

The scatter, residual correlations and curvature of the SPARC baryonic Tully–Fisher relation

Harry Desmond^{1,2★}

¹Kavli Institute for Particle Astrophysics and Cosmology, Physics Department, Stanford University, Stanford, CA 94305, USA

²SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

Accepted 2017 August 18. Received 2017 August 15; in original form 2017 June 3

ABSTRACT

In recent work, Lelli et al. argue that the tightness of the baryonic Tully–Fisher relation (BTFR) of the SPARC galaxy sample, and the weakness of the correlation of its residuals with effective radius, pose challenges to Λ cold dark matter cosmology. In this Letter, we calculate the statistical significance of these results in the framework of halo abundance matching, which imposes a canonical galaxy–halo connection. Taking full account of sample variance among SPARC-like realizations of the parent halo population, we find the scatter in the predicted BTFR to be 3.6σ too high, but the correlation of its residuals with galaxy size to be naturally weak. Further, we find abundance matching to generate BTFR curvature in 3.0σ disagreement with the data, and a fraction of galaxies with non-flat rotation curves somewhat larger than observed.

Key words: galaxies: formation – galaxies: fundamental parameters – galaxies: haloes – galaxies: kinematics and dynamics – galaxies: statistics – dark matter.

1 INTRODUCTION

Among galaxy scalings, the correlation of baryonic mass with rotation velocity [baryonic Tully–Fisher relation (BTFR)] stands out. In addition to possessing a very small intrinsic scatter, the BTFR is an almost perfect power law over six decades of mass, describes galaxies with a wide range of morphologies, and has residuals systematically uncorrelated with other galaxy variables. This makes it at once a strong test of galaxy formation theories and an important source of information on their degrees of freedom.

Recently, Lelli et al. (2016a, hereafter L16) have presented the BTFR of the SPARC sample (Lelli et al. 2016b), a compilation of 175 galaxies with high-quality H I rotation curves (RCs) and *Spitzer* imaging at $3.6\ \mu\text{m}$.¹ The authors claim two features of the SPARC BTFR to be very difficult for Λ cold dark matter (Λ CDM)-based models to account for: its small intrinsic scatter s_{BTFR} (~ 0.11 dex in baryonic mass) and the negligible correlation ρ of its residuals with galaxy size. In standard galaxy formation, s_{BTFR} should receive contributions from the halo mass–concentration and halo mass–galaxy mass relations, both of which themselves have scatter 0.1–0.2 dex, and a simple application of Kepler’s laws may be expected to yield an anticorrelation of velocity and size residuals.

Although valid, these arguments lack the statistical evidence required to claim a significant discrepancy. The aim of this work is to supply that evidence. In particular, we will calculate the expectation for s_{BTFR} and ρ in a vanilla Λ CDM model described by abundance matching (AM) by constructing mock data sets identical to SPARC in all baryonic variables and analysed in precisely the same way. We

will find two additional effects to be important, neither of which have previously been considered in detail: (1) sample variance in BTFR statistics between SPARC-like realizations of the full galaxy–halo population, and (2) the falloff with galactocentric radius of the sensitivity of the baryonic component of the RC to galaxy size. We will show that when these effects are accounted for in a complete and fully self-consistent comparison with the SPARC data, the discrepancies in s_{BTFR} and ρ are $\sim 3.6\sigma$ and 2.2σ , respectively. In addition, we investigate two further statistics that are constraining for galaxy formation models: the BTFR curvature and the fraction of RCs that are flat. The significance levels at which the SPARC values for these quantities differ from those of the model are $\sim 3.0\sigma$ and 2.2σ . We conclude that the BTFR statistics of the SPARC data pose a challenge to AM models that is moderately statistically significant.

Our work builds on a number of studies aimed at assessing the consistency between the observed BTFR and the Λ CDM prediction, which have been carried out within both AM (e.g. Trujillo-Gomez et al. 2011; Dutton 2012; Di Cintio & Lelli 2016) and hydrodynamical (e.g. Santos-Santos et al. 2016; Dutton et al. 2017) frameworks. Despite a broad consensus that the general shape of the relation is compatible with Λ CDM, the precise extent of this agreement – as well as the significance of more detailed features such as intrinsic scatter, curvature and residual correlations – remains unclear. We intend our focused work on SPARC to pave the way for more general statistical analyses in the future.

2 METHOD

Before detailing our procedure, we describe the twofold novelty of our approach. First, by using mock galaxies with baryonic properties identical to those of SPARC and sampled at the same radii

★ E-mail: harryd2@stanford.edu

¹ <http://astroweb.cwru.edu/SPARC/>

in the same ways, we eliminate systematic error in the comparison of BTFR statistics and ensure that any differences with the observed dynamics are due solely to the distribution of dark matter. Secondly, by thoroughly sampling the set of halo properties that may be associated with a given galaxy by AM, we robustly calculate the sample variance of each BTFR statistic in the model. Simple frequentist methods will then allow us to determine the significance of differences with the corresponding statistics in the data. Our method is similar to that of Desmond (2017), on which it is based. The steps are the following.

(i) From the full SPARC data set, remove starburst dwarfs and galaxies with $i < 30^\circ$ or quality flag 3. These are all the selection criteria of L16, except for a cut on RC flatness which we will come to shortly. We denote the resulting sample, containing 150 galaxies, as ‘SPARC’ hereafter.

(ii) Estimate the true stellar and gas masses of each SPARC galaxy by scattering the measured values (assuming $M_*/L = 0.5$ for the disc, $M_*/L = 0.7$ for the bulge and $M_{\text{gas}} = 1.33 M_{\text{HI}}$; Lelli, McGaugh & Schombert 2016b) by the measurement uncertainties calculated using L16, equation (5). Use the stellar mass to assign each galaxy a halo by the technique of AM (Kravtsov et al. 2004; Conroy, Wechsler & Kravtsov 2006). In particular, we will use the AM model that Lehmann et al. (2017) find to reproduce best the correlation function of SDSS, and match to haloes in the DARKSKY-400 simulation (Skillman et al. 2014), a $(400 \text{ Mpc}^{-1})^3$ box with 4096^3 particles run with the 2HOT code (Warren 2013). We identify haloes using ROCKSTAR (Behroozi, Wechsler & Wu 2013).

(iii) Use an NFW profile with the concentration and mass (subtracting the baryon fraction) of the assigned N -body halo to calculate the velocity due to the dark matter at each of the radii at which the RC of each SPARC galaxy was probed (Lelli et al. 2016b). Add in quadrature the fixed baryonic contribution (imported directly from the SPARC data) to calculate the total velocity, then scatter by the corresponding uncertainty (L16, equation 3) to model observational error.

(iv) Use the algorithm of L16 (equations 1–2) to determine whether a given model galaxy has a flat RC, and if so to calculate the corresponding velocity V_f . If the RC is not flat, discard the galaxy. Denote by N_f the total number of galaxies in the mock data set removed in this way.

(v) Use V_f and the baryonic masses (M_b) of the remaining galaxies to calculate the BTFR statistics. Begin by fitting to the BTFR and M_b – R_{eff} relation quadratic curves in log–log space, with Gaussian scatter in $\log(V_f)$ and $\log(R_{\text{eff}})$, respectively, by maximizing the corresponding likelihood model. Subtract the V_f and M_b measurement uncertainties in quadrature from the total scatter to estimate the intrinsic scatter s_{BTFR} . Take the best-fitting coefficient of the quadratic term, q , as a measure of the BTFR curvature.

(vi) Calculate the velocity and radius residuals as

$$\Delta V_f \equiv V_f - \langle V_f | M_b \rangle \quad (1)$$

and

$$\Delta R_{\text{eff}} \equiv R_{\text{eff}} - \langle R_{\text{eff}} | M_b \rangle, \quad (2)$$

where $\langle Y | M_b \rangle$ denotes the expectation for Y at fixed M_b given the fit to the full relation, and calculate ρ as the Spearman’s rank coefficient of their correlation. This completes the treatment of a single mock data set.

(vii) Repeat steps (ii)–(vi) for 2000 mock data sets, in each case randomly drawing for each SPARC galaxy a different DARKSKY-400 halo consistent with the AM model. This generates distributions

of s_{BTFR} , ρ , N_f and q that fully capture the sample variance of the model predictions.

(viii) Calculate the significance of the difference between model and data for each of $X \equiv \{s_{\text{BTFR}}, \rho, N_f, q\}$ as

$$\sigma_X \equiv (\langle X \rangle - X_d) / s_X, \quad (3)$$

where $\langle X \rangle$ is the mean of X over all mock data sets, s_X is the standard deviation of the distribution and X_d is the corresponding value in the SPARC data.

3 RESULTS

We present the predicted versus observed SPARC BTFR in Fig. 1(a), the correlation of V_f and R_{eff} residuals in Fig. 1(b), and the distribution of each statistic X in Fig. 2. Table 1 lists the mean and standard deviations of these distributions, along with the significances of offsets from the data. Here, we describe these results: Section 3.1 focuses on ρ , Section 3.2 on s_{BTFR} , Section 3.3 on N_f and Section 3.4 on curvature q .

3.1 The ΔR_{eff} – ΔV_f correlation

We begin with the correlation of R_{eff} and V_f residuals. In Fig. 1(b), we stack ΔR_{eff} and ΔV_f of all mock data sets to form a contour plot, on which we overlay the SPARC data. In Fig. 2(a), we compare the distribution of ρ in the mock data to the corresponding value in the real data, and in the first row of Table 1 we report $\langle \rho \rangle$, s_ρ , ρ_d and σ_ρ .

It is clear that the model prediction is not particularly discrepant with the data: neither shows a strong ΔR_{eff} – ΔV_f correlation. This may be understood as follows. V_f is calculated from the flat part of the RC defined by the final measured points. As this is typically several times beyond R_{eff} (~ 2 – 10 for the SPARC sample), the velocity contribution due to the galaxy is effectively that of a point mass at its centre. Variations in galaxy size at fixed M_b provide only a small perturbation to this leading order term, rendering $\langle \rho \rangle$ negligible. (Similar results obtain replacing R_{eff} by the scalelength of the stellar or gas disc.) Note that this is very different to the results of Desmond & Wechsler (2015) (their figs 6 and 7), which L16 cite as evidence for the expectation $\rho \ll 0$. This is because Desmond & Wechsler (2015) use the velocity at the radius enclosing 80 per cent of the i -band light, where the baryonic contribution to the RC is not only larger but depends much more sensitively on R_{eff} .

In fact, ρ_d is *more* negative than $\langle \rho \rangle$, indicating a *stronger* residual anticorrelation in the data than predicted by the model. Although only 2.2σ significant, this provides evidence within our framework for a second component of the galaxy–halo connection: an anticorrelation of R_{eff} with halo mass M_{vir} or concentration c at fixed M_b . This would give smaller galaxies on average more dark matter within R_f , and hence larger V_f . Such a correlation has already been suggested by Desmond (2017) on the basis of the correlation of the residuals of the mass discrepancy–acceleration relation (MDAR) with galaxy size. The red histogram in Fig. 2(a) shows the result for the best-fitting correlation found there, $\Delta R_{\text{eff}} \sim -0.4 \Delta c$; that this model gives a good fit to ρ_d suggests the BTFR and MDAR to contain similar information in this regard. This correlation may, however, be in disagreement with the observed ΔR – ΔV correlation when V is measured further in (Desmond & Wechsler 2015).

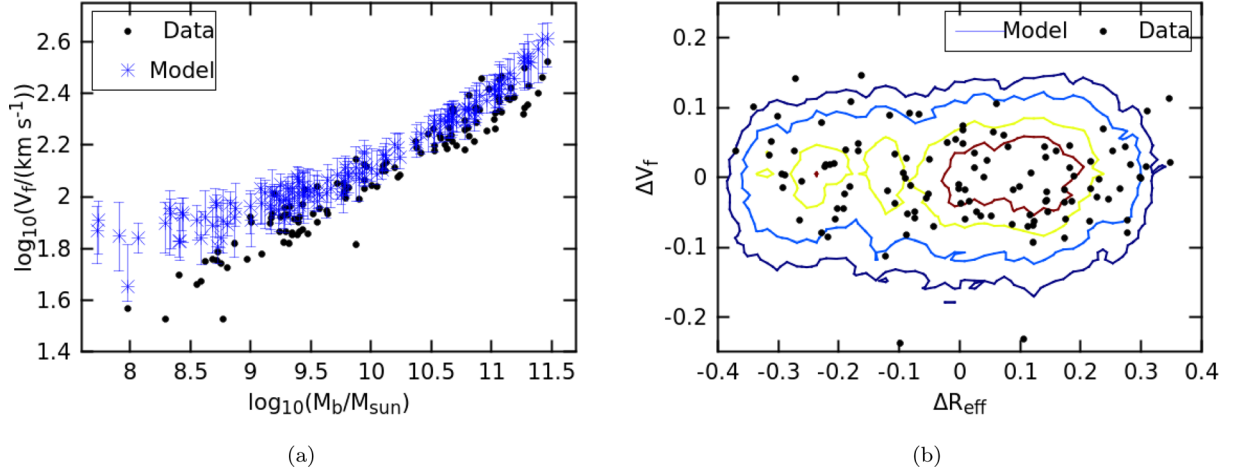


Figure 1. The prediction of AM applied to the SPARC sample for the BTFR (Fig. 1a) and $\Delta R_{\text{eff}}-\Delta V_f$ correlation (Fig. 1b), compared to the data itself. Blue stars show the modal V_f over all mock data sets for each SPARC galaxy, and error bars show the 1σ variation. While the model BTFR is curved and has higher scatter than is observed, its residuals are correctly uncorrelated with galaxy size.

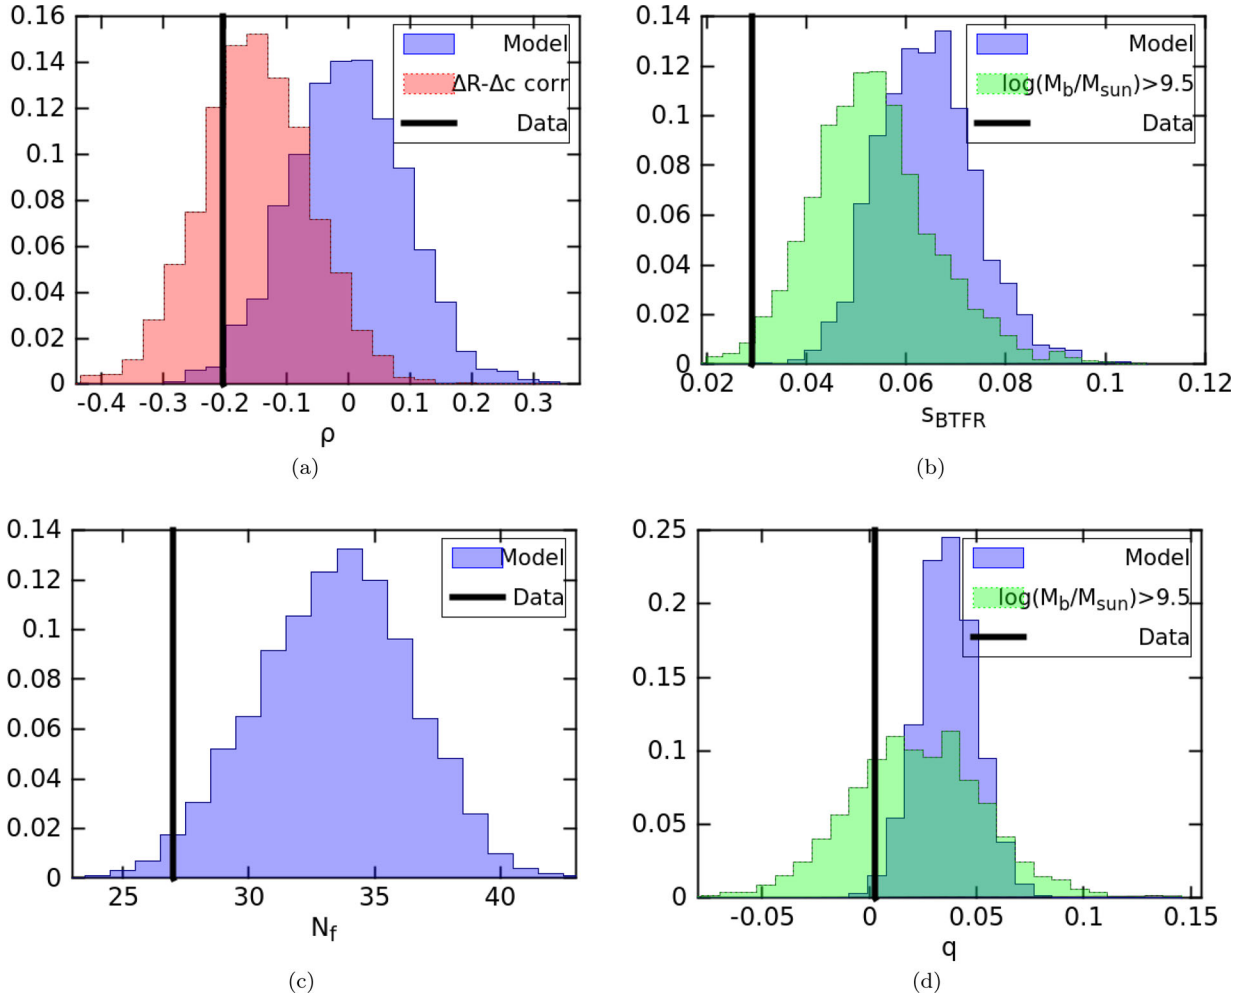


Figure 2. The distributions of four key BTFR statistics predicted by AM – the strength of the $\Delta R_{\text{eff}}-\Delta V_f$ correlation ρ , the intrinsic scatter S_{BTFR} , the number N_f of galaxies with non-flat RCs, and the curvature q – compared to the values in the SPARC data. The results are quantified in Table 1. ‘ $\Delta R-\Delta c$ corr’ in Fig. 2(a) denotes an anticorrelation of R_{eff} and c residuals, as described in Section 3.1.

Table 1. Comparison of SPARC and model BTFR statistics. ρ is Spearman’s rank coefficient of the $\Delta R_{\text{eff}} - \Delta V_f$ correlation, s_{BTFR} is the intrinsic BTFR scatter, N_f is the number of galaxies with non-flat RCs and q is the quadratic BTFR curvature. The third row shows s_{BTFR} for $M_b > 10^{9.5} M_\odot$ galaxies only.

Statistic	SPARC	Model mean	Discrepancy/ σ
ρ	−0.20	0.00	2.2
s_{BTFR} (dex)	0.029	0.064	3.6
" ($M_b > 10^{9.5} M_\odot$)	0.027	0.053	2.1
N_f	27	33.4	2.2
q	0.003	0.039	3.0

3.2 Scatter

We now proceed to s_{BTFR} . The blue histogram in Fig. 2(b) shows the distribution of this statistic in the model, and the second row of Table 1 lists $\langle s_{\text{BTFR}} \rangle$, s_s , s_d and σ_s . As anticipated by L16, and in agreement with Dutton (2012) and Di Cintio & Lelli (2016), we find the predicted BTFR scatter to be upwards of 0.15 dex in M_b , with typical mock data sets having ~ 0.25 dex.² Given the spread among mock data sets, this is 3.6σ discrepant with the SPARC value of ~ 0.11 dex. This is significant – none of our 2000 mock data sets have $s_{\text{BTFR}} < s_d$ – but not phenomenally so. s_{BTFR} can be reduced to a small degree by tightening the galaxy–halo connection: adopting an AM scatter of 0 reduces $\langle s_{\text{BTFR}} \rangle$ to 0.061, with a corresponding discrepancy of 3.2σ .

It is evident from Fig. 1(a) that the predicted BTFR scatter rises towards lower M_b . To quantify this effect, we show in green in Fig. 2(b) the s_{BTFR} distribution with all $M_b < 10^{9.5} M_\odot$ galaxies removed; this reduces the discrepancy to 2.1σ (Table 1, row 3). The model prediction may be unreliable for $M_b \lesssim 10^{9.5} M_\odot$, as the stellar mass function requires extrapolation, AM cannot be directly tested with clustering, and low-mass haloes may not be fully resolved. We note also that gas mass fractions rise rapidly below $M_b \sim 10^{9.5} M_\odot$, amplifying any potential error incurred by performing AM with stellar as opposed to total baryonic mass. An M_b -based AM would correlate M_b more strongly with M_{vir} and c , likely reducing s_{BTFR} .

3.3 Rotation curve flatness

Since we eliminate galaxies in each mock data set with non-flat RCs, Figs 1, 2(a), 2(b) and 2(d) pertain only to a subset of the full sample. An orthogonal statistic with which to compare model and data, therefore, is the number of galaxies out of the original 150 that fail the flatness cut, which we denote N_f . Fig. 2(c) shows the distribution of N_f over all the mock data sets compared to the value in SPARC (27), and the corresponding statistics are shown in the fourth row of Table 1. We find a larger fraction of our model galaxies to have rising RCs at the last measured point than in the data (a 2.2σ discrepancy), reflecting the fact the NFW density profile falls as $\sim r^{-1}$ out to large radius. This quantifies the longstanding ‘disc–halo conspiracy’ (van Albada & Sancisi 1986), and deserves attention in future studies.

3.4 Curvature

A final significant feature of the model BTFR is its curvature, which is $\sigma_q = 3.0\sigma$ discrepant with the data (Fig. 2d and Table 1, final

row). We caution, however, that this prediction is sensitive to the low- M_b model uncertainties described in Section 3.2, and removing the $M_b < 10^{9.5} M_\odot$ region reduces σ_q below 1σ (Fig. 2d, green histogram).

In the context of AM, BTFR curvature follows from the well-known mismatch between the slopes of the stellar and halo mass functions, and is therefore present to some extent in all AM-based studies (e.g. Trujillo-Gomez et al. 2011; Desmond 2012). The precise magnitude of q , however, depends on the details of the AM: lower curvature follows from a shallower bright end to the SMF (e.g. the photometry of Bernardi et al. 2013, used here), a halo proxy with less concentration dependence, and a lower AM scatter. Clustering studies have begun to set strong constraints on these variables, and we find that q varies by only ~ 5 per cent as the halo proxy and AM scatter span the ranges allowed by Lehmann et al. (2017). This suggests that spatial statistics provide sufficient information on AM for the shape of the predicted BTFR to follow almost uniquely. Older AM schemes that match to halo mass directly – as used, for example, in Di Cintio & Lelli (2016) – produce a straighter BTFR, although that change alone reduces q by 15 per cent at most and cannot lower σ_q below 2.5σ . A further reduction requires a preferential decrease in halo M_{vir} or c at the faint and/or bright ends (e.g. by baryonic feedback), biases from selection effects, or systematic errors in stellar mass measurements. Some hydrodynamical simulations incorporating these effects have achieved a straighter BTFR (e.g. Santos-Santos et al. 2016; Dutton et al. 2017). q likely depends in addition on the velocity measure: when sampled much beyond R_{max} , V will be lowered for high-mass galaxies with falling RCs, reducing the upturn in the BTFR at the bright end. This may also contribute to the lower curvature of Di Cintio & Lelli (2016), who measure V at $8R_d$.

4 DISCUSSION AND CONCLUSIONS

A range of approaches have been developed in the past two decades to elucidate the nature and origin of the BTFR, but only recently have data and models become sufficiently sophisticated for statistically rigorous analysis to be possible. To advance this programme, we calculate the significance levels at which four key statistics of a state-of-the-art observational BTFR data set, SPARC, differ from those expected in a modern Λ CDM AM model. We create mock data sets with precisely the baryonic properties of SPARC, and analyse them in an identical fashion to the real data (L16). Any differences between model and measured galaxies must therefore derive solely from differences in the distribution of dark matter, and hence be attributable to the galaxy–halo connection. Our findings are the following:

(i) When defined using the flat part of the RC, the BTFR’s residuals would *not* be expected to anticorrelate with galaxy size in Λ CDM; a significant test of galaxy formation requires that velocity be measured at a *smaller* radius, where the contribution of the baryonic mass depends more strongly on its concentration. On the other hand, the fact that the baryonic part of V_f depends little on R_{eff} makes their correlation more sensitive to a second global galaxy–halo correlation (after $M_* - f(M_{\text{vir}}, c)$ imposed by AM), viz the relation between galaxy size and halo properties. We find $\sim 2\sigma$ evidence for an anticorrelation of R_{eff} with M_{vir} or c at fixed M_b .

(ii) The predicted BTFR scatter is 3.6σ larger than observed, and cannot be appreciably lowered by tightening the galaxy–halo connection. However, simulation and model uncertainties may impact

² The scatters in M_b may be approximately obtained from the quoted scatters in V_f by dividing by the BTFR slope, ~ 0.25 .

the prediction at $M_b \lesssim 10^{9.5} M_\odot$, and excising this region reduces the discrepancy to 2.1σ .

(iii) A further BTFR statistic with significant constraining power for models of the galaxy–halo connection is the quadratic curvature, which we find to be overpredicted by AM at the 3.0σ level. This may indicate mass-dependent baryonic effects on the dark matter haloes (e.g. Di Cintio et al. 2014; Sawala et al. 2016) or a correlation of TFR selection criteria with dynamical halo properties.

(iv) For BTFR studies that remove galaxies with non-flat RCs (e.g. SPARC), the fraction of such galaxies in a given data set is orthogonal to other BTFR statistics and provides an additional handle on RC shape. We find AM to overpredict the fraction of galaxies with rising RCs at 2.2σ , suggesting that its application to N -body halo density profiles does not fully satisfy the observed ‘disc–halo conspiracy’.

(v) The significance of discrepancies between theoretical and observed BTFR statistics is set by sample variance among model realizations, which scales inversely with the size of the data set. As statistical tests are rarely performed in the literature, the importance of sample size is often overlooked (but see Sorce & Quan 2016). Assuming BTFR statistics to obey the central limit theorem, the widths of mock data distributions – and hence significance levels – will vary with $\sim\sqrt{N}$, suggesting that increasing sample size may be preferable for increasing the power of statistical tests than applying stringent cuts on data quality. Because our model explicitly includes the size of the SPARC sample (in addition to baryonic galaxy properties imported directly from the observations), our results should strictly be considered to apply only to the SPARC BTFR, not to the BTFR per se.

We propose three directions in which this work could be taken: (1) seek the features of the BTFR – or galaxy dynamics more generally – with most constraining power for galaxy formation, and calculate the significance of their deviations from AM predictions; (2) modify the model to alleviate the aforementioned discrepancies, for example, by complexifying the galaxy–halo connection or introducing new degrees of freedom for baryonic effects; and (3) incorporate improvements in simulation resolution and survey depth to strengthen the AM prediction for $M_b < 10^{9.5} M_\odot$. This regime is not only critical for the statistical power of BTFR tests, but also connects to topical galaxy formation issues at the dwarf scale.

ACKNOWLEDGEMENTS

I thank Federico Lelli for guidance with the SPARC data, and Federico Lelli, Stacy McGaugh, Risa Wechsler and an anonymous referee for comments on the manuscript.

This work used the DarkSky simulations, made using an INCITE 2014 allocation on the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory. I thank the DarkSky collaboration for creating and providing access to these simulations and Sam Skillman and Yao-Yuan Mao for running ROCKSTAR and CONSISTENT TREES on them.

This work received support from the US Department of Energy under contract number DE-AC02-76SF00515.

REFERENCES

- Behroozi P. S., Wechsler R. H., Wu H.-Y., 2013, *ApJ*, 762, 109
 Bernardi M., Meert A., Sheth R. K., Vikram V., Huertas-Company M., Mei F., Shankar F., 2013, *MNRAS*, 436, 697
 Conroy C., Wechsler R. H., Kravtsov A. V., 2006, *ApJ*, 647, 201
 Desmond H., 2012, preprint ([arXiv:1204.1497](https://arxiv.org/abs/1204.1497))
 Desmond H., 2016, *MNRAS*, 464, 4160
 Desmond H., Wechsler R. H., 2015, *MNRAS*, 454, 322
 Di Cintio A., Lelli F., 2016, *MNRAS*, 456, L127
 Di Cintio A., Brook C. B., Macciò A. V., Stinson G. S., Knebe A., Dutton A. A., Wadsley J., 2014, *MNRAS*, 437, 415
 Dutton A. A., 2012, *MNRAS*, 424, 3123
 Dutton A. A. et al., 2017, *MNRAS*, 467, 4937
 Kravtsov A. V., Berlind A. A., Wechsler R. H., Klypin A. A., Gottlober S., Allgood B., Primack J. R., 2004, *ApJ*, 609, 35
 Lehmann B. V., Mao Y.-Y., Becker M. R., Skillman S. W., Wechsler R. H., 2017, *ApJ*, 834, 37
 Lelli F., McGaugh S. S., Schombert J. M., 2016a, *ApJ*, 816, L14 (L16)
 Lelli F., McGaugh S. S., Schombert J. M., 2016b, *AJ*, 152, 157
 Santos-Santos I. M., Brook C. B., Stinson G., Di Cintio A., Wadsley J., Domínguez-Tenreiro R., Gottlober S., 2016, *MNRAS*, 455, 476
 Sawala T. et al., 2016, *MNRAS*, 457, 1931
 Skillman S. W., Warren M. S., Turk M. J., Wechsler R. H., Holz D. E., Sutter P. M., 2014, preprint ([arXiv:1407.2600](https://arxiv.org/abs/1407.2600))
 Sorce J. G., Quan G., 2016, *MNRAS*, 458, 2667
 Trujillo-Gomez S., Klypin A., Primack J., Romanowsky A. J., 2011, *ApJ*, 742, 16
 van Albada T. S., Sancisi R., 1986, *Phil. Trans. Roy. Soc. Lon. Ser. A*, 320, 447
 Warren M. S., 2013, 2HOT: An Improved Parallel Hashed Oct-Tree N-Body Algorithm for Cosmological Simulation. ACM, New York, p. 72

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