(v_{\text{rot}})$ are assumed on the rotation curve measurements. The points mark the median of the distribution across the MCMC samples with the 16th–84th percentile intervals shown by the error bars. The dashed line shows the 1:1 relation. A majority of the points lie above this line, indicating a preference for the pISO halo model. We select all galaxies whose rotation curve shapes can be well captured by at least one of the two halo models – where at least one achieves a median $\chi_r^2 < 2$, shown as the shaded region.

and only 3 per cent of $Q = 1$ galaxies are also removed, providing some reassurance that the discarded galaxies are generally less suitable for mass modelling overall. Our final sample comprises 134 galaxies.

4 RESULTS

We next proceed to apply the same analysis as for DDO 154 above (Section 3.4) to our selection from the SPARC sample. Details of all models are included as Supplementary Material. Our overall impression is that the uncertainties in the GP mass models are more representative of what can plausibly be inferred from the data (i.e. they are larger) than those in the nGP models. We begin by examining the quality of the fits, parametrized by the χ_r^2 statistic (Section 4.1), then consider the posterior probability distributions for the two additional parameters (a_k and s_k) in the GP case (Section 4.2).

4.1 Goodness of fit

Fig. 4 shows the median χ_r^2 value for the pISO halo model compared to that for the NFW halo model for both the nGP and GP cases (upper left panel). The upper right panel provides a closer look at the $\chi_r^2 < 3$ region. Since we have selected only galaxies where at least one model can provide a good description of the shape of the rotation curve (see Section 3.5), the cases with large χ_r^2 for both models necessarily come from galaxies where there is a good match to the overall shape but the uncertainties on individual points are very small. These only occur in the nGP models. The GP model responds to this situation by increasing a_k and/or s_k until a statistically good fit is achieved – χ_r^2 is never

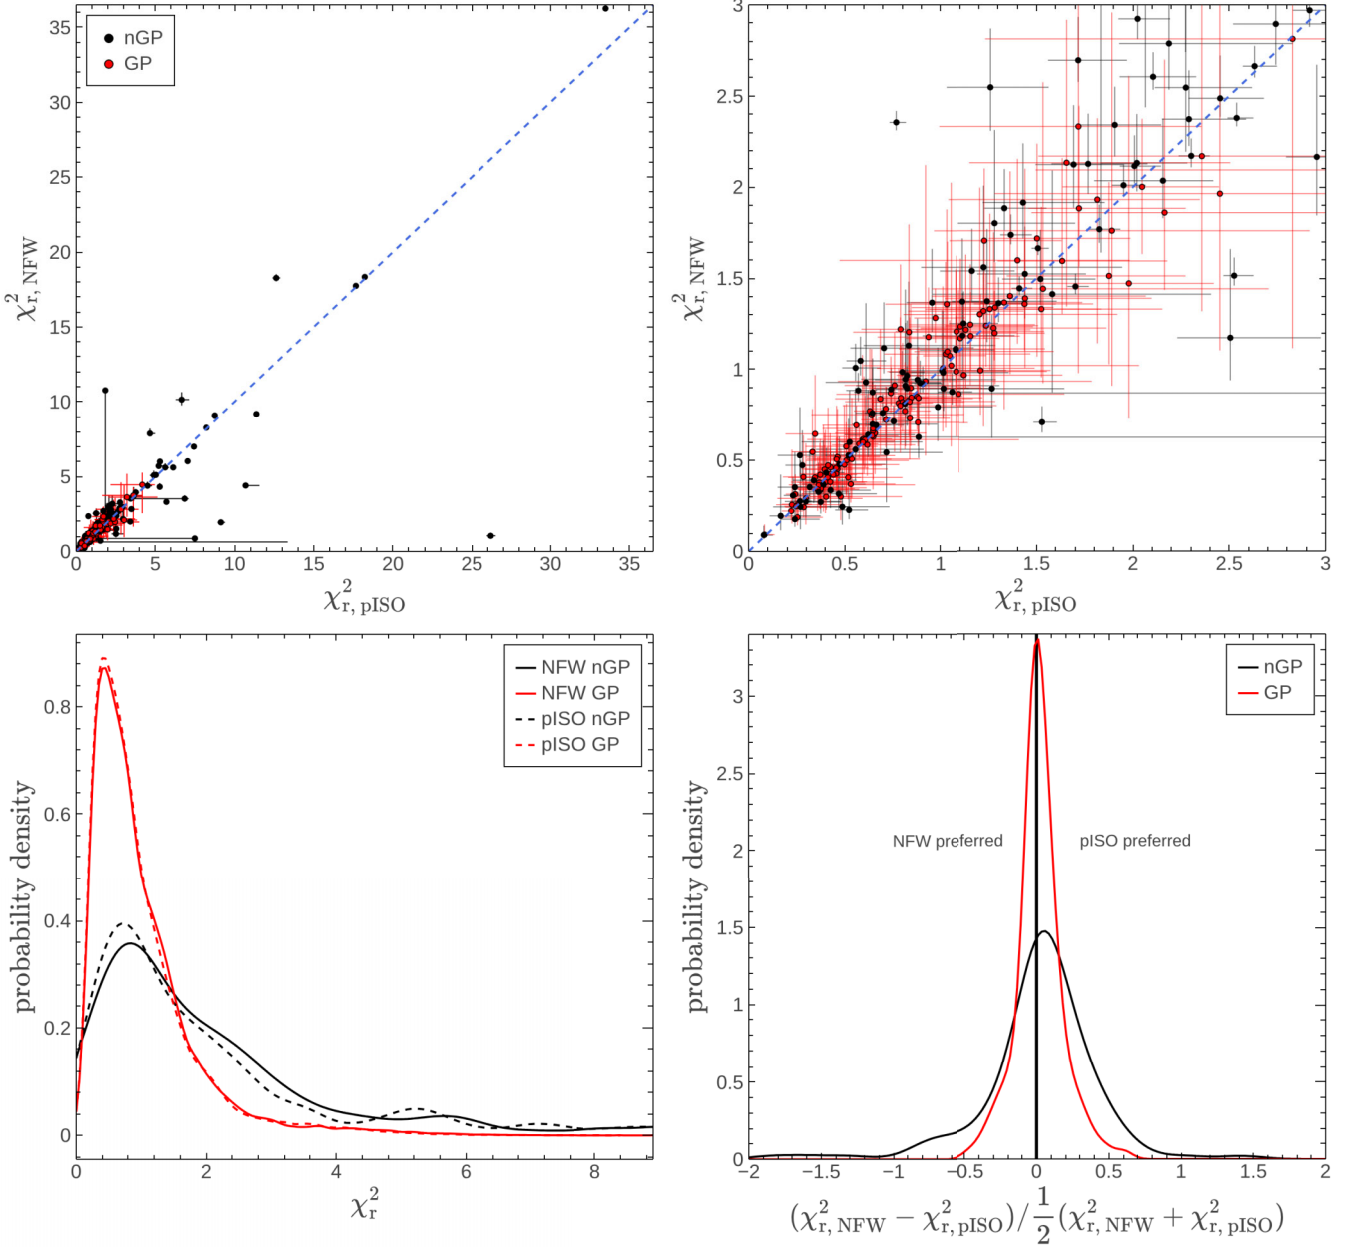


Figure 4. *Upper left:* Comparison of the χ_r^2 values for the pISO (abscissa) and NFW (ordinate) halo models with (red) and without (black) the GP model for radial correlations in the rotation curves. The points mark the median of the distribution across the MCMC samples with the 16th – 84th percentile intervals shown by the error bars. The dashed red line shows the 1:1 relation. The values are tabulated in Table D1. *Upper right:* As upper left but cropped to the $\chi_r^2 < 3$ region. In the nGP case most points lie above the 1:1 line indicating a preference for the pISO model, while in the GP case the points are evenly distributed around it, indicating no preference for either model. *Lower left:* Distribution of χ_r^2 summed across all sampled galaxies, smoothed with Gaussian KDE. The same slight preference for the pISO model in the nGP case as noted in the description of the upper left panel is visible here as a slightly sharper peak near $\chi_r^2 = 1$. In the GP case, the distributions for the two halo models are almost identical and much more sharply peaked near $\chi_r^2 = 1$ than their nGP counterparts. *Lower right:* The distribution of the fractional difference between the median χ_r^2 values for each galaxy. The distribution for the nGP case is slightly offset towards $\chi_{r,\text{pISO}}^2 < \chi_{r,\text{NFW}}^2$, again indicating a slight preference for the pISO halo model. This confirms that the preference seen in the other panels is not driven only by a very strong preference for the pISO model in a small number of galaxies. In the GP case, the distribution is centred about 0, indicating no preference for one halo model over the other.

more than ~ 2 for this selection of galaxies. We recall that our sample includes galaxies where at least one of the two models can broadly match the rotation curve shape, the other model may fail catastrophically. The most extreme example of this is UGC 02885 with $\chi_{r,\text{pISO}}^2 \sim 26$ but $\chi_{r,\text{NFW}}^2 \sim 1$ (in the nGP case).

Statistical preference for one model over the other at the galaxy population level can be qualitatively assessed by comparing the χ_r^2 distributions summed over the entire sample. This is shown in the lower left panel of Fig. 4. In the nGP case, the distribution for the pISO model is slightly more strongly peaked around

$\chi_r^2 = 1$ than that for the NFW model, indicating a very slight preference for the pISO over the NFW halo model (on average, $\chi_{r, \text{pISO}}^2 < \chi_{r, \text{NFW}}^2$). The lower right panel of the same figure shows that this is not due to a few outliers distorting the distribution but is also true on a galaxy-by-galaxy basis. The panel shows that the fractional difference

$$\frac{\chi_{r, \text{NFW}}^2 - \chi_{r, \text{pISO}}^2}{\frac{1}{2}(\chi_{r, \text{NFW}}^2 + \chi_{r, \text{pISO}}^2)} \quad (11)$$

evaluated galaxy by galaxy is slightly shifted towards a preference for the pISO halo model (88/134 galaxies; 65.7 per cent). That this preference is so weak came as a surprise in light of the claims of a clear preference for the pISO model of the NFW model in the literature (e.g. W. J. G. de Blok et al. 2008; P. Li et al. 2020) – we will return to this point in Section 5.1.

In the GP models, the slight preference vanishes: the pISO and NFW halo models fare equally well. The overall χ_r^2 distributions are essentially indistinguishable and peak near unity, and the fractional difference distribution is very nearly symmetric around 0. For the NFW halo model, 75 galaxies (56 per cent) have improved fits compared to the nGP case. Similarly, 76 galaxies (57 per cent) are improved in the GP case for the pISO halo model. The fits for the remaining galaxies do not get worse – these already had good fits in the nGP case ($\chi_r^2 \approx 1$).

A list of all galaxies in the sample and their median χ_r^2 for all four cases are given in Appendix D. We also include a complete set of figures including mass models (as in Fig. 1), corner plots (as for DDO 154 in Appendix B), and χ_r^2 distributions (as in Fig. 2) for all 134 galaxies in our sample as Supplementary Material.

4.2 Correlation amplitudes and scales

Fig. 5 shows the posterior distributions of the GP models’ characteristic correlation amplitude a_k (upper panel) and scale s_k (lower panel) for both halo models summed across all galaxies in our sample. The overall distributions are very similar. Interestingly there is a very slight preference for stronger correlations in the pISO model than in the NFW model. We might have expected that the disappearance in the GP models of the slight preference for the pISO over the NFW model found in the nGP models (Section 4.1) was due to fitting slightly stronger correlations for the NFW model to compensate for the poorer fits in the nGP models – this is not the case. The equal preference for both halo models in the GP case therefore seems to be a general consequence of modelling the correlations, rather than a systematic effect caused by the NFW halo GP model preferring stronger correlations to achieve better fits.

The GP models settle on preferred values for a_k and s_k that seem realistic (mostly $a_k \lesssim 30 \text{ km s}^{-1}$ and $s_k \lesssim 10 \text{ kpc}$), although the MCMC chains for some galaxies (usually ones with especially small uncertainties on the rotation curve) do sample much larger values. Only one galaxy has a peak (most likely) value for a_k above 56 km s^{-1} . We recall that we have selected galaxies where at least one of the halo models can reproduce the shape of the rotation curve. Attempting to fit the GP models to SPARC galaxies excluded from our sample often leads to artificially good fits with unphysically large values for a_k and s_k and very broad posterior probability distributions for the dark halo parameters.

Fig. 6 shows the correlation amplitude and scale as a function of maximum rotation velocity V_{max} , radius of the outermost measurement R_{out} , and SPARC data quality flag Q . There are




Figure 5. Posterior probability distributions for a_k (upper panel) and s_k (lower panel) summed across all 134 galaxies in our sample. Distributions are smoothed using Gaussian KDE. Both halo models show very similar distributions for correlation amplitude and scale length; the pISO model has a preference for very slightly stronger and further-reaching correlations.

no strong trends with any of these metrics, suggesting that the inference is not dominated by galaxy properties or data quality. There is weak structure in a_k with V_{max} and R_{out} . This is not unexpected: A galaxy with higher V_{max} may span a larger range in velocity across the rotation curve, providing a physical motivation for the larger correlation amplitudes inferred for some galaxies. Similarly, galaxies with larger R_{out} encapsulate a larger disc, where multiple or larger perturbations and other dynamical effects may contribute to increased inferred size of the correlation amplitudes. However, the very weak trends and large scatter in all

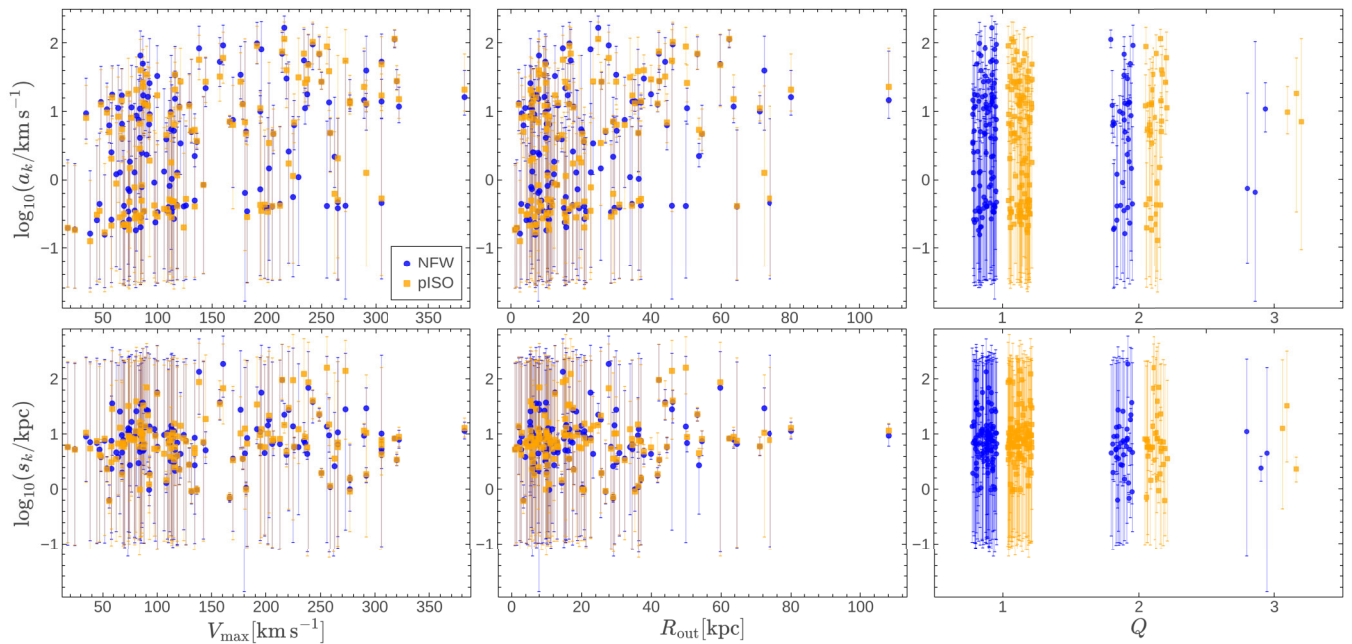


Figure 6. Correlation amplitude (a_k , top row) and scale (s_k , bottom row) dependence on sampling and resolution indicators across all 134 galaxies in our sample as a function of maximum rotation velocity (V_{\max} , left column), the radius of the outermost rotation curve measurement (R_{out} , middle column) and SPARC quality flag (Q , right column); a small amount of horizontal scatter has been added for clarity. Points mark the 50th percentile of the posterior probability distribution, with 16th–84th percentile interval shown with error bars. There are no strong trends with these sampling metrics.

panels of Fig. 6 imply that such effects are not the main drivers of the inferred values for a_k and s_k .

5 DISCUSSION

5.1 The preference for the pISO halo

Our finding that there is no strong statistical preference for the cored pISO halo model over the cuspy NFW halo model in the nGP case was unexpected (Section 4.1). P. Li et al. (2020) previously fit mass models to all SPARC galaxies and reported χ_r^2 values. This provides a helpful touchstone for comparison since that study used the same velocity measurements and baryonic mass profiles (from SPARC) and the same dark halo models (pISO and NFW, among others), therefore differing only in the model fitting methodology. The models are similarly parametrized, with their v_{200c} taking the place of our M_{200c} (we do not expect this to introduce any appreciable bias), and exact analogues of our c , Υ_{bulge} , and Υ_{disc} . In contrast to our approach, distance and inclination were also included as free parameters, with Gaussian priors centred on the values tabulated in the SPARC compilation, using the uncertainty as the standard deviation. Their priors⁷ on $\log_{10}(\Upsilon_{\text{disc}}/M_{\odot} L_{\odot}^{-1})$ and $\log_{10}(\Upsilon_{\text{bulge}}/M_{\odot} L_{\odot}^{-1})$ differ from ours: they also adopt normal distributions centred on -0.3 for the disc and -0.15 for the bulge (following S. S. McGaugh, F. Lelli & J. M. Schombert 2016; F. Lelli et al. 2017), but a narrower width of 0.1 dex (we used 0.2 dex).

⁷They also consider cases where a prior is imposed on the joint distribution of V_{200} and C_{200} (their ‘ Λ CDM’ prior), but we compare with their ‘flat’ prior case that is analogous to our choice of priors.

We show the distribution of the median χ_r^2 value of galaxies in our sample for both halo models and the GP and nGP cases in Fig. 7, and compare with the same distribution as tabulated in P. Li et al. (2020). Their fits are directly comparable with our nGP fits, with only the differences in methodology noted above. Our distribution of χ_r^2 across the galaxy sample agrees very closely with theirs for the NFW halo model, but they find systematically lower χ_r^2 values for the pISO halo model. The only possible sources for this difference are the added freedom in distance and inclination (effectively allowing independent multiplicative scaling of the radial and velocity axes of the rotation curve and its components) and the tighter priors on the mass-to-light ratio parameters. Re-running our model fits with either modified priors or extended to include distance and inclination as free parameters with similar priors as in P. Li et al. (2020) reveals that the latter is the primary driver of their preference for the pISO over the NFW model where we find none.⁸ This highlights the sensitivity of such model comparisons to what could reasonably be called minor changes to the modelling approach.

⁸Re-running our nGP model fits to be as similar as possible to the NFW-Flat and pISO-Flat models of P. Li et al. (2020) – i.e. including distance and inclination as free parameters and matching all priors – we find that we qualitatively reproduce their stronger preference for the pISO over the NFW halo model. If we use exactly the same model in the GP case, both models achieve a good fit and there is no preference for one model over the other. We treat distance and inclination as fixed parameters throughout our analysis to focus attention on the influence of the model for radial correlations in the rotation curves.




Figure 7. The distribution of the median χ_r^2 values for galaxies in our sample for the NFW (solid lines) and pISO halo model (dashed lines) when radial correlations in the rotation curves are modelled (GP; red) or not (nGP; black). This is similar to the lower left panel of Fig. 4 but plotting the distribution of median values instead of the summed distribution allows us to compare with distributions of the χ_r^2 values reported by P. Li et al. (2020) – these are shown with the yellow lines. We use only the 134 galaxies from our sample in constructing all distributions, and all distributions are Gaussian KDE smoothed. P. Li et al. (2020) found a stronger preference for the pISO over the NFW halo model than we do; see Section 5.1 for a discussion.

5.2 The halo mass–concentration relation

We recall that we did not impose a prior on the mass–concentration relation for the NFW model (Section 3.2) to enable a fairer comparison with the pISO halo model. However, we can check whether our posterior probability distributions are consistent with the expectation for Λ CDM haloes. The two-dimensional marginalized posterior probability distribution for M_{200c} and c_{NFW} summed over our entire sample is shown in Fig. 8. The distribution keeps the same overall shape in both the nGP and GP cases, but more weight shifts onto the peak on the A. A. Dutton & A. V. Macciò (2014) relation near $M_{200c} = 10^{11.5} M_\odot$ in the GP case. Allowing for possible correlations in the rotation curve data therefore seems to lead to a greater preference for the kinds of dark haloes predicted by Λ CDM N -body simulations. There are, however, still many samples well off the relation in both cases. The overall shape of the distribution in the nGP is quite similar to that found for the NFW model (flat priors) case by P. Li et al. (2020, see their fig. 3).

6 CONCLUSIONS

The points in measured rotation curves are necessarily correlated with each other. These correlations arise for reasons including (i) the definition of the circular velocity as an integral quantity; (ii) perturbative processes affecting galaxies; (iii) beam smearing; and likely other sources that are difficult to model. The amplitudes and length scales associated with such correlations are difficult to determine *a priori* (Section 1). We allow for correlations from any source but described by a parametric covariance ma-

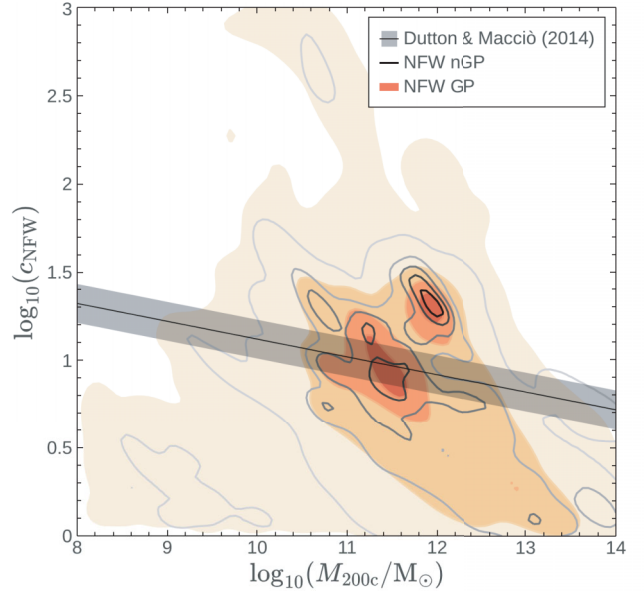


Figure 8. Two-dimensional marginalized posterior probability distributions for the c_{NFW} and M_{200c} parameters of the NFW halo model, summed across all galaxies in our sample. The nGP and GP cases are shown with open and filled linearly spaced contours, respectively. The $M_{200c} - c_{\text{NFW}}$ relation of A. A. Dutton & A. V. Macciò (2014) and its scatter are shown with the black line and shaded region. The probability density is slightly more concentrated around the $M_{200c} - c_{\text{NFW}}$ relation in the GP case, but several galaxies in the sample are still incompatible with the relation in both cases.

trix while constructing mass models to fit the rotation curves of galaxies from the SPARC compilation (Section 3), following the data-driven approach of L. Posti (2022). In our view, the resulting mass models have uncertainties more representative of the constraining power of the data than when correlations are neglected. The correlation amplitudes and scale lengths that our models prefer for SPARC galaxies are physically plausible, on the order of 20 km s^{-1} and 5 kpc , respectively (Section 4.2). While correlations must be present at some level – and our models will capture this – it is somewhat indiscriminate: Other effects, such as underestimated uncertainties on the rotation curve measurement, could be absorbed into inflated estimates of the correlation amplitude and scale length. We therefore consider our constraints on these parameters as upper limits.

Allowing for correlations erases the statistical preference for cored dark halo models (e.g. pISO) over ones with a central density cusp (e.g. NFW) that we find in our models that do not account for correlations and also reported in other work (e.g. W. J. G. de Blok et al. 2001; R. Kuzio de Naray et al. 2006; M. Spano et al. 2008; W. J. G. de Blok et al. 2008; S.-H. Oh et al. 2015; H. Katz et al. 2017; P. Li et al. 2020). This also results a statistically good fit ($\chi_r^2 \approx 1$) for both models in many galaxies where neither model achieves a good fit ($\chi_r^2 \gg 1$) when correlations are ignored (Section 4.1). This would be unsurprising if stronger correlations were needed to achieve an acceptable fit for the previously less favoured NFW model than for the pISO model, but this does not seem to be the case – if anything, the pISO model prefers (slightly) stronger correlations.

The weakened preference for one model over the other when correlations are allowed for implies weakened constraints on the central density profiles of haloes where the NFW and pISO mod-

els differ most. In more flexible models (J. Einasto 1965; A. Di Cintio et al. 2014; A. Dekel et al. 2017), the constraint on the parameter(s) controlling the central density slope will be weaker when correlations are allowed for, much like the constraints on the mass and concentration parameters explored in this work.

Comparison with the mass models of P. Li et al. (2020) reveals that the statistical preference for one model over another (e.g. the pISO over the NFW) can be very sensitive to the choice of priors (Section 5.1). Our mass models that do not account for correlations in the rotation curves are very similar to theirs, except that we treat the distances and inclinations of galaxies as fixed to the values recorded in the SPARC tables whereas they allow some freedom within the recorded uncertainties. This apparently small difference is enough to change the preference for the pISO model over the NFW from very slight (in our models) to much stronger (in theirs). Allowing for radial correlations in the rotation curves results in no preference for one model over the other for any choice of these priors that we explored. This implies that careful handling of correlations in rotation curve data are essential to correctly interpret them, including when evaluating the evidence for dark matter cusps or cores.

Allowing for correlations also seems to bring the haloes inferred for many of the galaxies that we modelled closer to the locus in halo mass and concentration expected from cosmological N-body simulations in the Λ CDM cosmogony (Section 5.2). Many galaxies in our sample still fall well off this locus however, and we also omit 41 galaxies from the SPARC compilation that neither the NFW nor the pISO model can fit well even when correlations in the rotation curves are allowed for (Section 3.5). Other interpretations must be sought for these galaxies, and those may feed back into the interpretation of correlations in rotation curves of all galaxies.

Finally, we remark that the form assumed for the correlations in the covariance matrix used when constructing the mass models is amongst the simplest possible – described by a single constant amplitude and scale length for each galaxy, with a squared-exponential kernel (see also S. Aigrain & D. Foreman-Mackey 2023 for a discussion of other kernel choices – we also find that this choice usually has little influence on the parameter inference). Any correlations between velocity measurements at different radii are certainly not constant across the discs of galaxies, so our analysis should be taken as an assessment of the overall influence that accounting for this feature of the data could have on mass models of galaxies. We conclude that they are plausibly as influential as several other sources of systematic uncertainty in mass models and bear further investigation in future work.

ACKNOWLEDGEMENTS

HC was supported through the Air Cadet Development Scheme (ref: ACDS-26-005) and a Royal Astronomical Society summer studentship. KAO, DD, and KEH acknowledge support by the Royal Society through a Dorothy Hodgkin Fellowship (DHF/R1/231105) held by KAO. KAO also acknowledges support by the European Research Council (ERC) through an Advanced Investigator grant to C. S. Frenk, DMIDAS (GA 786910), and the UK Science & Technology Facilities Council (STFC) through Grant No. ST/T000244/1. This research has made use of NASA's Astrophysics Data System.

For the purpose of open access, the author has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising from this submission.

Software: The following software packages were used in this work: ARVIZ (R. Kumar et al. 2019), ASTROPY (Astropy Collaboration 2022), BOKEH,⁹ CONTOURPY,¹⁰ CORNER (D. Foreman-Mackey 2016), JAX,¹¹ NUMPY (C. R. Harris et al. 2020), NUMPYRO (D. Phan et al. 2019), SCIPY,¹² and TINGP (D. Foreman-Mackey et al. 2024).

DATA AVAILABILITY

The SPARC data used in this work (F. Lelli et al. 2016) are publicly available from the Strasbourg astronomical Data Center¹³ (CDS) with DOI 10.26093/cds/vizier.51520157. The data from P. Li et al. (2020) are available from the CDS with DOI 10.26093/cds/vizier.22470031. The fitting routines of L. Posti (2022) are available at https://lposti.github.io/MLPages/gaussian_processes/2022/11/02/gp_rotcurves.html; our modified version of these is included with this article as Supplementary Material.

REFERENCES

- Aigrain S., Foreman-Mackey D., 2023, *ARA&A*, 61, 329
 Allaert F., Gentile G., Baes M., 2017, *A&A*, 605, A55
 Astropy Collaboration 2022, *ApJ*, 935, 167
 Bertone G., Hooper D., 2018, *Rev. Mod. Phys.*, 90, 045002
 Burkert A., 1995, *ApJ*, 447, L25
 de Blok W. J. G., 2010, *Adv. Astron.*, 2010, 789293
 de Blok W. J. G., McGaugh S. S., Rubin V. C., 2001, *AJ*, 122, 2396
 de Blok W. J. G., Walter F., Brinks E., Trachternach C., Oh S. H., Kennicutt R. C., Jr, 2008, *AJ*, 136, 2648
 Dekel A., Ishai G., Dutton A. A., Maccio A. V., 2017, *MNRAS*, 468, 1005
 Di Cintio A., Brook C. B., Maccio A. V., Stinson G. S., Knebe A., Dutton A. A., Wadsley J., 2014, *MNRAS*, 437, 415
 Dutton A. A., Maccio A. V., 2014, *MNRAS*, 441, 3359
 Einasto J., 1965, *Tr. Inst. Astroz. Alma-Ata*, 5, 87
 Fernández-Hernández L. M., Montiel A., Rodríguez-Meza M. A., 2019, *MNRAS*, 488, 5127
 Flores R. A., Primack J. R., 1994, *ApJ*, 427, L1
 Foreman-Mackey D., 2016, *J. Open Source Softw.*, 1, 24
 Foreman-Mackey D. et al., 2024, *dfm/tinygp: The tiniest of Gaussian Process libraries*, Zenodo
 Gunn J. E., Gott J. R., III, 1972, *ApJ*, 176, 1
 Harris C. R. et al., 2020, *Nature*, 585, 357
 Hernquist L., 1990, *ApJ*, 356, 359
 Katz H., Lelli F., McGaugh S. S., Di Cintio A., Brook C. B., Schombert J. M., 2017, *MNRAS*, 466, 1648
 Klypin A., Yepes G., Gottlöber S., Prada F., Heß S., 2016, *MNRAS*, 457, 4340
 Ko E., Yasin T., Desmond H., Stiskalek R., Jarvis M. J., 2025, *MNRAS*, 544, 4288
 Kormendy J., Freeman K. C., 2004, in Ryder S., Pisano D., Walker M., Freeman K., eds, *IAU Symp. Vol. 220, Dark Matter in Galaxies*. Astronomical Society of the Pacific, 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., Bosma A., 2006, *ApJS*, 165, 461
 Lazar A. et al., 2020, *MNRAS*, 497, 2393
 Lelli F., McGaugh S. S., Schombert J. M., 2016, *AJ*, 152, 157

⁹docs.bokeh.org/en/latest/.

¹⁰github.com/contourpy/.

¹¹github.com/jax-ml/jax/.

¹²docs.scipy.org/doc/.

¹³<https://cds.unistra.fr>.

- Lelli F., McGaugh S. S., Schombert J. M., Pawlowski M. S., 2017, *ApJ*, 836, 152
- Li P., Lelli F., McGaugh S., Schombert J., 2020, *ApJS*, 247, 31
- Ludlow A. D., Bose S., Angulo R. E., Wang L., Hellwing W. A., Navarro J. F., Cole S., Frenk C. S., 2016, *MNRAS*, 460, 1214
- Manju M., Majumdar S., 2023, preprint (arXiv:2310.20272)
- McGaugh S. S., Lelli F., Schombert J. M., 2016, *Phys. Rev. Lett.*, 117, 201101
- Meidt S. E. et al., 2014, *ApJ*, 788, 144
- Merritt D., Graham A. W., Moore B., Diemand J., Terzić B., 2006, *AJ*, 132, 2685
- Moore B., 1994, *Nature*, 370, 629
- Navarro J. F., Frenk C. S., White S. D. M., 1996, *ApJ*, 462, 563
- Navarro J. F., Frenk C. S., White S. D. M., 1997, *ApJ*, 490, 493
- Navarro J. F. et al., 2010, *MNRAS*, 402, 21
- Oh S.-H. et al., 2015, *AJ*, 149, 180
- Oman K. A., Riley A. H., 2024, *MNRAS*, 532, L48
- Phan D., Pradhan N., Jankowiak M., 2019, preprint (arXiv:1912.11554)
- Pineda J. C. B., Hayward C. C., Springel V., Mendes de Oliveira C., 2017, *MNRAS*, 466, 63
- Posti L., 2022, *Res. Notes Am. Astron. Soc.*, 6, 233
- Posti L., Fraternali F., Marasco A., 2019, *A&A*, 626, A56
- Read J. I., Agertz O., Collins M. L. M., 2016a, *MNRAS*, 459, 2573
- Read J. I., Iorio G., Agertz O., Fraternali F., 2016b, *MNRAS*, 462, 3628
- Rhee G., Valenzuela O., Klypin A., Holtzman J., Moorthy B., 2004, *ApJ*, 617, 1059
- Sales L. V., Wetzel A., Fattahi A., 2022, *Nat. Astron.*, 6, 897
- Sancisi R., 2004, in Ryder S., Pisano D., Walker M., Freeman K., eds, IAU Symp. Vol. 220, Dark Matter in Galaxies. Astronomical Society of the Pacific, San Francisco, p. 233
- Spano M., Marcelin M., Amram P., Carignan C., Epinat B., Hernandez O., 2008, *MNRAS*, 383, 297
- Swaters R. A., Sancisi R., van Albada T. S., van der Hulst J. M., 2009, *A&A*, 493, 871
- Valenzuela O., Rhee G., Klypin A., Governato F., Stinson G., Quinn T., Wadsley J., 2007, *ApJ*, 657, 773

SUPPORTING INFORMATION

Supplementary data are available at *MNRAS* online.

Appendix E fit_nfw_iso.py

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

APPENDIX A: GOODNESS-OF-FIT STATISTIC

Throughout our analysis, we use the distribution of the χ_r^2 statistic evaluated across MCMC samples as a measure of goodness of fit for our models. For the GP models, it is not obvious *a priori* that this is a good statistic because the covariance matrix K used to evaluate it depends on the model parameters a_k and s_k . We therefore investigate an alternative statistic to quantify the goodness of fit based on whitened residuals. We use the Cholesky transformation of K to obtain a whitening transformation $K^{-1/2}$, apply this to the model residuals as $K^{-1/2}V_{\text{res}}$ for each MCMC sample, and assess the resulting (flattened) distribution of whitened residuals. This should follow a Gaussian distribution with a standard deviation of 1 for a good model fit. The distributions for the GP models fit to some galaxies closely resemble such Gaussian distributions, while others differ in width or shape (e.g. skew or kurtosis), with extreme cases having multiple widely spaced peaks, reflecting the range of goodness of fit of the models for different galaxies – broadly this gives a similar impression to the distributions of the χ_r^2 statistic.

Fig. A1 shows the mean and standard deviation of the $K^{-1/2}V_{\text{res}}$ distribution for each galaxy against the median and 16th–84th




Figure A1. Comparison of goodness-of-fit statistics. For each galaxy in the full SPARC sample, a point marks the location of the median of the distribution of the χ_r^2 statistic for the NFW (GP) model (left panel) and pISO (GP) model (right panel), and the mean of the distribution of whitened residuals $K^{-1/2}V_{\text{res}}$ for the same model. Horizontal error bars mark the 16th–84th percentile interval of the distribution of the χ_r^2 statistic, while vertical error bars mark the standard deviation of $K^{-1/2}V_{\text{res}}$. For $\chi_r^2 \lesssim 1$, the standard deviation of $K^{-1/2}V_{\text{res}}$ is strongly correlated with χ_r^2 , while for larger values it approximately saturates at 1. The horizontal broken lines mark the expected mean (0) and standard deviation (± 1) of the distribution of $K^{-1/2}V_{\text{res}}$ for a good fit, while the vertical broken line marks $\chi_r^2 = 1$, also the expectation for a good fit. This provides reassurance that χ_r^2 is a useful goodness-of-fit statistic for our purposes.

percentiles of its χ_r^2 distribution. The mean of $K^{-1/2}V_{\text{res}}$ is usually close to 0, while its standard deviation is related to χ_r^2 : Smaller χ_r^2 values correspond to smaller $K^{-1/2}V_{\text{res}}$ values below $\chi_r^2 \sim 1$. Above this threshold, the standard deviation of $K^{-1/2}V_{\text{res}}$ saturates, rarely much exceeding 1. We also investigated the skew and kurtosis of the $K^{-1/2}V_{\text{res}}$ distribution as possible metrics, but the correspondence between the standard deviation and χ_r^2 is stronger. Taking this together with considering visualizations similar to Fig. 1, we conclude that χ_r^2 is a useful statistic for assessing goodness of fit for our purposes and therefore

use it as our fiducial statistic for this purpose throughout our analysis.

APPENDIX B: DDO 154: MARGINALIZED POSTERIOR PROBABILITY DISTRIBUTIONS

The one- and two-dimensional marginalized posterior probability distributions for our nGP and GP model fits to galaxy DDO 154 are shown in Figs B1 and B2 for the NFW and pISO halo models, respectively.

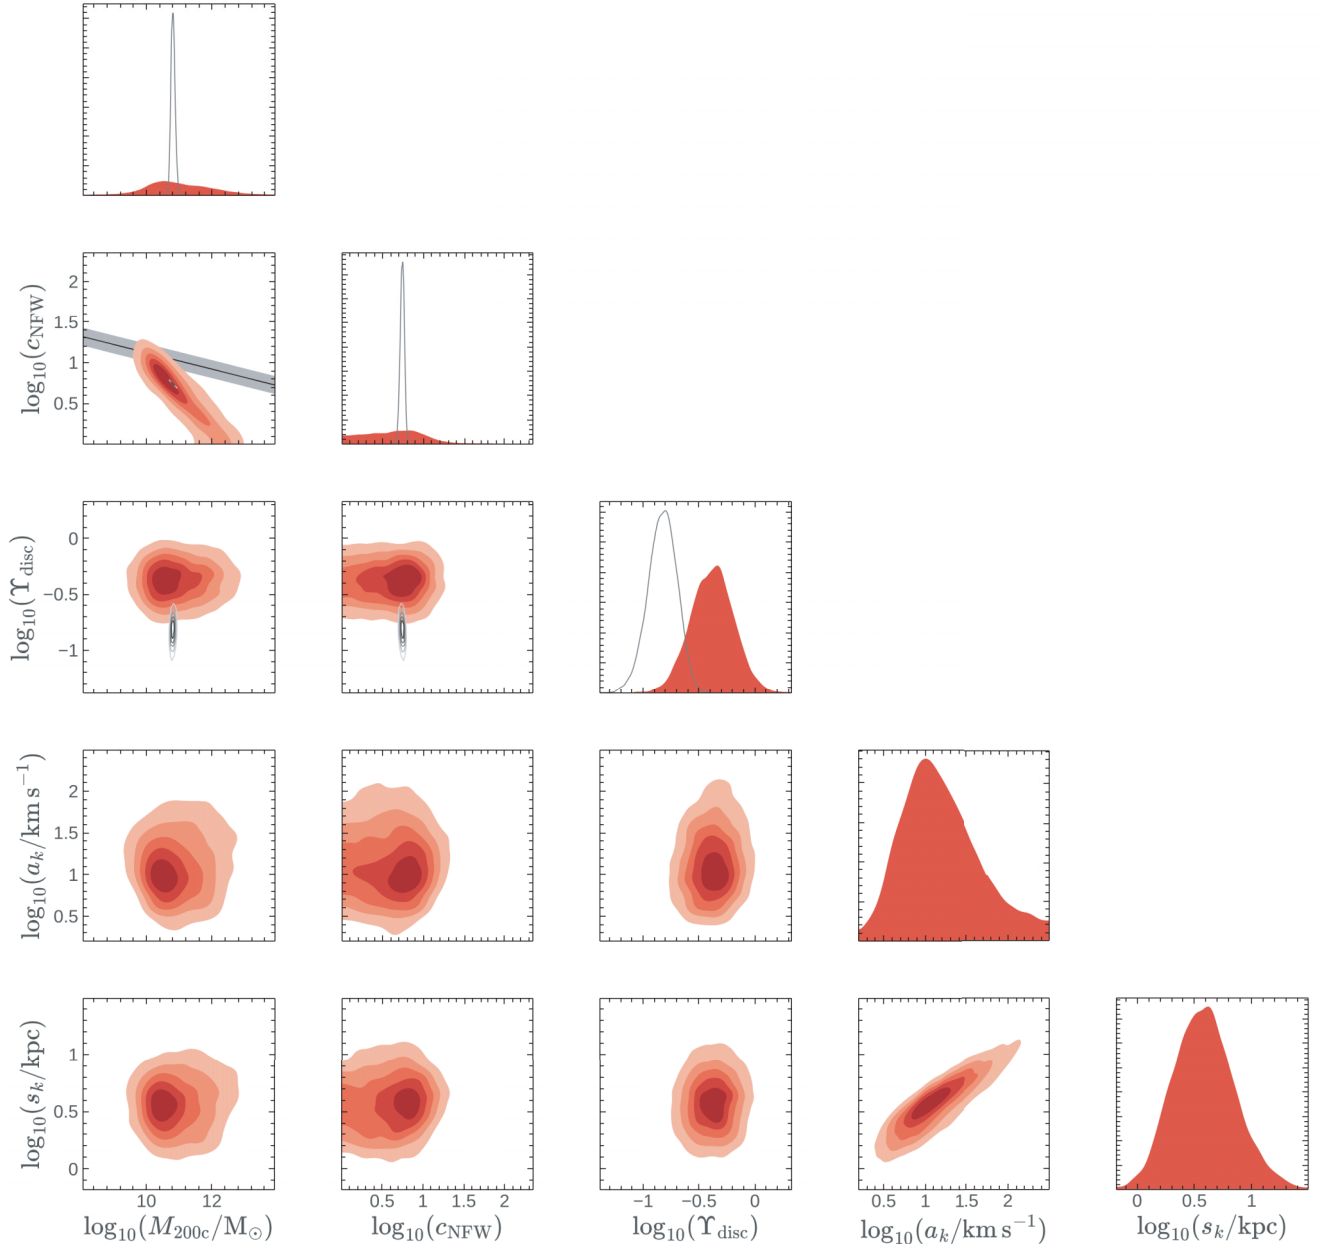


Figure B1. Corner plot showing the one- (diagonal panels) and two-dimensional (off-diagonal panels) marginalized posterior probability distributions of the parameters of our nGP (black lines and greyscale open contours) and GP (red filled contours and curves) models using the NFW dark halo model. The nGP models do not include the a_k and s_k parameters. The halo mass–concentration relation of (A. A. Dutton & A. V. Macciò 2014) and its scatter is shown in the M_{200c} – c_{NFW} panel as a black line.

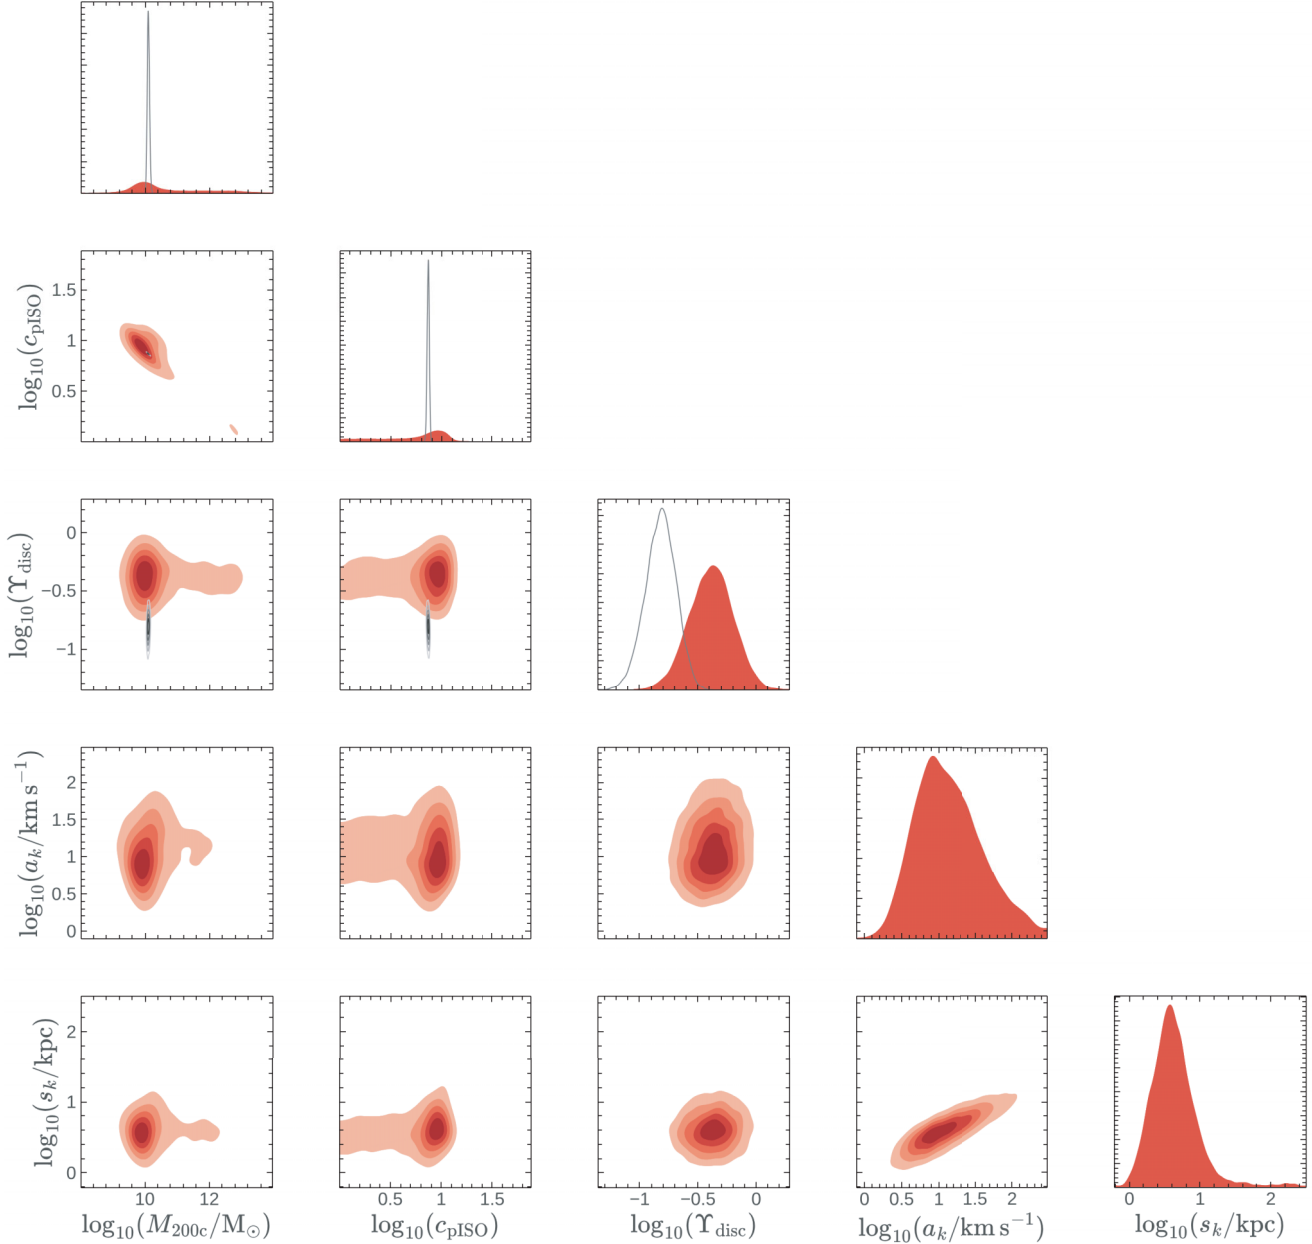


Figure B2. Same as Fig. B1 but for the pISO halo model. There is no analogue to the halo mass–concentration relation for the pISO dark halo model.

APPENDIX C: ROBUSTNESS AGAINST VARIATIONS IN METHODOLOGY

C1 Influence of the choice of kernel function

Our main analysis was carried out using a Gaussian kernel parametrization for the covariance matrix. We have repeated our model fits using a Matérn- $\frac{3}{2}$ kernel instead. Individual model fits using the two kernels sometimes differ appreciably, but in most cases the differences are small. We reach the same overall conclusions using the Matérn- $\frac{3}{2}$ kernel as with the Gaussian kernel; in particular, the lower panels of Fig. 4 are barely distinguishable. We include a version of Fig. 5 for this alternate kernel choice in Fig. C1. The constraints on the a_k and s_k parameters across the

galaxy sample are similar to with the Gaussian kernel, but the correlation amplitude is more sharply constrained to be $0.1 < a_k/\text{km s}^{-1} < 25$, whereas with the Gaussian kernel there was a tail extending to $> 100 \text{ km s}^{-1}$ – in this sense the outcome with the Matérn- $\frac{3}{2}$ kernel is perhaps even more physically plausible.

C2 Influence of priors

We experimented with alternative choices of priors including:

- (i) varying the bounds on the a_k and s_k ;
- (ii) adopting non-uniform (but still broad) priors on a_k and s_k ;




Figure C1. Same as Fig. 5, but using a Matérn- $\frac{3}{2}$ kernel to construct the covariance matrix rather than a Gaussian kernel. The parameter constraints at the population level are similar for the two kernels, except that the models with the Matérn- $\frac{3}{2}$ kernel lack the tails at $\log_{10}(a_k/\text{km s}^{-1}) \gtrsim 1.4$ and $\log_{10}(a_k/\text{km s}^{-1}) \lesssim -1$ that occurred when the Gaussian kernel was used.

(iii) adopting the posterior probability distribution of the nGP model fits as priors on the M_{200c} , c and Υ parameters for the GP model fits;

(iv) adopting the A. A. Dutton & A. V. Macciò (2014) halo mass-concentration relation as a prior on c_{NFW} (this particular choice raises the question of what, if anything, to impose for c_{pISO});

(v) combinations thereof.

We concluded that all of these choices can lead to situations where a model that would not converge to a physically plausible solution (e.g. preferring s_k much larger than the disc) with uniform priors would settle to a compromise between the prior and the area of parameter space otherwise favoured. However, this makes identifying when no physically plausible solution exists more difficult. Since this likely signals a fundamental incompatibility between the halo model(s) and observed rotation curve, we found that allowing the model to fail ‘catastrophically’ in these cases by adopting wide, uniform priors (except on the Υ parameters) led to the most readily interpreted posterior probability distributions. None of our experiments with alternative priors led to a clear preference for one halo model over the other.

APPENDIX D: GOODNESS-OF-FIT FOR ALL MODELS

Representative χ_r^2 values for nGP and GP model fits using the NFW and pISO halo models for all galaxies in our sample are tabulated in Table D1.

Table D1. Goodness of fit for our four mass models (NFW nGP; pISO nGP; NFW GP; pISO GP) for all galaxies in our subsample from SPARC. The values are the medians of the distribution of χ_r^2 values across MCMC samples.

Name	χ_r^2				Name	χ_r^2			
	NFW nGP	pISO nGP	NFW GP	pISO GP		NFW nGP	pISO nGP	NFW GP	pISO GP
D631-7	5.1	4.9	0.6	0.6	NGC 5985	2.0	3.4	0.6	0.6
DDO 064	0.7	0.7	0.8	0.8	NGC 6015	8.3	8.2	0.8	0.9
DDO 154	10.1	6.7	1.2	1.3	NGC 6195	1.9	9.1	1.4	1.4
DDO 161	0.9	0.6	0.4	0.4	NGC 6503	1.5	2.5	1.2	1.1
DDO 170	2.9	2.7	4.5	4.2	NGC 6674	2.1	1.7	1.5	1.9
ESO079-G014	4.0	3.8	1.0	1.0	NGC 6946	1.7	1.5	0.4	0.4
ESO116-G012	2.7	1.7	1.0	1.1	NGC 7331	0.9	1.1	0.4	0.4
ESO563-G021	17.7	17.7	1.2	1.2	NGC 7793	1.0	0.8	0.4	0.3
F563-1	1.3	1.1	1.3	1.2	NGC 7814	0.6	0.9	0.7	0.9
F563-V2	1.5	1.4	1.9	1.8	UGC 00128	3.3	5.7	1.2	1.1
F568-1	1.1	1.1	1.3	1.3	UGC 00191	3.7	3.6	3.6	3.2
F568-3	3.5	3.4	0.5	0.5	UGC 00731	0.5	0.3	0.6	0.3
F568-V1	0.4	0.2	0.4	0.3	UGC 01230	1.4	1.0	1.7	1.2
F571-8	3.2	2.3	1.0	1.1	UGC 02259	1.2	2.5	1.9	2.2
F574-1	1.9	1.3	0.6	0.6	UGC 02487	4.4	4.5	0.8	0.8
F579-V1	0.3	0.4	0.4	0.5	UGC 02885	1.0	26.2	1.2	1.3
F583-1	1.7	1.4	0.4	0.4	UGC 02916	10.8	1.8	0.3	0.4
F583-4	0.3	0.3	0.3	0.4	UGC 02953	5.6	6.1	0.2	0.2
IC 2574	36.3	33.5	0.6	0.6	UGC 03205	2.8	3.5	0.3	0.4
IC 4202	18.3	12.6	0.8	0.8	UGC 03546	1.0	0.8	0.8	0.8
KK98-251	2.1	2.0	0.7	0.7	UGC 03580	2.2	2.3	0.6	0.6
NGC 0024	0.4	0.4	0.5	0.4	UGC 04278	1.4	1.3	0.4	0.4
NGC 0055	2.3	1.9	0.4	0.4	UGC 04325	3.0	2.1	2.1	3.0
NGC 0100	0.9	0.9	0.4	0.4	UGC 04483	1.0	1.0	1.6	1.6
NGC 0247	1.8	1.8	0.4	0.4	UGC 04499	1.6	1.2	2.1	1.7
NGC 0289	2.0	2.0	0.7	0.6	UGC 05005	0.4	0.4	0.5	0.5
NGC 0300	0.8	0.7	0.8	0.7	UGC 05253	2.4	0.8	0.4	0.4
NGC 0801	5.6	5.6	3.7	3.6	UGC 05716	2.1	1.8	1.5	2.0
NGC 0891	2.5	1.3	1.2	1.1	UGC 05721	1.0	0.6	0.8	0.6
NGC 1003	2.4	2.5	0.5	0.5	UGC 05750	1.2	1.1	1.4	1.3
NGC 1090	2.6	2.1	1.2	0.9	UGC 05764	7.9	4.7	2.0	2.5
NGC 1705	0.2	0.5	0.3	0.5	UGC 05829	0.2	0.2	0.3	0.2
NGC 2366	3.3	2.8	0.4	0.4	UGC 05918	0.5	0.7	0.9	1.1

Table D1 – continued

Name	X_r^2				Name	X_r^2			
	NFW nGP	pISO nGP	NFW GP	pISO GP		NFW nGP	pISO nGP	NFW GP	pISO GP
NGC 2403	9.2	11.4	0.6	0.6	UGC 05986	5.7	5.2	0.5	0.5
NGC 2683	1.4	1.2	1.9	1.7	UGC 06399	0.9	0.9	1.3	1.3
NGC 2841	1.5	1.7	0.4	0.5	UGC 06446	0.3	0.4	0.4	0.4
NGC 2903	6.0	7.0	1.3	1.0	UGC 06614	0.6	0.5	0.7	0.6
NGC 2915	0.9	0.7	0.5	0.5	UGC 06667	1.8	1.3	2.3	1.7
NGC 2955	3.5	6.8	1.4	1.4	UGC 06786	0.7	1.5	0.7	0.6
NGC 2976	0.9	0.9	0.5	0.5	UGC 06787	18.3	18.2	0.6	0.6
NGC 2998	2.2	3.0	1.1	1.1	UGC 06917	1.1	0.8	1.4	1.0
NGC 3109	9.1	8.7	0.6	0.5	UGC 06930	0.6	0.6	0.8	0.8
NGC 3198	1.4	1.4	0.4	0.4	UGC 06973	1.4	1.6	2.2	2.4
NGC 3521	0.3	0.3	0.3	0.3	UGC 06983	0.8	0.6	0.9	0.7
NGC 3726	2.4	2.3	1.8	1.9	UGC 07089	0.4	0.3	0.5	0.4
NGC 3741	0.5	0.5	0.5	0.5	UGC 07125	1.1	0.7	1.2	0.8
NGC 3769	0.8	1.0	1.0	1.2	UGC 07151	2.5	2.3	1.4	1.5
NGC 3893	0.9	0.6	1.2	0.8	UGC 07323	0.9	0.8	1.1	1.0
NGC 3917	3.0	2.9	0.8	0.8	UGC 07399	1.4	1.1	1.7	1.5
NGC 3953	1.0	1.0	1.6	1.4	UGC 07524	0.9	0.6	0.3	0.4
NGC 3972	1.5	1.5	2.0	2.0	UGC 07577	0.4	0.3	0.5	0.5
NGC 3992	0.9	1.3	1.3	1.5	UGC 07603	1.9	1.4	0.9	0.9
NGC 4010	2.5	2.5	1.0	1.1	UGC 07608	0.7	0.6	1.1	1.0
NGC 4013	0.8	0.8	0.8	0.8	UGC 08286	2.0	2.2	0.8	0.7
NGC 4088	0.6	0.6	0.8	0.8	UGC 08490	0.2	0.2	0.2	0.3
NGC 4100	1.0	0.6	0.6	0.6	UGC 08550	0.9	0.8	1.2	1.1
NGC 4157	0.5	0.5	0.6	0.6	UGC 08699	0.7	0.8	0.7	0.7
NGC 4183	0.3	0.4	0.3	0.4	UGC 09037	1.5	1.2	0.7	0.8
NGC 4214	0.9	1.0	0.9	0.8	UGC 09133	7.0	7.4	0.7	0.6
NGC 4217	2.1	2.0	0.5	0.5	UGC 11455	5.1	5.0	0.8	0.9
NGC 4559	0.3	0.2	0.3	0.2	UGC 11820	0.9	0.8	1.3	1.2
NGC 5005	0.2	0.3	0.3	0.3	UGC 11914	0.9	7.5	0.2	0.2
NGC 5033	2.9	2.0	0.9	0.8	UGC 12506	0.2	0.5	0.2	0.3
NGC 5055	2.7	2.6	1.1	1.0	UGC 12632	0.5	0.3	0.5	0.3
NGC 5371	4.4	10.7	1.2	1.2	UGC 12732	0.3	0.5	0.4	0.5
NGC 5585	6.0	5.3	1.2	1.2	UGCA 442	2.8	2.2	2.8	2.8
NGC 5907	4.3	5.3	1.4	1.4	UGCA 444	0.1	0.1	0.1	0.1

This paper has been typeset from a TeX/LaTeX file prepared by the author.