

FAINT GALAXIES AND SMALL HALOS:
PROBES OF GALAXY FORMATION AND DARK MATTER

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Ethan Oliver Nadler
June 2021

© 2021 by Ethan Oliver Nadler. All Rights Reserved.
Re-distributed by Stanford University under license with the author.



This work is licensed under a Creative Commons Attribution-
Noncommercial 3.0 United States License.
<http://creativecommons.org/licenses/by-nc/3.0/us/>

This dissertation is online at: <http://purl.stanford.edu/qc509cs6632>

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Risa Wechsler, Primary Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Tom Abel

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Roger Blandford

Approved for the Stanford University Committee on Graduate Studies.

Stacey F. Bent, Vice Provost for Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Preface

Our best scientific model of the Universe on large scales— Λ CDM cosmology—predicts that most of the matter in the cosmos is *dark*. Without this gravitating substance, known as cold dark matter (CDM), the cosmic structures we inhabit including galaxies, stars, and planets would never form.

Λ CDM cosmology does not hold the same status as physical theories like general relativity and quantum mechanics because it is *descriptive*. This is highlighted by its two building blocks: dark matter and dark energy. In particular, the CDM particle is assumed to be cold (i.e., to move slowly on cosmological scales) and to have feeble—if any—non-gravitational interactions with both itself and Standard Model particles. This begs the question: how cold and how feebly interacting is dark matter?

The modern picture of galaxy formation in the Λ CDM context also sows its own foundational questions. A vast and compelling body of evidence teaches us that *galaxies live in dark matter halos*, and that more massive halos tend to host brighter galaxies. Open questions remain at the limits of these generic statements: What is the smallest galaxy that can form? Does every dark matter halo host a galaxy? And, are the answers to these questions influenced by non-gravitational dark matter physics, or only by astrophysical processes?

This thesis addresses the intersection of dark matter and galaxy formation physics by modeling and analyzing the faintest observable galaxies in the Universe in a cosmological context. Following an introduction (Chapter 1), it proceeds in two parts. Part I (Chapters 2–5) develops an empirical model for the connection between faint galaxies and small dark matter halos in a CDM context, and applies this model to state-of-the-art observations of the Milky Way satellite galaxy population to infer the properties of the dark matter halos that host dwarf galaxies. Part II (Chapters 6–10) expands these analyses for several modifications to the CDM paradigm, and thereby places new and stringent constraints on dark matter particle properties. Chapter 11 combines dark matter constraints from dwarf galaxy abundances and strong gravitational lensing, paving the way for next-generation surveys to deliver unprecedented insights into dark matter physics. Finally, Chapter 12 discusses the outlook for these upcoming surveys and Chapter 13 concludes.

Acknowledgments

I count myself extremely fortunate to have the opportunity to pursue a Ph.D. in cosmology. Reflecting on the happy circumstances that allowed this to happen, my deepest gratitude starts at the beginning: to my parents, for emphasizing the importance of science and encouraging me to take the road less traveled, and to my brother, for reminding me that physicists often need to be tempered with a good dose of philosophical reasoning.

Although my family played the most important role in my early education, my experiences in community college and as an undergraduate were the best preparation I could have asked for before graduate school. To my teachers and mentors at the Santa Fe Community College (particularly Jerry Friedman and Robert Shankin) and at the University of California, Santa Barbara and its College of Creative Studies (particularly Sathya Guruswamy and Peng Oh)—thank you for everything! I wouldn't be writing this thesis without you.

There are many amazing scientists who contributed their time and energy to my growth as a graduate student. Above all, my advisor Risa Wechsler was incredibly supportive and dedicated to our science while consistently providing me with opportunities to excel. No amount of thanks is sufficient, but I hope the joy we both derived from doing the science in this thesis counts!

Each project in this thesis benefited from collaborations with postdocs and professors, many of whom have been associated with KIPAC. I am deeply grateful to Yao-Yuan Mao for easing my transition into Risa's group, providing advice and technical expertise on our satellite galaxy model, and helping frame several of the chapters below. I am likewise grateful to Keith Bechtol and Alex Drlica-Wagner for their efforts over many years that allowed the model to be interfaced with data; to Kimberly Boddy and Vera Gluscevic for their theoretical guidance that allowed us to place qualitatively new constraints on dark matter interactions; to Susmita Adhikari and Arka Banerjee for their insights that helped connect my work to dark matter self-interactions; to Simon Birrer for working together to bridge the gap between Milky Way satellite and strong lensing observations; to Gregory Green for helping develop the statistical framework underlying our satellite model; and to Subinoy Das for his ideas and expertise on late-forming dark matter. I would also like to emphasize my gratitude to members of the Dark Energy Survey Milky Way Working Group, the Dark Energy Science Collaboration Dark Matter Working Group, and the Satellites Around Galactic Analogs

Survey for their efforts to facilitate our scientific collaborations.

Others at KIPAC and throughout Stanford have also provided crucial guidance. I am particularly grateful to Tom Abel for his continual mentorship, Roger Blandford for asking the hard questions, and Leonardo Senatore for teaching me to think like an effective field theorist. Importantly, all of the graduate students, undergraduates, and postdocs associated with Risa’s group throughout the years have been a constant source of knowledge and support. Special thanks go to Joe, Sean, Chun-Hao, Nick, Sebastian, Elise, Richie, Sid, and those I have mentored (or at least, tried my best to!), including Abby, Nicel, Veronica, and Deveshi. Lastly, thanks to my interdisciplinary research colleagues for keeping me grounded and curious.

Like the contents of the Universe, the visible side of my graduate career is only the surface. I am lucky to have a group of true friends who made disconnecting from work enjoyable while challenging me to push myself in realms outside of physics. Those connections mean the world to me.

Contents

Preface	iv
Acknowledgments	v
1 Introduction	1
1.1 Preliminaries	2
1.2 Dark Matter and Hierarchical Structure Formation	3
1.3 Galaxy Formation in Dark Matter Halos	4
1.4 Dwarf Galaxies around the Milky Way	5
1.5 Connecting Dark Matter to Dwarf Galaxies	6
2 Modeling the Impact of Baryons on Subhalo Populations	8
2.1 Paper Status and External Contributions	9
2.2 Introduction	9
2.3 Simulation Data	12
2.4 Random Forest Classification	14
2.4.1 Overview	14
2.4.2 Choice of Subhalo Features	15
2.4.3 Training and Validation	19
2.4.4 Importance of Subhalo Features	22
2.4.5 Model Limitations	22
2.5 Results	23
2.5.1 Predictions for Dark Matter-only Simulations of the FIRE Halos	23
2.5.2 Predictions for a Suite of Dark Matter-only Zoom-in Simulations	30
2.6 Conclusions and Discussion	36
3 Modeling the Subhalo–Satellite Connection	38
3.1 Paper Status and External Contributions	38
3.2 Introduction	39

3.3	Simulations	41
3.4	Model	42
3.4.1	Host Halo Properties	42
3.4.2	Satellite Luminosities	43
3.4.3	Satellite Locations	46
3.4.4	Satellite Sizes	47
3.4.5	Subhalo Disruption due to Baryonic Effects (\mathcal{B})	48
3.4.6	Orphan Satellites (\mathcal{O})	51
3.4.7	Comparison to Recent Models	53
3.5	Comparison to Observed Satellites	54
3.5.1	Qualitative Comparison	54
3.5.2	Quantitative Comparison	56
3.6	Results and Discussion	60
3.6.1	Derived Constraints	60
3.6.2	Predictions for Future Surveys	62
3.6.3	Implications for the Connection between Low-mass Subhalos and Faint Satellites	63
3.6.4	Caveats and Future Work	67
3.7	Summary	68
4	Galaxy–Halo Connection Constraints Including the LMC	70
4.1	Paper Status and External Contributions	70
4.2	Introduction	71
4.3	Analysis Overview	73
4.4	Simulations	73
4.4.1	General Description	73
4.4.2	Host Halo and LMC Analog Selection	75
4.5	Galaxy–Halo Connection Model	76
4.5.1	Satellite Luminosities	76
4.5.2	Satellite Sizes	77
4.5.3	Subhalo Disruption Due to Baryonic Effects	78
4.5.4	Galaxy Formation Efficiency	78
4.5.5	Orphan Satellites	78
4.6	Observational Selection Functions	79
4.7	Statistical Framework	80
4.8	Results	81
4.8.1	Observed Satellite Populations	82
4.8.2	The Impact of the LMC	83
4.8.3	The Total MW Satellite Population	85

4.8.4	Galaxy–Halo Connection Model Constraints	88
4.8.5	Properties of Halos that Host the Faintest Satellites	90
4.8.6	Implications for Dark Matter Microphysics	92
4.9	Theoretical Uncertainties	93
4.10	Conclusions	95
5	Placing the Milky Way in Context	96
5.1	Paper Status and External Contributions	96
5.2	Satellites of Milky Way Analogs and the SAGA Survey	97
5.3	SAGA Stage II: Implications for the Galaxy–Halo Connection	98
5.3.1	The Model	98
5.3.2	Comparing Model Predictions with SAGA Results	99
5.3.3	Implications	100
5.4	Forecasts for SAGA Stage III	104
6	Dark Matter Constraints. I. DM–SM Interactions	108
6.1	Paper Status and External Contributions	108
6.2	Introduction	109
6.3	Analytic Estimate from Individual Halos	110
6.3.1	Minimum Halo Mass in an Interacting Cosmology	110
6.3.2	Limits on the Interaction Cross Section	113
6.4	Satellite Population Likelihood Analysis	113
6.5	Implications for Other Non-CDM Models	116
6.6	Conclusions and Discussion	117
7	Dark Matter Constraints. II. Warm, Interacting, & Fuzzy DM	119
7.1	Paper Status and External Contributions	119
7.2	Introduction	120
7.3	Analysis Overview	121
7.3.1	Transfer Function	121
7.3.2	Subhalo Mass Function	123
7.3.3	MW Satellite Model	123
7.3.4	Fitting Procedure	124
7.4	WDM Analysis	125
7.5	IDM Analysis	126
7.6	FDM Analysis	127
7.7	Results	129
7.8	Discussion	130

8 Dark Matter Constraints. III. DM Formation Epoch	133
8.1 Paper Status and External Contributions	133
8.2 Introduction	134
8.3 Late-forming Dark Matter Models	136
8.3.1 Self-interacting LFDM	137
8.3.2 Free-streaming LFDM	138
8.4 Linear Perturbations	139
8.4.1 Free-streaming LFDM	139
8.4.2 Self-interacting LFDM	140
8.5 Transfer Functions	140
8.5.1 Self-interacting LFDM	141
8.5.2 Free-streaming LFDM	143
8.6 Constraints from Milky Way Satellites	143
8.6.1 Self-interacting LFDM	146
8.6.2 Free-streaming LFDM	147
8.7 Conclusion	148
9 Dark Matter Self-interactions in Milky Way-mass Halos	149
9.1 Paper Status and External Contributions	150
9.2 Introduction	150
9.3 SIDM Model	151
9.4 Simulations	154
9.4.1 General Description	154
9.4.2 SIDM Implementation	154
9.4.3 Subhalo Definitions and Resolution Cuts	155
9.4.4 Subhalo Matching Procedure	156
9.5 SIDM Effects on the MW Host Halo	157
9.6 SIDM Effects on MW Subhalos	160
9.6.1 Pre-infall Subhalo Evolution	160
9.6.2 Post-infall Subhalo Evolution	161
9.6.3 Surviving Subhalo Populations	165
9.7 Challenges	169
9.7.1 Numerical Effects	169
9.7.2 Sample Variance	170
9.7.3 Baryonic Effects	170
9.8 Comparison to Previous Studies	172
9.9 Prospects for SIDM Constraints	173
9.10 Conclusions	174

10 Disentangling Dark Matter and Baryonic Physics	177
10.1 Paper Status and External Contributions	177
10.2 Introduction	178
10.3 Methods	179
10.3.1 MW-like CDM Simulations	179
10.3.2 SIDM Simulations	181
10.3.3 WDM Model	181
10.3.4 Disk Disruption Model	183
10.4 Results	183
10.4.1 Subhalo Velocity Function	183
10.4.2 Subhalo Radial Distribution	185
10.4.3 Spatial Anisotropy in the MW Subhalo Population	186
10.5 Discussion and Conclusion	187
11 Combining Milky Way Satellites and Strong Lensing	189
11.1 Paper Status and External Contributions	190
11.2 Introduction	190
11.3 Dark Matter Substructure Model	192
11.3.1 Projected subhalo mass function	192
11.3.2 CDM host mass and redshift dependence	193
11.3.3 Subhalo disruption efficiency Due to baryons	194
11.3.4 Warm dark matter subhalo mass function	196
11.3.5 Warm dark matter mass–concentration relation	197
11.4 Milky Way Substructure Modeling	198
11.4.1 Milky Way Satellite Data	198
11.4.2 Milky Way Satellite Model	198
11.4.3 Constraints from Milky Way Satellite Observations	200
11.5 Strong Lens Substructure Modeling	202
11.5.1 Strong Lensing Data	202
11.5.2 Strong Lensing Model	204
11.5.3 Constraints from Strong Lensing Observations	204
11.6 Joint Analysis Methodology	205
11.6.1 Probe Combination Procedure	205
11.6.2 Inferring Σ_{sub} from Milky Way Satellites	207
11.6.3 Probe Combination Statistics	210
11.7 Results	211
11.7.1 Conventions for WDM Constraints	212
11.7.2 Fiducial WDM Constraints	213

11.7.3	Impact of the Differential Subhalo Disruption Efficiency due to Baryons	214
11.7.4	Projected Subhalo Number Density Constraints	216
11.8	Discussion	216
11.8.1	Systematics	217
11.8.2	Comparison to Recent Studies	218
11.9	Conclusions	220
12	Forecasts for Future Surveys	223
12.1	Paper Status and External Contributions	223
12.2	Forecasts for the Minimum Halo Mass Inferred from LSST Milky Way Satellite Observations	224
12.3	Half-mode Mass Forecasts Including the Large Magellanic Cloud Satellite Population	226
13	Conclusions	229
A	Random Forest Scatter, Feature Selection, & Resolution	232
B	Satellite Model Statistical Methods	236
B.1	Convergence of Improved Likelihood to the Poisson Distribution	236
B.2	Bayesian Evidence for Radial Scaling, Size Reduction, and Orphan Satellites	237
C	Satellite Model Data Comparison Details	239
C.1	Galaxy–Halo Connection Model Details	239
C.1.1	MW Host Halo Mass	239
C.1.2	Mass-dependent Scatter	240
C.1.3	Radial Scaling	241
C.1.4	Tidal Stripping	242
C.1.5	Concentration-dependent Satellite Sizes	242
C.1.6	Orphan Satellite Contribution	244
C.2	Statistical Framework Details	244
C.2.1	Poisson Process Likelihood	245
C.2.2	Priors	245
C.2.3	Bayes Factor Calculation	245
C.2.4	Caveats and Future Work	246
C.3	Robustness to Observational Systematics	247
C.3.1	Kinematically Unconfirmed Satellites	247
C.3.2	Satellite Size Criterion	248
C.3.3	Biases in Measured Satellite Properties	248
C.4	Resolution and Sample Variance	248

C.5	Observed Satellite Data Vectors	249
D	Warm, Interacting, & Fuzzy Dark Matter Constraints Supplement	252
D.1	Milky Way Satellite Model	252
D.2	Fitting Procedure Details and Posterior Distributions	253
D.3	Resonantly-Produced Sterile Neutrino Constraints	254
D.4	FDM Subhalo Mass Functions	257
E	SIDM Resolution Tests	258
E.1	Simulation Resolution Tests	258
E.2	Subhalo Resolution Threshold	259
E.2.1	Subhalo Abundance for a Fixed V_{peak} Threshold	259
E.2.2	Choice of Fiducial V_{peak} Threshold	260
E.2.3	Severity of Subhalo Disruption for a Less Conservative V_{peak} Threshold	260
E.3	Subhalos near the Host Center	261
F	Probe Combination Details	262
F.1	Milky Way Zoom-in Simulations	262
F.2	Comparing Milky Way Zoom-in Simulations to <i>Galacticus</i>	265
F.3	Host-to-host, Poisson, and projection scatter in $\Sigma_{\text{sub,MW}}$	266
F.4	Subhalo Disruption Efficiency Prior	267
	Bibliography	269

List of Tables

2.1	Performance metrics for five different random forest classifiers trained on disrupted and surviving subhalos from the <code>m12i</code> and <code>m12f</code> FIRE simulations with $V_{\text{peak}} > 10 \text{ km s}^{-1}$. The first column lists the subhalo features used to train each classifier. The second column lists the out-of-bag classification score, which is the percentage of subhalos in the training data identified correctly when each tree does not vote on subhalos in its own training set. The third and fourth columns list the percentage of disrupted and surviving subhalos in the test set that are identified correctly by each classifier, averaged over 100 test-training splits. The test set is the collection of subhalos from the <code>m12i</code> and <code>m12f</code> FIRE simulations with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ that are not included in the training set. We also indicate the standard deviation of each classification accuracy. Note that the ratio of disrupted to surviving subhalos in our fiducial halo catalog is roughly 3:2.	17
3.1	Summary of the physical ingredients, underlying assumptions, and parameterizations of various processes that enter our model for the subhalo–satellite connection. The final column indicates whether each component of the model is held fixed or allowed to vary for our fit to the luminosity distribution of classical and SDSS-identified satellites in Section 3.5. Bold values correspond to parameters that are varied in our fit to the observed luminosity function.	40
3.2	Prior distributions for the parameters varied in our fit to the luminosity distribution of classical and SDSS-identified MW satellites.	59
4.1	Galaxy–halo connection model constraints derived from our fit to the DES and PS1 satellite populations. Asterisks mark prior-driven constraints.	87

7.1 Constraints on the WDM, IDM, and FDM paradigms from observations of MW satellite galaxies. Limits for each non-CDM model are derived by assuming that it constitutes the entirety of the DM. The first column lists the DM paradigm, the second column describes the particle physics parameters constrained by this analysis, the third column lists the corresponding constraints at 95% confidence, the fourth column describes the derived property constrained for each DM model, and the fifth column lists constraints on the derived parameters. Limits on the DM-proton scattering cross sections depend on the DM particle mass, m_χ (see Figure 7.2); for simplicity, we present our constraint for $m_\chi = 100$ MeV.	125
9.1 Summary of SIDM model variants and simulation results. SIDM model variants considered in this work and the main qualitative results of our zoom-in simulation of an MW-mass system for each case. The first column lists the DM model, the second and third columns list the characteristic velocity scale (in km s^{-1} and amplitude of the self-interaction cross section (in $\text{cm}^2 \text{ g}^{-1}$) for our SIDM model variants, the fourth (fifth) column lists whether the host halo (subhalos) are cored by self-interactions, the sixth column lists whether subhalos are affected by ram pressure stripping due to self-interactions with the host, and the seventh column lists the fraction of subhalos among our matched subhalo populations that survive to $z = 0$ relative to the number of surviving subhalos with $V_{\text{peak}} > 20 \text{ km s}^{-1}$ in our CDM simulation.	155
10.1 Impact of SIDM, WDM, and the Galactic disk on the MW subhalo population. The first column lists the MW subhalo population model. The second and third columns describe the suppression of the subhalo velocity function and radial distribution. The fourth column describes whether the suppression of the subhalo population is anisotropic. The fifth column lists the fraction of subhalos with $V_{\text{peak}} > 20 \text{ km s}^{-1}$ in the hemisphere that points from the MW halo center towards the LMC, averaged over our two MW-like simulations. The final column describes the main physical effect that impacts the MW subhalo population. The asterisk indicates that the shape of the WDM radial distribution may change for subhalos below our V_{peak} threshold.	179

11.1 95% confidence and 20:1 likelihood ratio upper limits on M_{hm} and corresponding lower limits on m_{WDM} for our multidimensional probe combination for various differential subhalo disruption efficiency values q , and for an analysis that combines the fully marginalized one-dimensional M_{hm} distributions. $q = 0.5$ corresponds to twice as efficient subhalo disruption due to baryons in the MW relative to strong lenses, $q = 1$ (our fiducial model) corresponds to equally efficient subhalo disruption due to baryons, and $q = 2$ corresponds to twice as efficient subhalo disruption due to baryons in strong lenses.	213
C.1 Prior distributions for the parameters varied in our fiducial eight-parameter fit to the DES and PS1 satellite populations. Here $\mathcal{N}(\mu, \sigma)$ denotes a normal distribution with mean μ and standard deviation σ	244
C.2 MW satellites used in our analysis. Properties of confirmed and candidate DES and PS1 satellites used in our analysis, listed in order of detection significance (Paper I). Asterisks mark kinematically unconfirmed systems. ^a Eridanus II is not included because it lies outside our fiducial 300 kpc heliocentric distance cut. ^b Cetus II is detected in both PS1 and DES; in our analysis, we only count this system in the observed DES population.	251
D.1 Galaxy–halo connection and DM model parameters varied in our thermal relic WDM and FDM fits to the MW satellite population. Note that M_0 is constrained in a separate fit that yields similar confidence intervals for the eight galaxy–halo connection parameters. Asterisks mark prior-driven constraints. See [333] for details on our galaxy–halo connection model.	254

List of Figures

2.1	Normalized joint and marginal distributions of pericentric distance and scale factor at first pericentric passage after accretion for surviving (blue) and disrupted (red) subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ in the <code>m12i</code> and <code>m12f</code> FIRE simulations. We select disrupted subhalos starting at $a = 0.25$ ($z = 3$).	16
2.2	Normalized joint and marginal distributions of maximum circular velocity and virial mass at accretion for surviving (blue) and disrupted (red) subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ in the <code>m12i</code> and <code>m12f</code> FIRE simulations. We select disrupted subhalos starting at $a = 0.25$ ($z = 3$); note that $V_{\text{acc}} < V_{\text{peak}}$ for subhalos that are stripped prior to infall (e.g., see [46]).	18
2.3	True- vs. false-positive classification rate for our fiducial five-property random forest classifier, which is trained on subhalos from the <code>m12i</code> and <code>m12f</code> FIRE simulations with $V_{\text{peak}} > 10 \text{ km s}^{-1}$. These classification rates apply to subhalos that are not included in the training set. The AUC is equal to 1 for a perfect classifier (red), 0.93 for our random forest classifier (blue), and 0.5 for a random classifier (black).	20
2.4	Feature importance scores for the five subhalo properties used to classify disrupted and surviving subhalos in the <code>m12i</code> and <code>m12f</code> FIRE simulations. The colored bars above each property indicate the feature importance scores averaged over 100 test-training splits when that property is added to the training features. Thus, the columns correspond to the five different classifiers in Table 2.1. For a given classifier, each property's score indicates its relative importance for classifying disrupted and surviving subhalos. Here d_{peri} and a_{peri} are the pericentric distance and scale factor at first pericentric passage after accretion, and a_{acc} , M_{acc} , and V_{acc} are the scale factor, virial mass, and maximum circular velocity at accretion.	21

2.9 Mass accretion histories for the suite of DMO zoom-in simulations of MW-mass host halos presented in [315]. The black line shows the mean mass accretion history for the 45 hosts, and the shaded area shows the associated $\pm 1\sigma$ standard deviation. Mass accretion histories for the `m12i` and `m12f` FIRE simulations are shown in orange and green, respectively. 32

2.10 Velocity functions (top) and radial distributions (bottom) for the suite of 45 zoom-in simulations of MW-mass host halos presented in [315]. The thick lines show the mean number of subhalos predicted by the DMO simulations (black) and by our random forest classifier (blue), which is trained on the `m12i` and `m12f` FIRE simulations; the shaded areas show the $\pm 1\sigma$ standard deviation of these predictions. The thin lines show the DMO result and the most probable random forest prediction for each host. The thick dotted lines show the mean DMO velocity function and radial distribution scaled by a factor of 2/3 for visual comparison, and the orange and green lines show the results for `m12i` and `m12f`, respectively. Note that the scaled DMO line in the top panel is mostly obscured by the random forest prediction. The inset in the bottom panel shows the number of predicted surviving subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ and within 300 kpc of their respective host versus the number of such subhalos in the corresponding DMO simulations. The thick dotted line in the inset shows the constant fraction of surviving subhalos corresponding to the scaled DMO curves, and the thin dash-dotted line shows a 1 : 1 relationship for comparison. The vertical line at $V_{\text{max}} = 9 \text{ km s}^{-1}$ in the top panel represents a conservative resolution limit for these simulations. 33

2.11 Luminosity functions for the DMO zoom-in simulations presented in [315] (black), inferred using the V_{peak} –luminosity abundance-matching relation tuned to the r -band luminosity function from the GAMA survey [295]. We do not apply scatter in the $V_{\text{peak}} - M_{r,o}$ relation to highlight the host-to-host variability. The blue line shows our mean prediction for the luminosity function of surviving satellites, and shaded areas show $\pm 1\sigma$ and $\pm 2\sigma$ standard deviations. We also plot luminosity functions for `m12i` (orange) and `m12f` (green). Here $M_{r,o}$ is the observed r -band luminosity, and the vertical line corresponds to the completeness limit of the SAGA survey. 35

3.1 Example of the relation between satellite luminosity and subhalo peak circular velocity used in our model. We treat the faint-end power-law luminosity function slope α and the constant lognormal scatter σ_M as free parameters. Dark (light) shaded areas show $\pm 1\sigma_M$ ($\pm 2\sigma_M$) scatter. The dashed black line indicates a conservative upper bound on the resolution limit of our simulations, and the red lines indicate the absolute magnitude of the faintest spectroscopically confirmed MW satellite (Segue I; $M_V = -1.5$ mag) and the mean inferred V_{peak} of its subhalo for $\alpha = -1.3$ and $\sigma_M = 0.2$ dex. 44

3.2 Size distributions for mock satellites in our six MW-like host halos, with satellite sizes set by subhalo sizes at accretion ($\beta = 0$) and at $z = 0$ using our prescription for satellite size reduction due to tidal stripping ($\beta = 1$). The distributions are weighted by survival probability using our subhalo disruption model with $\mathcal{B} = 1$ 49

3.3 V_{\max} function (left panel) and radial distribution (right panel) of subhalos with $V_{\text{acc}} > 10 \text{ km s}^{-1}$ and $V_{\max} > 10 \text{ km s}^{-1}$ in one of our zoom-in simulations, shown with (dashed lines) and without (solid lines) our fiducial orphan model ($\mathcal{O} = 1$). The dark blue lines show results from a high-resolution simulation of this host halo, and the blue bands show the Poisson scatter for our prediction that includes all disrupted subhalos tracked to $z = 0$ in the fiducial-resolution simulation. The bottom panels show the ratio of the number of subhalos to the number of subhalos in the corresponding simulation including orphans for our fiducial-resolution (blue) and high-resolution (dark blue) runs. Including orphans brings the fiducial-resolution V_{\max} function and radial distribution into fairly good agreement with the high-resolution results. 50

3.4 Visualizations of the predicted satellite population in one of our MW-like host halos (blue symbols) and the observed population of classical and SDSS-identified MW satellites (black stars) in the absolute magnitude–half-light radius plane (top panels) and the absolute magnitude–heliocentric distance plane (bottom panels) for $\mathcal{M}_{\min} = 5 \times 10^8 \text{ M}_\odot$ (left panels) and $\mathcal{M}_{\min} = 10^8 \text{ M}_\odot$ (right panels). All classical satellite analogs ($M_V \leq -8.8$ mag) and all systems in a mock SDSS footprint that pass both the SDSS surface brightness and completeness radius detection limits (dashed lines; [271, 273]) are plotted as blue stars, while blue circles show systems in the mock SDSS footprint that do not pass both detection criteria. The color bar indicates satellite disruption probability. This realization uses a faint-end slope of $\alpha = -1.3$, a luminosity scatter of $\sigma_M = 0.2$ dex, and a disruption parameter of $\mathcal{B} = 1$; we fix the remaining parameters according to Table 3.1. The error bars in the top-left panel show characteristic uncertainties for M_V and $r_{1/2}$. Pisces I and Pegasus III, which were discovered in SDSS using methods that do not adhere to our assumed detection criteria, are plotted as unfilled stars. 57

3.5 Posterior distribution from our fit to the luminosity distribution of classical and SDSS-identified satellites. Dark (light) shaded contours show 68% (95%) confidence intervals, and shaded areas in the marginal distributions show 68% confidence intervals. Note that σ_M is reported in dex and \mathcal{M}_{\min} is reported as $\log(\mathcal{M}_{\min}/\text{M}_\odot)$ in this plot. 61

3.6 Our prediction for the observed luminosity function of classical and SDSS satellites inferred from our fit to the absolute magnitude distribution of these systems. Dark (light) shaded areas show 68% (95%) confidence intervals. 64

3.7	Left panel: the radial distribution of classical and SDSS satellites inferred from our fit to the observed luminosity distribution of these systems. Dashed red lines show 68% confidence intervals for a fit with satellite radii scaled inward relative to subhalo radii by a factor of 0.8. Right panel: the corresponding size distribution, calculated by setting satellite sizes based on subhalo properties at accretion ($\beta = 0$) with a constant lognormal scatter of $\sigma_R = 0.01$ dex. In both panels, dark (light) shaded areas show 68% (95%) confidence intervals.	65
3.8	Left panel: the total number of satellites within 300 kpc of the MW inferred from our fit to the classical-plus-SDSS luminosity distribution (blue) compared to results from previous studies and to all observed MW satellites including candidate systems (black). Right panel: the total number of satellites with $M_V < -1.5$ mag within 300 kpc of the MW as a function of limiting observable surface brightness. The red line shows our prediction for the total number of satellites with $M_V < -1.5$ mag independent of surface brightness; dashed red lines show 68% Poisson confidence intervals. In both panels, dark (light) shaded areas indicate 68% (95%) confidence intervals.	65
3.9	Joint distribution of subhalo mass and satellite absolute magnitude predicted by our best-fit model with $\mathcal{M}_{\min} = 10^8 M_{\odot}$. The dark (light) contour corresponds to peak (present-day) subhalo virial mass, and the contours indicate satellite number density.	66
4.1	Visualization of our MW satellite modeling framework. In the first step, we perform high-resolution zoom-in simulations of MW-like halos selected from a larger cosmological volume (Section 4.4); in the second step, we paint galaxies onto subhalos using a parametric model for the galaxy–halo connection (Section 4.5); in the third step, we use the observational selection functions derived in Paper I to compute the probability that these satellites would be observed in DES or PS1 imaging data (Section 4.6); and in the final step, we calculate the likelihood of producing the true DES and PS1 satellite populations given many mock satellite population realizations at fixed galaxy–halo connection model parameters (Section 4.7). We then iterate this process to constrain our galaxy–halo connection model.	74
4.2	Predicted DES and PS1 satellite luminosity functions resulting from a joint fit to these satellite populations. Dark (light) blue bands correspond to 68% (95%) confidence intervals from our fiducial eight-parameter galaxy–halo connection model, dashed red lines show the 68% confidence interval for a model using host halos without LMC analogs (“No LMC”), and black lines show the observed luminosity functions within each survey footprint. Our fiducial model, which includes realistic LMC analogs, is decisively favored over the No LMC scenario, with a Bayes factor of $\sim 10^4$	81

4.3	Size distributions derived by fitting our galaxy–halo connection model to the DES and PS1 satellite populations. Dark (light) blue bands correspond to 68% (95%) confidence intervals from our fiducial eight-parameter model, dashed red lines show the 68% confidence interval for a model using host halos without LMC analogs (“No LMC”), and black lines show the observed size distributions.	82
4.4	Left panel: total MW satellite luminosity function inferred from our joint fit to the DES and PS1 satellite populations (blue) compared to the current census of confirmed and candidate MW satellites (black) and the empirical estimate derived in Paper I (gray), which assumes an isotropic satellite distribution and a cored NFW radial satellite distribution. The 68% confidence intervals from hydrodynamic simulations of the Local Group using the FIRE feedback prescription are shown in red [188]. Luminosity function slopes predicted from hydrodynamic simulations with (solid green line) and without (dashed green line) H_2 -based star formation are shown for comparison [330]; these predictions do not account for subhalo disruption due to the Galactic disk. Note that the Paper I prediction (gray) differs from the “All Known Satellites” curve (black) at the bright end because it does not include the LMC, SMC, or Sagittarius. Right panel: The surface brightness distribution of MW satellites with $M_V < 0$ mag and $r_{1/2} > 10$ pc as a function of the limiting observable surface brightness of an all-sky survey. Arrows indicate approximate detection limits for current surveys. Note that LSST Y1 is expected to have similar detection sensitivity to HSC [237, 440, 214, 336].	84
4.5	Posterior distribution from our fit to the DES and PS1 satellite populations. Dark (light) shaded contours represent 68% (95%) confidence intervals. Shaded areas in the marginal distributions and parameter summaries correspond to 68% confidence intervals. Note that σ_M , σ_{gal} , and $\sigma_{\log R}$ are reported in dex, \mathcal{M}_{50} is reported as $\log(\mathcal{M}_{50}/M_\odot)$, and \mathcal{A} is reported in pc.	86
4.6	Left panel: fraction of halos that host galaxies, inferred from our fit to the DES and PS1 satellite populations. The solid line shows the median inferred galaxy occupation fraction, and dark (light) shaded contours represent 68% (95%) confidence intervals. The resolution limit of our simulations is indicated by the dashed vertical line. Right panel: SMHM relation inferred from our fit to the DES and PS1 satellite populations. An extrapolation of the mean SMHM relation derived from more luminous field galaxies is shown in gray [45]. Stars illustrate the mean of the predicted $\mathcal{M}_{\text{peak}}$ range corresponding each observed DES and PS1 satellite, and top ticks indicate the corresponding present-day virial masses of the halos that host these systems.	88

4.7	Impact of modeling assumptions on the minimum subhalo mass inferred from the observed DES and PS1 satellite populations. The first three models match the number of subhalos to the number of confirmed DES and PS1 satellites, and the last two models populate subhalos with galaxies to fit the position-dependent MW satellite luminosity function and size distribution.	93
5.1	Distribution of satellite number in SAGA hosts. The black stars show the data from the SAGA Stage II sample of 36 complete hosts, and the error bars indicate incompleteness corrections as described in [316] (note that the incompleteness-corrected N_{sat} measurements are highly correlated). Dark blue (light blue) contours indicate the predicted 68% (95%) confidence intervals based on our simulation and galaxy–halo connection model. The MW (M31) is shown as a dashed (dashed–dotted) gray line.	101
5.2	SAGA Stage II satellite luminosity functions and incompleteness corrections (colored lines and bands) compared to predictions from a cosmological dark matter–only simulation populated with galaxies using the empirical satellite model in [336, 333], which has been fit to the MW satellite population. Dark blue lines indicate the mean prediction for each satellite population, and dark blue (light blue) contours indicate 68% (95%) confidence intervals, which include the effects of host galaxy–halo abundance-matching scatter, uncertainty in our galaxy–halo connection model (see Section 5.3.1), and projection effects.	102
5.3	Comparison of predicted and observed radial satellite distributions, normalized to the number of satellites within a projected distance of 300 kpc. The predicted mean and confidence interval is identical to Figure 5.2, and the observed radial distributions are computed using all satellites above our $M_{r,o} < -12.3$ absolute magnitude limit. Hence, no observed data are shown for systems that do not have any satellites with $M_{r,o} < -12.3$. We do not correct the SAGA radial distributions because the incompleteness model in [316] predicts that missing satellites should have the same distribution as confirmed ones.	103
5.4	Forecasts for the average number of satellites per MW analog detected by a SAGA Stage III-like survey. Dark (light) blue bands show the predicted 68% (95%) confidence intervals for the average number of satellites predicted for 100 mock SAGA observations as a function of one galaxy–halo connection model parameter. In each panel, the remaining parameters are fixed to their best-fit value from our MW satellite analysis in Chapter 4 [333] and vertical dashed lines show the corresponding 95% confidence intervals inferred from MW satellites.	107

6.1	Upper limits on the velocity-independent DM–proton scattering cross section as a function of DM particle mass. The blue shaded region is excluded by the population of classical and SDSS-discovered MW satellites with 95% confidence by our likelihood analysis, which marginalizes over relevant astrophysical uncertainties (Section 6.4). The dashed line shows analytic upper limits derived from the existence of the lowest-mass halos hosting satellites (Section 6.3). Green contours show cosmological constraints from the CMB [67, 198] and the Lyman- α forest [489]. Gray contours show experimental constraints from cosmic-ray scattering [88], the X-ray Quantum Calorimeter (XQC; [162]), and direct-detection experiments including CRESST-III [118], the CRESST 2017 surface run [27], and XENON1T [31], as interpreted by [159]. Limits from Galactic center gas clouds [61, 463] overlap with parts of the XQC, CMB, and Lyman- α contours for $10^{-3} \text{ GeV} \lesssim m_\chi \lesssim 100 \text{ GeV}$, and are omitted for clarity.	111
6.2	Ratio of the linear matter power spectrum in a DM–baryon scattering cosmology to that in CDM (solid lines), for a range of interaction cross sections (for 1 MeV DM particles). Dashed lines show the same quantity for WDM models with matching half-mode scales (denoted as black stars). Vertical lines indicate the critical scale discussed in Section 6.3.	115

7.1 Left panel: Transfer functions for the WDM (orange), IDM (blue), and FDM (magenta) models that are ruled out by our analysis at 95% confidence, corresponding to $m_{\text{WDM}} = 6.5 \text{ keV}$, $\sigma_0 = 8.8 \times 10^{-29} \text{ cm}^2$ (for DM particle mass $m_\chi = 100 \text{ MeV}$), and $m_\phi = 2.9 \times 10^{-21} \text{ eV}$, respectively. These constraints are marginalized over our MW satellite model and the properties of the MW system. Middle panel: SHMF suppression relative to CDM for each ruled-out non-CDM model. The vertical dashed line indicates the 95% confidence upper limit on the lowest-mass halo inferred to host MW satellite galaxies [333]. Note that the IDM SHMF is assumed to be identical to the WDM SHMF in our analysis, and is offset slightly for visual clarity. Right panel: Predicted MW satellite galaxy luminosity functions for each ruled-out non-CDM model compared to DES and PS1 observations, evaluated at the best-fit MW satellite model parameters from Ref. [333]. The shaded band illustrates the uncertainty of our WDM prediction due to the stochasticity of our galaxy-halo connection model and the limited number of simulations used in our analysis; the size of this uncertainty is very similar to that in CDM and the other alternative DM models shown. This panel is a simple one-dimensional representation of our MW satellite and DM model fit to the luminosity, size, and spatial distribution of satellites in the DES and PS1 survey footprints. The comparison of our CDM model to data is described in Ref. [333], and full posterior distributions for our non-CDM analyses are provided in Supplemental Material [188, 30, 219].

122

7.2 Exclusion regions for WDM and IDM models from our analysis of MW satellites observed with DES and PS1 (red) compared to previous constraints from classical and SDSS satellites [335] (blue) and other experimental results. Left panel: Constraints on the mass and mixing angle of resonantly produced sterile neutrino DM. These constraints are derived by finding mass and mixing angle combinations that suppress the linear matter power spectrum more strongly than the $m_{\text{WDM}} = 6.5 \text{ keV}$ thermal relic ruled out at 95% confidence by our analysis. The black point with error bars shows the sterile neutrino interpretation of the 3.5 keV x-ray line [82]. The dark gray region is ruled out by dwarf galaxy internal dynamics [83], and the gray contour shows x-ray constraints [225, 363, 134]. Solid black lines indicate regions of parameter space in which resonantly produced sterile neutrinos cannot constitute all of the DM in the neutrino minimal standard model [34, 414]. Right panel: Constraints on the interaction cross section and DM mass for velocity-independent DM-proton scattering. Green contours show cosmological limits from the CMB [67, 198] and the Lyman- α forest [489]. Light gray contours show experimental limits from the x-ray quantum calorimeter [314] and direct detection results as interpreted by Ref. [159].

128

7.3	Constraints on ultralight axion particle mass versus axion-photon coupling from our analysis of the MW satellite population (red). Limits from CMB polarization washout [170] and the Lyman- α forest [233] are shown in green, and haloscope limits are shown as gray vertical bands. Experimental constraints from the CAST experiment [32], the lack of a γ -ray signal from SN1987A [357], and the x-ray transparency of the intracluster medium [382] are shown in gray and do not require that the ultralight axion makes up all of the DM. The dashed lines [386] span canonical QCD axion models [263, 140].	131
8.1	Linear matter power spectra (left) and transfer functions (right) for self-interacting (magenta) and free-streaming (cyan) late-forming dark matter models, compared to cold dark matter (dashed black) and thermal relic warm dark matter (dashed red). Both LFDM models are shown with a transition redshift of $z_T = 1.5 \times 10^6$, corresponding to a comoving wave number of $k_T = 7h \text{ Mpc}^{-1}$. LFDM power spectra are suppressed relative to CDM at wave numbers greater than k_T , and they exhibit dark acoustic oscillations on even smaller scales, beginning at $\sim 6k_T$ ($\sim 2k_T$) for SI (FS) LFDM. The cutoff in the SI LFDM power spectrum is very similar to that in WDM, until the onset of DAOs.	142
8.2	Transfer functions for self-interacting (left) and free-streaming (right) late-forming dark matter models, compared to cold dark matter (dashed black) and thermal relic warm dark matter (dashed red). SI LFDM models are shown for a range of transition redshifts, with the highest transition redshift corresponding to the SI LFDM model that is ruled out by the abundance of Milky Way satellites at 95% confidence: $z_{T,SI} > 5.5 \times 10^6$. The light-blue FS LFDM model corresponds to the transition redshift that is conservatively ruled out by our analysis: $z_{T,FS} > 2.1 \times 10^6$. Vertical dashed lines show the comoving scale that approximately corresponds to the mass of the smallest halo inferred to host observed MW satellite galaxies, $3.2 \times 10^8 \text{ M}_\odot$ [333]. In the left panel, WDM transfer functions are slightly shifted horizontally for visual clarity. . .	144

8.3 Constraints on the transition redshift for self-interacting late-forming dark matter, versus the corresponding thermal relic warm dark matter mass based on the half-mode mass relation in Equation 8.13. Our Milky Way satellite constraint on $z_{T,SI}$ and the lower limit on the thermal relic WDM mass of 6.5 keV from which we derive this limit [332] are shown by the shaded purple region. Limits on the SI LFDM transition redshift from the cosmic microwave background (green), Sloan Digital Sky Survey galaxy clustering (dashed blue [407]), the high-redshift galaxy luminosity function (dot-dashed blue [115]) and the Lyman- α forest (dotted blue [407]) are shown as vertical lines. Vertical lines indicate constraints derived specifically for LFDM, and do not indicate other recent WDM constraints from small-scale structure probes. LFDM must transition to CDM between matter-radiation equality ($z \approx 3 \times 10^3$) and big bang nucleosynthesis ($z \approx 10^{10}$), which are schematically indicated by arrows. .	145
9.1 Momentum transfer cross sections for our SIDM model variants as a function of relative scattering velocity. Each model variant is labeled by w , the velocity scale above which the SIDM cross section falls off as v^{-4} . The velocity scale relevant for interactions among host halo particles and between host halo and subhalo particles is indicated by the ‘‘MW Host Halo’’ band. Shaded bands indicate characteristic velocities for DM particles within the subhalos expected to host classical and ultrafaint MW satellite galaxies.	153
9.2 SIDM effects on the host halo. Left panel: host halo density profiles for our SIDM model variants. Right panel: corresponding velocity dispersion profiles. As the SIDM cross section at the characteristic velocity scale of the MW host halo increases, the inner regions of the host become increasingly cored and thermalized. Note that the host halo density and velocity dispersion profiles are nearly indistinguishable in CDM and w_{10}	156

9.3 DM phase-space distributions for our zoom-in simulations of an MW-mass host halo in CDM and SIDM. The density of DM particles in bins of radial velocity and radial distance from the center of the host is plotted for CDM (top left) and for three of our SIDM model variants: w_{10} (top right), w_{100} (bottom left), and w_{500} (bottom right). These distributions qualitatively illustrate several of our main findings. For example, the host halo (labeled A) has a very similar phase space distribution in CDM and w_{10} , while subhalos in w_{10} (e.g., subhalo B) are somewhat less dense because of the large self-interaction cross section at low relative velocities in this case (see Figure 9.1). On the other hand, particles near the center of the host in w_{200} and w_{500} are preferentially scattered onto tangential orbits, and many of the low-mass subhalos that survive in CDM (e.g., subhalo C) are disrupted in these SIDM model variants owing to a combination of ram pressure stripping caused by self-interactions with the host and tidal stripping.	158
9.4 SIDM effects on subhalos before infall. Left panel: distributions of peak maximum circular velocity for surviving and disrupted subhalos in each SIDM model variant (unfilled histograms) vs. CDM (filled histogram). Right panel: cumulative distributions of the ratio of maximum circular velocity evaluated at the time of each subhalo's accretion onto the host divided by the peak maximum circular velocity along the main branch of the subhalo. Although subhalo assembly is statistically identical in CDM and SIDM, subhalos are mildly stripped by self-interactions prior to accretion onto the host halo in our SIDM simulations.	159
9.5 Properties of surviving and disrupted subhalos. Top left panel: distribution of the distance of closest approach to the host at first pericentric passage for surviving subhalos in each of our SIDM model variants (open histograms) vs. CDM (filled histogram). Top right panel: same as the top left panel, but for disrupted subhalos (surviving and disrupted subhalos are defined in Section 9.4.3). Bottom left panel: distribution of the maximum relative velocity with respect to the host halo evaluated along the orbit of each surviving subhalo. Bottom right panel: same as the bottom left panel, but for disrupted subhalos. Many subhalos that survive in our CDM simulation disrupt during early pericentric passages in our SIDM simulations, because ram pressure stripping caused by self-interactions with the host at large relative velocities makes subhalos more susceptible to tidal disruption.	164

9.6	Surviving subhalo populations. Left panel: peak velocity function of subhalos in our CDM simulation and in each of our SIDM model variants. Right panel: corresponding radial subhalo distributions in units of the host halo virial radius in each simulation. The abundance of surviving subhalos is reduced in SIDM, and the strength of this effect increases with w owing to more significant ram pressure stripping caused by self-interactions with the host. Subhalo disruption in our SIDM simulations is particularly severe in the inner regions of the host halo.	165
9.7	Orbital anisotropy profile of subhalos in our SIDM simulations. Surviving subhalos in SIDM model variants with larger values of w occupy tangentially biased orbits relative to CDM. This occurs because a combination of ram pressure and tidal stripping preferentially disrupts subhalos on radial orbits.	167
9.8	DM profiles of a matched set of surviving subhalos in our CDM and SIDM simulations. Density profiles defined by the initial set of bound particles are shown for the same subhalo with $V_{\text{peak}} \approx 40 \text{ km s}^{-1}$ and $z_{\text{acc}} \approx 1$ at the time of accretion onto the host (left panel) and at $z = 0$ (middle panel). Subhalos in model variants with larger values of w have lower-amplitude, flatter inner density profiles owing to self-interactions with the host halo. The light-blue line in the middle panel shows the density profile for a subhalo with similar V_{peak} and z_{acc} but with a <i>large</i> pericentric distance ($d_{\text{peri}} \approx 160 \text{ kpc}$). The discrepancy between this profile and that of the corresponding low-pericenter subhalo demonstrates that the impact of self-interactions on subhalo density profiles depends sensitively on their orbital properties. The right panel shows the corresponding circular velocity profile for each subhalo at $z = 0$. . .	169
10.1	Projections of the subhalo population in one of our MW-like simulations in CDM (top left) and SIDM (top right), and for subhalos above a survival probability threshold corresponding to the suppression of the peak velocity function in our WDM (bottom left) and disk disruption (bottom right) models applied to the same CDM simulation. Bold markers show subhalos with $V_{\text{peak}} > 20 \text{ km s}^{-1}$, light gray markers show subhalos below this V_{peak} threshold, and marker size is proportional to V_{peak} . In each panel, the black circle shows the virial radius of the MW host halo ($\sim 300 \text{ kpc}$), the diagonal line shows the plane defined by the current position of the MW (magenta star) and LMC (cyan star), and transparent cyan stars show the LMC at previous snapshots as it falls into the MW.	180

10.2 The effects of SIDM, WDM, and the Galactic disk on the velocity function and radial distribution of surviving subhalos in one of our MW-like simulations. Left panel: number of subhalos as a function of peak maximum circular velocity V_{peak} in our CDM simulations of MW-like systems (black dashed), SIDM simulations (blue), and predicted by applying our WDM (red) and disk disruption (green) models to our CDM simulations. Disruption due to the Galactic disk is approximately mass-independent, while SIDM and WDM preferentially disrupt less massive subhalos. Right panel: same as the left panel, but for the cumulative radial distribution of subhalos in units of the MW host halo virial radius. The suppression of subhalo abundance in WDM is not a strong function of Galactocentric radius, while SIDM and the disk preferentially disrupt subhalos at small radii. In the bottom panels, the dashed blue line shows the suppression of the subhalo population for the same SIDM model in a MW-mass system without a realistic LMC analog.	184
11.1 Posterior distribution of WDM half-mode mass versus baryonic disruption efficiency from our MW satellite analysis. $\mathcal{B} = 0$ corresponds to zero additional subhalo disruption relative to CDM, and larger values of \mathcal{B} correspond to more efficient subhalo disruption due to baryons. The color map shows the probability density normalized to its maximum value in this parameter space. Solid (dashed) white lines indicate 1σ (2σ) contours for a two-dimensional Gaussian distribution.	201
11.2 Projected SHMFs for MW-like host halos. Blue lines show the average results from the zoom-in simulations used in our MW satellite inference as a function of baryonic disruption efficiency \mathcal{B} ($\mathcal{B} = 0$ corresponds to CDM only and larger values of \mathcal{B} correspond to more efficient subhalo disruption due to baryons). Red lines show our analytic SHMF (Equation 11.5) using the host halo mass and redshift scaling predicted by Galacticus , evaluated at the average halo mass of our MW-like simulations with a slope of $\alpha = -1.92$. $\Sigma_{\text{sub},\text{MW}}$ is chosen such that the SHMF amplitude matches our simulations at the subhalo mass corresponding to the faintest observed MW satellites, \mathcal{M}_{min} (dashed vertical line). Dark (light) red contours show 68% (95%) confidence intervals from Galacticus for host halos with characteristics matched to our MW-like simulations. We impose the resolution cuts described in Appendix F.1 on the simulation and Galacticus results.	203

11.3 Left panel: posterior distribution of WDM half-mode mass versus projected subhalo number density at the strong lensing scale inferred from the MW satellite posterior, transformed according to the procedure in Section 11.6.2, with $q = 1$ (i.e., for equally efficient subhalo disruption due to baryons at the MW and strong lensing host halo mass and redshift scales). Right panel: posterior distribution in the same parameter space from the [195] strong lensing analysis. The vertical band labeled “ Σ_{sub} Prior” shows the range of Σ_{sub} inferred from the MW satellite posterior in our fiducial joint analysis (i.e., $0.015 \text{ kpc}^{-2} \leq \Sigma_{\text{sub}} \leq 0.03 \text{ kpc}^{-2}$). In both panels, color maps show the probability density normalized to its maximum value in each parameter space, and solid (dashed) white lines indicate 1σ (2σ) contours for a two-dimensional Gaussian distribution.	206
11.4 Joint marginal likelihood of WDM half-mode mass versus projected subhalo number density at the strong lensing scales from our combined MW satellite–strong lensing posterior, transformed according to the procedure in Section 11.6.2, with $q = 1$. The colormap shows the probability density normalized to its maximum value in this parameter space. Solid (dashed) white lines indicate 1σ (2σ) contours for a 2-dimensional Gaussian distribution.	208
11.5 Marginal distributions from our joint MW satellite–strong lensing likelihood (Figure 11.4) for projected subhalo number density at the strong lensing scale (left panel) and WDM half-mode mass (right panel), assuming equally efficient subhalo disruption due to baryons in the MW and strong lens systems ($q = 1$). The marginalized MW satellite posterior is shown in blue, the marginalized strong lensing posterior is shown in red, and results obtained from our probe combination and marginalized over the remaining dimension are shown in purple. In the left panel, the vertical band labeled “ Σ_{sub} Prior” shows the range of Σ_{sub} inferred from the MW satellite posterior in our fiducial joint analysis (i.e., $0.015 \text{ kpc}^{-2} \leq \Sigma_{\text{sub}} \leq 0.03 \text{ kpc}^{-2}$, slightly offset from the posteriors for visual clarity), and the dashed red line on the right panel shows the lensing half-mode mass posterior restricted to this range of Σ_{sub} values.	212

11.6 Left panel: the impact of systematics on the marginalized one-dimensional posterior distributions of projected subhalo number density at the strong lensing scale. The marginalized posterior distribution from our MW satellite analysis is shown in blue, the marginalized strong lensing posterior is shown in red, the dashed blue distributions indicate additional uncertainty in our MW satellite inference due to the mass of the MW halo, and the dotted-dashed green distribution illustrates the effects of systematic uncertainty in the differential efficiency of subhalo disruption due to baryons at the MW and strong lensing host halo scales. Right panel: joint marginal likelihood of WDM half-mode mass for our MW satellites plus strong lensing probe combination. Joint likelihoods are shown for equally efficient subhalo disruption in the MW and strong lens host halo mass and redshift regimes ($q = 1$, purple), twice as efficient disruption due to baryons in the MW relative to strong lens halos ($q = 0.5$, dotted-dashed green), and twice as efficient disruption in strong lens halos relative to the MW ($q = 2$, dashed green). The gray distribution shows the result of combining the fully marginalized one-dimensional M_{hm} posteriors derived from strong lensing and MW satellites.

215

12.1 Forecast for the minimum dark matter subhalo mass probed by LSST via observations of Milky Way satellites. The red band shows the 95% confidence interval from our MCMC fits to mock satellite populations as a function of the true peak subhalo mass necessary for galaxy formation. Note that we marginalize over the relevant nuisance parameters associated with the galaxy-halo connection—including the effects of baryons using the Chapter 2 model calibrated on subhalo disruption in hydrodynamic simulations [336]—in our sampling. We indicate the corresponding constraints on the warm dark matter mass assuming $M_{\text{hm}} = \mathcal{M}_{\text{min}}$

225

12.2 Forecasts for the galaxy formation and dark matter sensitivity of LSST Milky Way satellite observations using the satellite model from Chapters 4 and 7, which includes realistic Large Magellanic Cloud analog systems. The x -axis shows the 95% confidence upper limit on \mathcal{M}_{50} and the y -axis shows the 95% confidence upper limit on M_{hm} for warm dark matter from classical and SDSS-discovered satellites (blue star; Chapter 3 [336]), DES and Pan-STARRS1 satellites (red star; 4 [333]), and mock Milky Way satellite observations as they approach LSST sensitivity (blue-green triangles). The black dashed line shows the one-to-one relation and the gray dashed line shows the current lower limit on the warm dark matter mass from our joint Milky Way satellite and strong lensing analysis in Chapter 11. A cutoff in satellite abundances in Region I. must be astrophysical in origin, a cutoff in Region II. can be caused by either galaxy formation or dark matter physics, and the detection of halos below the lower limit of the \mathcal{M}_{50} axis at $\sim 5 \times 10^7 \text{ M}_\odot$ would provide evidence for the existence of completely dark halos.	228
A.1 Velocity functions for <code>m12i</code> (left) and <code>m12f</code> (right) predicted by the most probable realization of our random forest classifier when trained only on subhalos from <code>m12i</code> (orange) or <code>m12f</code> (green) with $V_{\text{peak}} > 10 \text{ km s}^{-1}$. Blue lines show 200 realizations of the prediction for our fiducial classifier, which is trained on subhalos from both hosts, and red lines show the FIRE results. While there is a difference between the total number of surviving subhalos predicted by classifiers trained only on <code>m12i</code> or <code>m12f</code> , the scatter about the most probable prediction for our fiducial classifier is small.	233
A.2 Velocity functions and radial distributions of subhalos in <code>m12i</code> (left) and <code>m12f</code> (right) predicted by the most probable realization of random forest classifiers trained on subhalos from both hydrodynamic simulations with $V_{\text{peak}} > 10 \text{ km s}^{-1}$. The classifiers use the features d_{peri} (blue); d_{peri} and a_{peri} (red); d_{peri} , a_{peri} , and a_{acc} (green); d_{peri} , a_{peri} , a_{acc} , and M_{acc} (orange); and d_{peri} , a_{peri} , a_{acc} , M_{acc} , and V_{acc} (cyan), corresponding to the rows of Table 2.1 and the columns of Figure 2.4. The solid red lines show the FIRE results.	234
A.3 Velocity functions (top) and radial distributions (bottom) for subhalos hosted by <code>m12i</code> (left) and <code>m12f</code> (right), predicted by the most probable realization of our random forest classifier trained on subhalos from <code>m12i</code> and <code>m12f</code> with $V_{\text{peak}} > 5 \text{ km s}^{-1}$ (blue). The FIRE (red), DISK (dot-dashed), and DMO (dashed) results are shown for comparison. We restrict these plots to subhalos within 300 kpc of their respective hosts at $z = 0$ and with $V_{\text{max}} > 5 \text{ km s}^{-1}$, where V_{max} is the maximum circular velocity evaluated at $z = 0$, to allow for a direct comparison with the results in [190].	235

B.1 Likelihood of observing n_i counts given \hat{N} mock observations $\hat{n}_{i,1}, \dots, \hat{n}_{i,\hat{N}}$ assuming that n_i and all \hat{n}_i are drawn from a Poisson distribution with rate parameter λ_i . We show results for $\hat{N} = 10$ (top row) and $\hat{N} = 1000$ (bottom row) mock observations given true rate parameters $\lambda_i = 1$ (left column), $\lambda_i = 5$ (middle column), and $\lambda_i = 10$ (right column), computed using (i) the Poisson likelihood given the true rate parameter (gray points); (ii) the likelihood used in our analysis (Equation 3.13), which marginalizes over λ_i given all \hat{n}_i simultaneously (blue triangles); and (iii) the likelihood used in [241], which averages the likelihoods obtained from multiple mock observations (red circles). Our likelihood converges to the underlying Poisson distribution in the limit of many mock observations, while the averaged version does not. Note that we rescale each version of $P(n_i \{\hat{n}_i\})$ by its maximum value in every panel.	237
C.1 Radial distributions derived from our fit to the DES and PS1 satellite populations. Our fiducial eight-parameter galaxy occupation fraction model is shown in blue. Dark (light) blue bands correspond to 68% (95%) confidence intervals, dashed red lines show the 68% confidence interval for a model using host halos without LMC analogs (“No LMC”), and black lines show the observed radial distributions. Dotted-dashed blue lines show the 68% confidence interval for a model with a radial scaling parameter of $\chi = 0.5$	243
C.2 Size distributions derived by fitting to the DES and PS1 satellite populations. Our fiducial eight-parameter galaxy occupation fraction model is shown in blue. Dark (light) blue bands correspond to 68% (95%) confidence intervals, dashed red lines show the 68% confidence interval for a model with a concentration-dependent galaxy–halo size relation, and dotted-dashed blue lines show the 68% confidence interval for a model with an extreme dependence of satellite size on tidal stripping.	243
C.3 Posterior distribution from our fit to the kinematically confirmed DES and PS1 satellite populations. Dark (light) shaded contours represent 68% (95%) confidence intervals. Shaded areas in the marginal distributions and parameter summaries correspond to 68% confidence intervals. Note that σ_M , σ_{gal} , and $\sigma_{\log R}$ are reported in dex, \mathcal{M}_{50} is reported as $\log(\mathcal{M}_{50}/M_\odot)$, and \mathcal{A} is reported in pc. Note that σ_{gal} is not constrained at 68% confidence in this fit.	250
D.1 Posterior distribution from our fit to the DES and PS1 satellite populations for thermal relic WDM. Dark (light) shaded contours represent 68% (95%) confidence intervals. Shaded areas in the marginal distributions and parameter summaries correspond to 68% confidence intervals. Note that σ_M , σ_{gal} , and $\sigma_{\log R}$ are reported in dex, \mathcal{M}_{50} and M_{hm} are reported as $\log(\mathcal{M}_{50}/M_\odot)$ and $\log(M_{\text{hm}}/M_\odot)$, \mathcal{A} is reported in pc, and α , \mathcal{B} , and n are dimensionless.	255

D.2	Same as Figure D.1, but for our FDM fit.	256
E.1	Simulation resolution study. Left panel: peak velocity functions of subhalos for our high-resolution CDM and w_{500} resimulations (labeled “16K”), compared to those from our fiducial simulations (labeled “8K”). Right panel: corresponding radial subhalo distributions in units of the host halo virial radius in each simulation.	259
E.2	Subhalo resolution threshold study. Left panel: total number of surviving subhalos in each SIDM model variant above a fixed V_{peak} threshold divided by the corresponding number of surviving subhalos in our CDM simulation. Middle panel: same as the left panel, but for the number of disrupted subhalos. Right panel: same as the previous panels, but for the number of surviving-plus-disrupted subhalos.	260
F.1	Left panel: projected subhalo mass function for MW-like host halos as a function of present-day subhalo virial mass. Blue lines show results from the zoom-in simulations used in our MW satellite inference for several values of the baryonic disruption efficiency parameter \mathcal{B} (models with more efficient subhalo disruption are shown in darker colors). Right panel: average radial subhalo distribution in units of the host halo virial radius for our MW-like simulations (blue). Dashed vertical lines approximately mark the radial range of observed MW satellite galaxies used in our analysis. In both panels, dark (light) red contours show 68% (95%) confidence intervals from Galacticus for a sample of halos with characteristics matched to our MW-like simulations (see Appendix F.2 for details). To calculate the Galacticus radial distributions, we only consider halos with $M_{\text{peak}} > 10^8 M_{\odot}$ in addition to the V_{peak} and V_{max} cuts described in Appendix F.1 to facilitate a direct comparison to our simulation results.	263
F.2	Left panel: projected subhalo mass functions versus peak subhalo virial mass. Lines correspond to individual zoom-in simulations from the [315] suite of MW-mass host halos and are colored according to their projected subhalo number density $\Sigma_{\text{sub,MW}}$; the two thickest lines correspond to the MW-like simulations used in our analysis. The gray band indicates the range of SHMFs from these simulations, and dark (light) red contours show 68% (95%) confidence intervals from Galacticus for a sample of halos with characteristics matched to our MW-like simulations. Right panel: relation between host halo mass, concentration, and $\Sigma_{\text{sub,MW}}$ for the same suite of zoom-in simulations. Stars show the two MW-like simulations used in our analysis, which include realistic LMC analog systems, triangles show simulations from this suite that have an LMC analog (i.e., a subhalo with $V_{\text{max}} > 55 \text{ km s}^{-1}$) anywhere within their virial radius, and circles show simulations that do not have an LMC analog. Colors indicate fractional differences relative to the average value of $\Sigma_{\text{sub,MW}}$ from the two MW-like simulations.	264

Chapter 1

Introduction

Λ CDM cosmology rests on assumptions about its fundamental constituents: dark energy is assumed to behave as a cosmological constant (Λ), filling space with a constant energy density that drives the late-time accelerated expansion of the Universe, and cold dark matter (CDM) is assumed to move slowly and interact feebly on cosmological time and length scales, providing the backbone for the gravitational formation and evolution of cosmic structure. A myriad of diverse data are consistent with these simple hypotheses, indicating that dark energy, cold dark matter, and Standard Model particles respectively constitute about 70%, 25%, and 5% of the Universe’s contents today [15].

New generations of cosmological and astrophysical observations promise to test the Λ CDM paradigm in distinct ways. Surveys of the distribution of large-scale structure traced by galaxies are enabling percent-level measurements of cosmological parameters in the late-time Universe, including their potential evolution (e.g., [130]). Meanwhile, probes of cosmic expansion at both early times and the present day test for consistency between the physics governing these epochs, pointing towards currently unresolved tensions (e.g., [458] and references therein).

One of the most compelling questions raised by the success of Λ CDM concerns the nature of dark matter. Modified theories of gravity tailored to explain dynamical observations on galactic scales fail to simultaneously match the detailed distribution of cosmic structure encoded in the matter power spectrum and reflected in the cosmic microwave background (CMB). Thus, it is usually assumed that CDM is composed of particles, but no known particle is cold or collisionless enough to account for (even a small fraction of) the dark matter. Several decades of theoretical and experimental efforts have been driven by the hope that discovering the microphysical properties of dark matter will lead to a discovery of physics beyond the Standard Model.

Galaxies form and grow in the gravitational potential wells provided by dark matter, tracing its distribution down to the smallest cosmic scales. The entire spectrum of galaxies therefore contains crucial information about dark matter physics that must be understood in concert with the large-scale observations described above. In turn, the formation and evolution of galaxies can in principle

be influenced by non-gravitational dark matter physics in addition to astrophysical processes, and the two must be addressed together in a self-consistent model.

This thesis studies the interface of galaxy formation and dark matter physics with an aim to simultaneously derive insights into both. In particular, by focusing on the faintest observable galaxies and their connection to the microphysical properties of dark matter, it explores the boundaries of galaxy formation theory where any imprints of new dark matter physics are expected to be most noticeable. This chapter provides a high-level overview of the underlying physics, including cosmological assumptions in Section 1.1, the hierarchical buildup of cosmic structure that scaffolds galaxy formation in Section 1.2, and the processes by which gas cools and forms stars in dark matter halos in Section 1.3. Section 1.4 gives an observational overview of Milky Way dwarf satellite galaxies, which inform many of the comparisons to data in this thesis, and Section 1.5 introduces the connection between dwarf galaxies, dark matter halos, and microphysical dark matter properties.

1.1 Preliminaries

This thesis operates in a flat, homogeneous, isotropic, and expanding spacetime, referred to as a Friedmann-Lemaître-Robertson-Walker (FLRW) cosmology. The metric for this spacetime,

$$ds^2 = dt^2 - a^2(t)d\mathbf{x}^2, \quad (1.1)$$

is characterized by the scale factor $a(t)$, which describes the size of the Universe relative to today, with $a(t_{\text{today}}) \equiv 1$. The redshift z , which describes the ratio by which a photon's wavelength is stretched as it travels through the expanding Universe, is related to the scale factor via

$$z = \frac{1}{a} - 1. \quad (1.2)$$

The evolution of the scale factor is governed by the Friedmann equation,

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho, \quad (1.3)$$

where the energy density can be written in units of the critical density today, ρ_c , as

$$\frac{\rho}{\rho_c} = \Omega_\Lambda + \Omega_m a^{-2} + \Omega_{\text{rad}} a^{-4}. \quad (1.4)$$

In Equation 1.4, Ω_Λ , Ω_m , and Ω_{rad} respectively represent the fraction of the Universe's energy density currently stored in a cosmological constant ($\Omega_\Lambda \approx 0.7$), nonrelativistic matter ($\Omega_m \approx 0.3$), and relativistic species ($\Omega_{\text{rad}} \approx 10^{-4}$), which sum to unity in a flat FLRW spacetime. The matter component is further decomposed into CDM and baryonic species, respectively denoted Ω_c and Ω_b .

Equation 1.4 implies that the contents of the Universe observed today were partitioned differently at early times. Extrapolating the energy densities backwards places the redshift of matter–radiation equality at $z \sim 3000$. Before this epoch—i.e., in the first 300,000 years after the Big Bang—radiation pressure slowed the gravitational growth of matter perturbations. Structure grew efficiently during a matter-dominated epoch over the next ~ 10 Gyr, and the remaining portion of the Universe’s 13.7 Gyr history was dominated by dark energy, which accelerates the late-time expansion rate.

The matter perturbations imprinted on the CMB provide clear evidence for the fundamental role that dark matter plays in structure formation. CMB temperature anisotropies, which trace overdensities of baryonic matter, are observed at the level of $(\delta T/T) \sim 10^{-5}$. However, matter perturbations grow linearly with the scale factor during the matter-dominated epoch. Requiring that overdensities become nonlinear by today in order to form galaxies therefore requires underlying fluctuations of $(\delta\rho/\rho) \sim 10^{-3}$ at the time of the CMB. These overdensities must be sourced by non-baryonic dark matter. Furthermore, they are consistent with being drawn from a homogeneous and isotropic Gaussian random field laid down at very early times, for example by a phase of rapid inflationary expansion.

1.2 Dark Matter and Hierarchical Structure Formation

The arguments above imply that dark matter scaffolds structure formation. But how does dark matter structure itself form and grow? Small initial overdensities attract matter and evolve into self-gravitating systems referred to as *dark matter halos*. Halos are roughly spherical distributions of dark matter with centrally concentrated density profiles; despite being approximately virialized, they are not in equilibrium because they continuously accrete matter from their surroundings.

Because gravitational interactions are scale free, dark matter halos assemble in a roughly self-similar fashion down to the mass scales impacted by dark matter microphysics. Unlike the CDM limit, realistic particle dark matter candidates have a nonzero free-streaming length, defined as the typical distance a dark matter particle travels before matter perturbations begin to grow efficiently after matter–radiation equality [416]:

$$\lambda_{\text{fs}} = \int_0^{t_{\text{MRE}}} \frac{v(t)}{a(t)} dt. \quad (1.5)$$

Halo formation is suppressed on mass scales smaller than those corresponding to λ_{fs} , moderately suppressed up to about one decade in mass above this scale (e.g., [415]), and asymptotes to CDM predictions at even higher masses. In linear theory, the mass M of a collapsed halo is sourced by

density perturbations on a scale λ with the relation

$$M = \frac{4\pi}{3} \rho_m \left(\frac{\lambda}{2} \right)^3 = \frac{4\pi}{3} \rho_m \left(\frac{\pi}{k} \right)^3, \quad (1.6)$$

where the cosmological wavenumber $k \equiv \lambda/2\pi$.

The free-streaming scale is sensitive to a variety of dark matter physics. The most commonly considered phenomenological model is that of thermal relic warm dark matter, in which dark matter is produced with a Maxwell-Boltzmann distribution when the Universe drops below a temperature corresponding to the particle mass m_{WDM} . Chapters 6–8 of this thesis describe how velocity-independent dark matter–baryon interactions and dark matter that forms relatively late in the radiation-dominated epoch mimic the free-streaming effects of thermal relic warm dark matter.

Halo formation is hierarchical well above the free-streaming scale: small halos form first and merge together to form larger halos, including populations of subhalos (i.e., halos embedded within larger halos). The number density of halos per mass interval roughly obeys a power law [494]

$$\frac{dn}{dm} \propto m^{-\alpha}, \quad (1.7)$$

with $\alpha = -1.9$, and internal halo structure is reasonably well-described by the double power-law Navarro-Frenk-White (NFW; [338]) density profile. The weak scale-dependence in the NFW profile, which sets the transition of its power-law slope and thus the concentration of dark matter halos, reflects characteristic halo formation times (e.g., [473]).

The mass function, spatial distribution, and internal properties of subhalos are affected by non-linear gravitational and astrophysical processes that are relevant for the satellite galaxy modeling in this thesis. Chapters 3–4 address some of this physics in the CDM context, and Chapters 9–10 describe the effects of dark matter self-interactions on subhalos. Subhalos within the Milky Way are crucial for understanding both the threshold of galaxy formation and the small-scale behavior of dark matter because they host the faintest observable galaxies.

1.3 Galaxy Formation in Dark Matter Halos

Before the epoch of the CMB, baryons were distributed in an opaque plasma of protons and electrons. As the Universe cools below the binding energy of hydrogen, these protons and electrons combine to form a diffuse, transparent neutral hydrogen gas. This gas falls into the gravitational potential wells provided by dark matter halos, gaining kinetic energy and shock-heating to the virial temperature

$$T_{\text{vir}} \approx 10^5 \text{ K} \left(\frac{V_{\text{vir}}}{100 \text{ km s}^{-1}} \right)^2, \quad (1.8)$$

where V_{vir} is the virial velocity, which is directly related to halo mass. Gas heated to sufficiently high temperatures $T_{\text{vir}} \gtrsim 10^4$ K efficiently cools through two-body radiative processes and sinks to the center of its halo. This cool, dense gas can then begin forming stars, initiating the life cycle of galaxy formation and evolution.

A rich set of astrophysical processes complicate this simple picture to shape the evolving relation between galaxy and halo properties and assembly histories over cosmic time. The efficiency of star formation peaks for systems with masses comparable to the Milky Way, corresponding to stellar masses of $M_* \approx 10^{10} M_\odot$ and halo masses of $M_{\text{halo}} \approx 10^{12} M_\odot$ [45]. Note that this percent-level conversion of mass to stars is significantly lower than the cosmic baryon fraction of $\sim 15\%$. Star formation efficiency drops sharply both above and below the Milky Way mass, largely due to feedback from active galactic nuclei at higher masses and supernovae at lower masses (e.g., see [474] and references therein).

Combined with Equation 1.8, the requirement that gas is shock-heated to $T_{\text{vir}} \gtrsim 10^4$ K implies that sufficiently low-mass halos are not expected to form stars. Precisely predicting and measuring this “galaxy formation threshold” represents an active area of study, and Chapters 3–4 in this thesis derive new upper limits of $\sim 10^8 M_\odot$ on the masses of the smallest halos that must form at least a few hundred stars in order to account for the abundance of faint galaxies orbiting the Milky Way. Late-time star formation in dwarf galaxies with masses just above the galaxy formation threshold is inhibited by the photoionizing background emitted by other galaxies at early times, further suppressing the star formation efficiencies in this regime (e.g., [98, 349, 56]).

1.4 Dwarf Galaxies around the Milky Way

The physics that dictates dwarf galaxy formation implies that the luminous counterparts of low-mass dark matter halos are extremely challenging to detect. In particular, the faintest dwarf galaxies can only be detected nearby as satellite galaxies orbiting within the Milky Way’s dark matter halo. Among these satellite galaxies, the Large and Small Magellanic Clouds have been known and recorded by indigenous peoples in the Southern Hemisphere throughout human history. In the 1900s, astronomers including Shapley [420] discovered 9 more “classical” satellites by visually identifying sub-degree scale overdensities on photographic plates. Classical satellites range in luminosity from the Large Magellanic Cloud, with $M_* \approx 10^9 M_\odot$, down to systems as faint as $M_* \approx 10^6 M_\odot$. Beyond the extent of the Milky Way’s dark matter halo (i.e., at distances beyond ~ 300 kpc), only a few comparably faint galaxies are known, mostly within the Local Group as satellites of Andromeda (M31; [319]) and other Local Volume central galaxies (e.g., [104]).

The Sloan Digital Sky Survey (SDSS; [21]) revolutionized the search for Milky Way satellites in the 2000s with systematic digital imaging of $\sim 1/3$ of the Northern sky. Automated dwarf galaxy searches run on SDSS data led to the discovery of 16 new Milky Way satellites. Many of

these galaxies are significantly fainter than the classical satellites, including systems like Segue I with $M_* \approx 300 \text{ M}_\odot$ [192, 426], which remains the faintest spectroscopically confirmed galaxy ever observed. Galaxies with $M_* \lesssim 10^5$ are referred to as “ultra-faint” dwarfs and exhibit uniformly ancient stellar populations (e.g., [425] and references therein). These systems are distinct from globular clusters of comparable luminosities because they are highly dark matter dominated, with mass-to-light ratios much larger than unity and as high as ~ 1000 inferred from stellar velocity dispersion measurements.

The last decade has seen a renaissance of Milky Way satellite galaxy discovery and science driven by new generations of photometric surveys including the Dark Energy Survey (DES; [130]) and Pan-STARRS1 (PS1; [108]). Combined with other community-led efforts using the Dark Energy Camera, deep surveys of small areas like the Hyper Suprime-Cam Subaru Strategic Program, and novel search techniques based on RR Lyrae clustering, our census of Milky Way satellites has doubled again, reaching ~ 60 confirmed and candidate dwarf galaxies [149]. In addition to the faintest known galaxy, this population contains the lowest surface brightness galaxy ever observed (Antlia 2; [444]).

1.5 Connecting Dark Matter to Dwarf Galaxies

The remarkable nature of the Milky Way’s dwarf satellite galaxy population begs for a theoretical explanation. Initial attempts to compare Λ CDM predictions to the classical satellite population (e.g., [266, 325]) seemed to severely overpredict the abundance of observable dwarf galaxies, termed the “missing satellites problem.” This apparent discrepancy inspired a plethora of solutions based in both astrophysics and modifications to the CDM paradigm (e.g., see [96] and references therein). For example, [98] noted that the observed Milky Way satellite luminosity function can be reconciled with theoretical predictions if reionization significantly suppresses star formation in dwarf galaxies. Meanwhile, many dark matter models that suppress the formation of dwarf galaxy halos were also proposed (e.g., using warm dark matter particles of $m_{\text{WDM}} \sim 1 \text{ keV}$; [200]), and such modifications continue to drive dark matter theory development today (e.g., [228]).

Although this thesis does not address the physics of the very early Universe, it is worth noting that the abundance of dwarf galaxy halos is also sensitive to the *primordial* matter power spectrum and thus to non-standard models of inflation or early matter domination (e.g., [252, 189, 322, 435]). Other potential problems concerning the inferred dark matter density profiles of dwarf galaxies (e.g., [126])—and particularly the Milky Way’s brightest satellites (e.g., [84, 242])—have inspired proposals for modifications to dark matter physics that deplete halos’ central densities at late times, including self-interaction cross sections of $\sim 1 \text{ cm}^2 \text{ g}^{-1}$ (e.g., [432]). In parallel, improved feedback prescriptions and resolution in hydrodynamic simulations subsequently revealed that astrophysical processes can also create dark matter cores (e.g., [369, 378] and references therein). A careful treatment of dwarf galaxy halo profiles and their connection to observations is beyond the scope of

this thesis, although the self-interacting dark matter simulations of Milky Way-like systems presented in Chapters 9–10 inform this question.

Returning to the missing satellites problem, Equation 1.6 implies that inhibiting the formation of dwarf galaxy halos (roughly taken to be systems with $M_{\text{halo}} \lesssim 10^{10} M_{\odot}$, which will be made precise in subsequent chapters) requires suppressing density perturbations on scales $\lambda \lesssim 1 \text{ Mpc}$. These scales correspond to cosmological modes with $k \gtrsim 10 h \text{ Mpc}^{-1}$ and enter the horizon during the radiation-dominated epoch at $z \gtrsim 10^5$, when the Universe is less than a few months old. Because observations on larger scales are consistent with CDM predictions, this sets the free-streaming scale for particle dark matter models that affect dwarf galaxy halos. These simple arguments foreshadow the upper limits derived on the free-streaming scale for warm dark matter-like models in Chapters 6, 7, and 8 based on the consistency of the Milky Way satellite galaxy population with CDM predictions demonstrated in Chapters 3 and 4.

Early comparisons between cosmological simulations and Milky Way satellite galaxy observations were hampered by resolution effects including artificial disruption and fragmentation (e.g., [467, 454]). Recent generations of N -body and hydrodynamic cosmological simulations using the “zoom-in” technique ([258, 347]) have only recently reached sufficient resolution to accurately resolve the halos expected to host dwarf galaxies as subhalos of a Milky Way-mass host. These simulations reveal that key physics was missing in the early comparisons to data underlying the missing satellites problem. In particular, central disks akin to the Milky Way galaxy tidally disrupt a large fraction of orbiting subhalos, reducing the expected number of nearby dwarf galaxies. Chapter 2 studies this phenomenon in detail using state-of-the-art hydrodynamic simulations.

Of equal importance, observational selection effects that dictate the probability dwarf galaxies can be detected as a function of their position on the sky and structural parameters were not developed in detail until the time of SDSS searches (e.g., [271, 465]). Chapter 3 uses these SDSS selection functions to compare satellite population predictions to Milky Way satellite galaxy data, and Chapter 4 uses new selection functions derived from DES and PS1 data [149], which explicitly incorporate the effects of satellite galaxy size and sky position on detectability for the first time.

In one sense, this thesis demonstrates that careful treatments of the galaxy–halo connection, observational selection effects, and subhalo disruption are sufficient to explain the missing satellites problem. More importantly, by developing a comprehensive modeling framework that unifies dark matter and galaxy formation physics, this work addresses the following questions:

1. *What is the quantitative relationship between dwarf galaxy and dark matter halo properties?*
2. *What is the mass of the halo that hosts the faintest observed galaxy, and what is the mass of the faintest galaxy that can form in a dark matter halo?*
3. *Does every dark matter halo host a galaxy?*
4. *What do the faintest galaxies teach us about the microphysical properties of dark matter?*

Chapter 2

Modeling the Impact of Baryons on Subhalo Populations

Abstract

We identify subhalos in dark matter-only (DMO) zoom-in simulations that are likely to be disrupted due to baryonic effects by using a random forest classifier trained on two hydrodynamic simulations of Milky Way (MW)-mass host halos from the Latte suite of the Feedback in Realistic Environments (FIRE) project. We train our classifier using five properties of each disrupted and surviving subhalo: pericentric distance and scale factor at first pericentric passage after accretion, and scale factor, virial mass, and maximum circular velocity at accretion. Our five-property classifier identifies disrupted subhalos in the FIRE simulations with an 85% out-of-bag classification score. We predict surviving subhalo populations in DMO simulations of the FIRE host halos, finding excellent agreement with the hydrodynamic results; in particular, our classifier outperforms DMO zoom-in simulations that include the gravitational potential of the central galactic disk in each hydrodynamic simulation, indicating that it captures both the dynamical effects of a central disk and additional baryonic physics. We also predict surviving subhalo populations for a suite of DMO zoom-in simulations of MW-mass host halos, finding that baryons impact each system consistently and that the predicted amount of subhalo disruption is larger than the host-to-host scatter among the subhalo populations. Although the small size and specific baryonic physics prescription of our training set limits the generality of our results, our work suggests that machine-learning classification algorithms trained on hydrodynamic zoom-in simulations can efficiently predict realistic subhalo populations.

2.1 Paper Status and External Contributions

This chapter is published in slightly modified form in the *Astrophysical Journal*, Volume 859, Issue 2, p.129-146 with the title, “Modeling the Impact of Baryons on Subhalo Populations with Machine Learning,” on which I am the corresponding author. It is the result of a collaboration with Yao-Yuan Mao and Risa Wechsler with contributions from co-authors Shea Garrison-Kimmel and Andrew Wetzel. Shea and Andrew published and provided us with the simulations used in this work. Yao provided initial analysis tools and he and Risa were involved in the development of the random forest algorithm and its applications. In addition, Yao and Risa made editorial contributions to the text. This work matured during the 2018 Kavli Institute for Theoretical Physics program “The Galaxy–Halo Connection Across Cosmic Time,” during which conference participants made helpful suggestions.

2.2 Introduction

The Λ CDM cosmological model provides a remarkably successful framework in which the observed large-scale distribution of galaxies can be understood in terms of the underlying distribution of dark matter halos. However, there are several outstanding “small-scale” problems associated with Λ CDM cosmology (see [97] and [129] for recent reviews). For example, dark matter-only (DMO) simulations predict large numbers of low-mass subhalos that contribute to an ever-rising low-mass end of the subhalo mass function. If these low-mass subhalos exist and host galaxies, we should observe many more dwarf satellites than currently detected around the Milky Way (MW) or the Andromeda Galaxy (M31); this is often dubbed the missing-satellites problem [266, 325]. Meanwhile, the “too big to fail” (TBTF; [84]) problem arises because the number of subhalos with high maximum circular velocities ($V_{\max} \gtrsim 15 \text{ km s}^{-1}$) found in DMO simulations of MW-mass systems substantially exceeds the number of such subhalos inferred to exist around the MW and M31. Equivalently, observational estimates for the masses of the subhalos that host the dwarf satellites of the MW and M31 fall below the masses predicted by DMO simulations [85].

While these small-scale problems present challenges to the Λ CDM paradigm, a number of promising astrophysical solutions to each problem have been proposed. For example, it is now understood that cosmic reionization suppresses star formation in low-mass subhalos, while supernova (SN) feedback can suppress star formation in more massive subhalos, potentially resolving the missing-satellites problem [98, 429]. Proposed solutions to the TBTF problem build on these ideas by invoking stellar feedback to soften central density cusps and deplete subhalos of dark matter [201, 369], along with enhanced subhalo disruption via tidal stripping or disk shocking, to destroy many of the high- V_{\max} subhalos found in DMO simulations. Several authors have suggested that these mechanisms can yield subhalo populations in agreement with those inferred observationally for the MW and M31 [502, 91, 90, 478, 411].

Indeed, recent high-resolution hydrodynamic simulations that self-consistently resolve star formation, stellar feedback, and the formation of central galactic disks indicate that the missing-satellites and TBTF problems can largely be mitigated for the subhalo populations of MW-mass host halos. For example, [478] and [190] (hereafter GK17) studied the subhalo populations of two MW-mass host halos from the Latte simulation suite of the Feedback in Realistic Environments (FIRE) project [223] using the ‘zoom-in’ simulation technique [258, 347]. These authors found that the total number of subhalos in each simulation is reduced by about a factor of two relative to corresponding DMO simulations with identical initial conditions, and they also found significantly fewer subhalos with high circular velocities in the hydrodynamic runs. Moreover, the subhalo populations in both of these systems are consistent with a variety of observational probes for the MW and M31, which suggests that the missing-satellites and TBTF problems can be resolved in these particular simulations [478]. [501] reached similar conclusions by comparing hydrodynamic zoom-in simulations of MW-mass host halos from the Aquarius Project [433] to DMO simulations of the same hosts.

These results rely on a limited number of high-resolution simulations of MW-mass host halos; unfortunately, studying a large, diverse sample of subhalo populations in hydrodynamic zoom-in simulations is currently infeasible. While many authors have justifiably focused on the subhalo populations of MW-mass host halos, since these are particularly relevant to the original TBTF problem, it is important to assess whether analogous TBTF problems arise for the subhalo populations of more massive host halos. In addition, understanding whether the TBTF problem is consistently mitigated in a range of simulations with different baryonic physics implementations is necessary in order to make robust conclusions. Quantifying the impact of baryonic physics on subhalo populations more generally will be important in order to interpret results from large-scale surveys, including the Dark Energy Spectroscopic Instrument [131] and the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST; [301]), and from targeted searches for satellites of MW-like galaxies outside the Local Group such as the Satellites Around Galactic Analogs Survey (SAGA; [193]).

Thus, models that can incorporate a variety of hydrodynamic simulations to predict realistic subhalo populations directly from DMO simulations are worth exploring. As a first step toward such a model, we present a machine-learning classification algorithm to identify subhalos in DMO zoom-in simulations of MW-mass host halos that are likely to be disrupted due to baryonic effects in hydrodynamic resimulations. In particular, we train a random forest classifier on disrupted and surviving subhalos from the FIRE zoom-in simulations presented in GK17, and we use the classifier to predict surviving subhalo populations in DMO zoom-in simulations. Our aim is to explore whether this algorithm can capture the effects of baryons in existing hydrodynamic simulations and how the particular baryonic physics in these simulations alters subhalo populations in independent DMO simulations. Rather than providing a detailed comparison of different classification algorithms, we show that a simple random forest classifier predicts subhalo populations in excellent agreement

with hydrodynamic results when applied to DMO simulations of the FIRE host halos. This technique is efficient, since a trained classifier can immediately predict surviving subhalo populations from relatively inexpensive DMO simulations. We view classification as a promising technique for predicting subhalo disruption because classifiers will become more robust as the number of high-resolution hydrodynamic simulations to train on increases. In particular, classification algorithms can be trained on a variety of zoom-in simulations to capture the impact of baryons on subhalo populations for a range of host halo masses, central galaxy types, formation histories, and subgrid physics prescriptions.

In addition to the practical utility of our results for predicting realistic subhalo populations, our work provides insights into subhalo disruption in hydrodynamic simulations and relates to the small-scale challenges described above. For example, our random forest classifier determines how strongly various subhalo properties correlate with disruption, which indicates the importance of different disruption mechanisms, including tidal effects and stellar feedback, given the specific baryonic physics prescription in these simulations. To explore the relative importance of these disruption mechanisms, we compare the surviving subhalo populations that we predict from DMO simulations of the FIRE host halos to the DMO-plus-disk simulations presented in GK17, which are designed to capture the dynamical effects of the central galactic disk that develops in each hydrodynamic simulation. In particular, by performing DMO zoom-in simulations of two systems with analytic disk potentials tuned to match the galactic disks that develop in the corresponding hydrodynamic simulations, GK17 found subhalo populations in good agreement with the hydrodynamic results, particularly in the innermost regions ($r \lesssim 100$ kpc). This result suggests that, for MW-mass halos with a central galactic disk, the tidal effects of the disk are largely responsible for disrupting both the low- V_{\max} subhalos relevant to the missing-satellites problem and the high- V_{\max} subhalos relevant to the TBTF problem. Our machine-learning predictions are consistent with the DMO-plus-disk simulations at low V_{\max} , but we find enhanced disruption for subhalos with $V_{\max} \gtrsim 15$ km s⁻¹ and our results match the FIRE simulations more closely for such subhalos. Interestingly, several authors have suggested that baryonic physics efficiently creates cored subhalo density profiles in this regime [109, 442, 176]. We therefore argue that baryonic effects within subhalos, such as stellar feedback, can help to relieve the tension between the subhalo populations predicted by DMO simulations and those inferred from observations of the Local Group.

Our work also has broader implications for studying the galaxy–halo connection. For example, by using our classifier to predict surviving subhalo populations for the suite of DMO zoom-in simulations of MW-mass host halos from [315], we find that the average amount of subhalo disruption due to baryonic effects is larger than the host-to-host scatter among various subhalo populations. Thus, models that utilize subhalo statistics from these simulations should account for enhanced subhalo disruption when marginalizing over the effects of baryonic physics. Several semianalytic models (e.g., [302, 303]) incorporate subhalo velocity functions predicted by DMO zoom-in simulations of

MW-mass host halos in order to constrain the properties of the MW satellite galaxies and their host halos, and it is plausible that the physical insights provided by these models could change when more realistic subhalo populations are used as input.

This paper is organized as follows. In Section 2.3, we describe the FIRE simulations that we use to train our random forest classifier, as well as the DMO and DMO-plus-disk simulations presented in GK17 to which we compare our results. In Section 2.4, we describe our training and cross-validation methods, and we test our classifier by predicting disrupted subhalos in two FIRE zoom-in simulations. We present our main results in Section 2.5. In Section 2.5.1, we predict surviving subhalo populations in DMO simulations of the FIRE host halos, and we present velocity functions and radial distributions for our predicted subhalo populations; in Section 2.5.2, we predict surviving subhalo populations for the suite of DMO zoom-in simulations from [315], and we discuss the implications for satellite searches. We address avenues for future work and summarize our conclusions in Section 2.6.

We adopt cosmological parameters consistent with each simulation that we analyze. In particular, we use $h = 0.702$, $\Omega_m = 0.272$, $\Omega_b = 0.0455$, and $\Omega_\Lambda = 0.728$ for our analysis of the FIRE simulations and $h = 0.7$, $\Omega_m = 0.286$, $\Omega_b = 0.047$, and $\Omega_\Lambda = 0.714$ for our analysis of the MW zoom-in simulation suite. Note that we express distances in physical kpc and velocities in km s^{-1} .

2.3 Simulation Data

We train our random forest classifier using subhalos from the hydrodynamic zoom-in simulations presented in GK17. These authors studied the subhalo populations of two MW-mass host halos, referred to as `m12i` ($M_{\text{vir}} = 1.1 \times 10^{12} M_\odot$) and `m12f` ($M_{\text{vir}} = 1.6 \times 10^{12} M_\odot$), which were simulated as part of the Latte suite from the FIRE project [223]. These simulations were performed using the FIRE-2 code [224], which includes the same radiative heating and cooling, star formation, and stellar feedback prescriptions as the original FIRE-1 code in addition to several numerical improvements. The simulations were run in the same cosmological volume (side length $60 h^{-1} \text{ Mpc}$) as the AGORA project [262]; the `m12i` and `m12f` zoom-in simulation regions each contain a single host halo at redshift $z = 0$ that has no MW-mass neighbors within 3 Mpc. The `m12i` simulation was originally presented in [478]; `m12f`, which was simulated using the same parameters and pipeline, was first presented in GK17. The baryonic mass resolution in these simulations is $\sim 7000 M_\odot$, while the dark matter particle mass is $3.5 \times 10^4 M_\odot$, corresponding to a subhalo mass resolution of $\sim 3 \times 10^6 M_\odot$. We refer the reader to GK17 and [224] for details on the initial conditions, gravitational-force softenings, and models for radiative heating/cooling, star formation, and stellar feedback in these simulations. Halo catalogs were created using AHF [268] and merger trees were generated using the `consistent-trees` merger code [48].

We will compare our results to the `m12i` and `m12f` subhalo populations from three sets of simulations: the hydrodynamic FIRE simulations described above, DMO simulations that were run with identical initial conditions, and the dark matter-plus-disk potential (DISK) simulations presented in GK17. The DISK simulations are identical to the corresponding DMO simulations, but they include gravitational potentials designed to capture the effects of the central disks in the hydrodynamic simulations. In particular, a disk potential is added to each DMO zoom-in simulation at $z = 3$, and its parameters and evolution are tuned to match the central disk that develops in the corresponding FIRE simulation. We refer the reader to GK17 for a detailed description of the DISK simulations.

Figure 1 in GK17 illustrates the dark matter substructure in `m12i` for each type of simulation. The visual differences between the FIRE and DMO subhalo populations qualitatively show that baryonic physics in the FIRE simulations lowers both the total number of surviving subhalos and the number of high- V_{\max} subhalos that contribute to the TBTF problem. This figure also shows that the DISK simulation captures the majority of the subhalo disruption in `m12i`, particularly in the innermost regions ($r \lesssim 100$ kpc), which implies that the central disk is largely responsible for the subhalo disruption in the corresponding hydrodynamic simulation. We have verified the quantitative results in GK17 by calculating velocity functions and radial distributions for the `m12i` and `m12f` subhalo populations in the FIRE, DISK, and DMO simulations. Note that, as in GK17, we scale all subhalo masses by a factor of $1 - f_b$ and all subhalo circular velocities by a factor of $\sqrt{1 - f_b}$ in our post-processing of the DMO and DISK halo catalogs, where $f_b = \Omega_b/\Omega_m \simeq 0.17$ is the cosmic baryon fraction. The mass correction accounts for the fact that the baryonic mass in the hydrodynamic simulations is included in the dark matter particles in the DMO simulations, and the circular velocity correction is an approximate way to account for reduced subhalo densities due to stellar feedback, similar to the prescription in [502]. Neither of these corrections affect our results.

To study disrupted subhalos in the FIRE simulations, we select subhalos that disappear from the `m12i` and `m12f` halo catalogs after $z = 3$. We choose this redshift in order to match the initial redshift of the DISK simulations in GK17; note that there are very few subhalos disrupted before $z = 3$ that pass our subsequent minimum circular velocity cuts. We restrict our analysis to first-order subhalos (i.e., we exclude subhalos of subhalos); thus, for a disrupted subhalo to be included in our catalog, it must contribute to the host halo at $z = 0$. Operationally, each disrupted subhalo must have a descendant ID equal to the ID of a main-branch progenitor of the final host halo. Meanwhile, we define surviving subhalos as those that remain in the halo catalog at $z = 0$ and have a parent ID that is equal to the host ID, which similarly excludes higher-order subhalos.

To ensure that we study well-resolved subhalos, we restrict both disrupted and surviving subhalos to those with peak circular velocity $V_{\text{peak}} > 10 \text{ km s}^{-1}$ in our fiducial model, where V_{peak} is defined as the largest maximum circular velocity a subhalo attains along its entire main branch. This is a conservative choice; for example, GK17 presented velocity functions using the cut $V_{\max} > 5 \text{ km s}^{-1}$, where V_{\max} is the maximum circular velocity at $z = 0$. However, this cut ensures that we train our

algorithm on subhalos that are consistent with those we will classify in a lower-resolution zoom-in simulation suite. By choosing a V_{peak} threshold rather than a V_{max} threshold, we also avoid biasing our subhalo selection with a redshift-dependent cut, since V_{peak} — unlike V_{max} — is not defined at a particular redshift. The $V_{\text{peak}} > 10 \text{ km s}^{-1}$ cut results in a combined total of 566 surviving subhalos and 872 disrupted subhalos from `m12i` and `m12f`, which we combine to form our fiducial training set. In Appendix A, we examine the impact of different training sets and minimum circular velocity cuts, and we present the results using the V_{max} cut employed in GK17 for comparison.

2.4 Random Forest Classification

2.4.1 Overview

We use the random forest algorithm from the package `Scikit-Learn` [362] to classify disrupted and surviving subhalos. We refer the reader to the `Scikit-Learn` documentation for a detailed description of the algorithm, but we outline the most important aspects here. A random forest is a collection of decision trees, each of which is tuned to classify objects based on their input properties. Each tree in the forest is trained on a random sample of the training data with replacement, using a random subset of the input features at each split in the learning process, with the goal of predicting the classes of the objects in the training set as accurately as possible according to some metric. For example, the default `Scikit-Learn` implementation minimizes the Gini impurity of the classifier’s prediction. The random forest prediction for a given object is the majority vote of the tuned decision trees, while the classification probability is equal to the fraction of trees that predict a certain class. In this work, we label subhalos as either surviving until $z = 0$ or disrupted at some earlier time; thus, our random forest objects are subhalos, and our decision trees vote for whether each subhalo is disrupted or survives until $z = 0$. Note that our model does not explicitly include enhanced mass stripping due to baryonic effects, since we simply label subhalos as disrupted or surviving.

We train our classifier using the disrupted and surviving subhalos from `m12i` and `m12f` described above. We train on subhalo properties that depend on the entire history of each subhalo to avoid biasing the classifier by using properties defined at specific redshifts — for example, at $z = 0$ for surviving subhalos or at the final available redshift for disrupted subhalos. In particular, since we aim to classify subhalos in DMO halo catalogs that have survived to $z = 0$ but are likely to be disrupted in hydrodynamic resimulations, training our classifier with only present-day properties results in too many surviving subhalos because of the systematic evolution of subhalo properties over time.

Thus, we train on the following properties: pericentric distance and scale factor at first pericentric passage after accretion (d_{peri} , a_{peri}), and scale factor, virial mass, and maximum circular velocity at accretion (a_{acc} , M_{acc} , V_{acc}). In principle, we could train the classifier on additional subhalo properties at pericenter or accretion; these properties could also include information about the host

halo, such as subhalo scale radius in units of the host halo’s scale radius. Indeed, random forests are well-suited to classifying objects using a large number of features because of the randomized nature of the training process, so we could even use *every* available subhalo property at pericenter and accretion to train the classifier. However, we will show that our five-property classifier performs very well, so we adopt this model to simplify our analysis and avoid overfitting the training data. In addition, we checked whether including the present-day properties V_{\max} and M_{vir} improves our classifier, finding that these properties are much less informative than features defined at pericenter or accretion. We discuss the correlations among the training features below, and we explore the feature selection in more detail in Appendix A.

We calculate the aforementioned subhalo features from the merger trees as follows. We define accretion as the last snapshot, working backward in time from $z = 0$ (for surviving subhalos) or from the redshift of disruption (for disrupted subhalos), at which a subhalo’s host ID is equal to the main halo’s ID. Physically, this occurs when a subhalo enters the virial radius of the host halo for the final time.¹ We then take a_{acc} , M_{acc} , and V_{acc} as the scale factor, virial mass, and maximum circular velocity at the time of accretion for each subhalo. We define pericenter as the first snapshot after accretion at which a subhalo reaches a local minimum in its three-dimensional distance from the center of the host halo. We inspected individual subhalo orbits and determined that selecting the distance from the center of the host halo at the first snapshot after accretion at which a subhalo’s separation from the host increases provides an accurate estimate of d_{peri} .² For subhalos that do not reach a local minimum in their separation from the host halo after accretion, we define d_{peri} as the instantaneous distance from the center of the host. In particular, for surviving subhalos on infalling orbits that have not experienced a pericentric passage by $z = 0$, we define d_{peri} as the distance from the host at $z = 0$. Analogously, for destroyed subhalos on infalling orbits that have not reached pericenter by the time of disruption, we define d_{peri} as the distance from the host at the time of disruption.

2.4.2 Choice of Subhalo Features

We choose the subhalo properties listed above because we expect them to correlate with subhalo disruption. Several of these properties are motivated by the results in GK17, which show that most of the subhalo disruption in `m12i` and `m12f` is caused by the central galactic disk in each simulation. For example, Figure 2.1 shows the joint and marginal distributions of d_{peri} and a_{peri} for disrupted and surviving subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ in `m12i` and `m12f`. Disrupted subhalos, shown in red, tend to have closer pericentric passages that occur at earlier times — or smaller values of a_{peri} — than their surviving counterparts, which are shown in blue. The d_{peri} distributions make sense

¹Note that a subhalo could have been contained within the host halo’s virial radius at an earlier time and later reaccreted; we select the final accretion event for each subhalo.

²Given a spacing of ~ 25 Myr between halo catalog snapshots and a generous subhalo orbital velocity of $\sim 300 \text{ km s}^{-1}$ at pericenter, the uncertainty in d_{peri} is only $\sim 8 \text{ kpc}$.

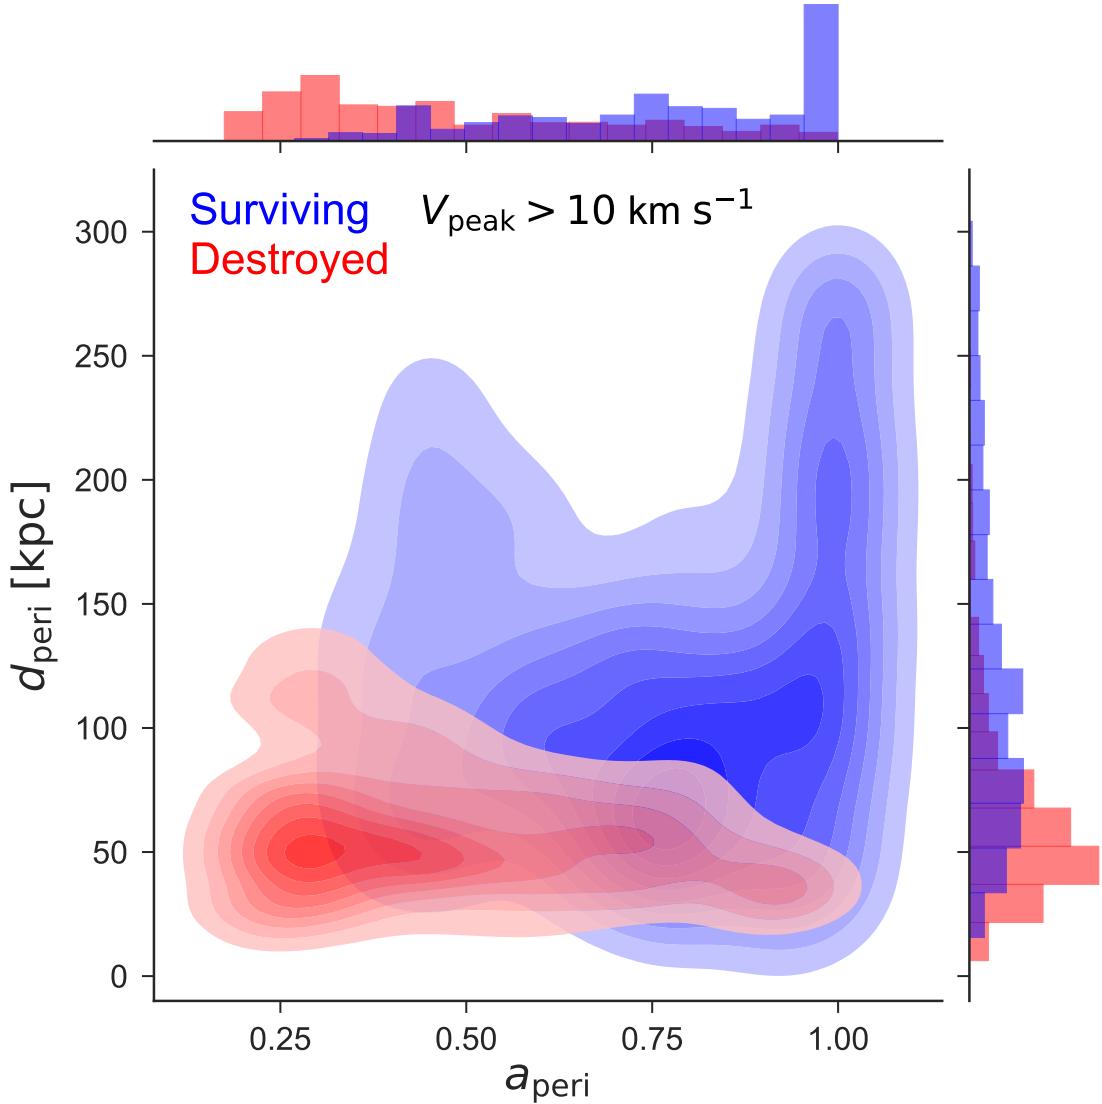


Figure 2.1: Normalized joint and marginal distributions of pericentric distance and scale factor at first pericentric passage after accretion for surviving (blue) and disrupted (red) subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ in the `m12i` and `m12f` FIRE simulations. We select disrupted subhalos starting at $a = 0.25$ ($z = 3$).

Training Features	OOB Score	Accuracy (Disrupted)	Accuracy (Surviving)
d_{peri}	72%	$80\% \pm 3\%$	$58\% \pm 4\%$
$d_{\text{peri}}, a_{\text{peri}}$	82%	$88\% \pm 2\%$	$72\% \pm 3\%$
$d_{\text{peri}}, a_{\text{peri}}, a_{\text{acc}}$	85%	$87\% \pm 2\%$	$82\% \pm 4\%$
$d_{\text{peri}}, a_{\text{peri}}, a_{\text{acc}}, M_{\text{acc}}$	85%	$88\% \pm 2\%$	$81\% \pm 3\%$
$d_{\text{peri}}, a_{\text{peri}}, a_{\text{acc}}, M_{\text{acc}}, V_{\text{acc}}$	85%	$89\% \pm 2\%$	$80\% \pm 4\%$

Table 2.1: Performance metrics for five different random forest classifiers trained on disrupted and surviving subhalos from the `m12i` and `m12f` FIRE simulations with $V_{\text{peak}} > 10 \text{ km s}^{-1}$. The first column lists the subhalo features used to train each classifier. The second column lists the out-of-bag classification score, which is the percentage of subhalos in the training data identified correctly when each tree does not vote on subhalos in its own training set. The third and fourth columns list the percentage of disrupted and surviving subhalos in the test set that are identified correctly by each classifier, averaged over 100 test-training splits. The test set is the collection of subhalos from the `m12i` and `m12f` FIRE simulations with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ that are not included in the training set. We also indicate the standard deviation of each classification accuracy. Note that the ratio of disrupted to surviving subhalos in our fiducial halo catalog is roughly 3:2.

physically; subhalos that pass close to the center of the host experience significant tidal forces due to the galactic disk and are therefore more likely to disrupt.³ Next, consider the a_{peri} dependence: subhalos that reach pericenter earlier have relatively low masses at pericenter and tend to experience more pericentric passages, both of which contribute to enhanced disruption. Although a_{peri} and a_{acc} are somewhat degenerate properties, we find that including a_{acc} improves our results, likely because subhalos accreted at higher redshifts are tidally stripped for longer periods of time, making them more susceptible to disruption.

Figure 2.2 illustrates the V_{acc} and M_{acc} distributions for disrupted and surviving subhalos in `m12i` and `m12f`. Interestingly, even though these features mainly contain information about internal rather than orbital subhalo properties, they are useful for identifying disrupted subhalos; as we show below, these properties account for 16% of the total feature importance score for our fiducial five-property classifier. At the low-mass end of the subhalo population, subhalos with lower values of V_{acc} are more likely to be disrupted. In particular, the survival of low-mass subhalos at fixed M_{acc} is dictated by tidal effects that preferentially disrupt lower-concentration subhalos, i.e., subhalos with smaller values of V_{acc} at fixed M_{acc} . However, at the high-mass end of the subhalo population, subhalos with larger values of V_{acc} are more likely to be disrupted. This behavior suggests that baryonic mechanisms, in addition to the tidal effects of the central disk, contribute to subhalo disruption in the FIRE simulations. Specifically, it is plausible that V_{acc} and M_{acc} encode information about stellar feedback, which can soften central density cusps. In particular, we expect high-mass subhalos

³GK17 found that the amount of disruption is largely insensitive to the shape and mass of the central disk, so subhalo disruption in these simulations is at least partly due to disk shocking rather than tidal stripping.

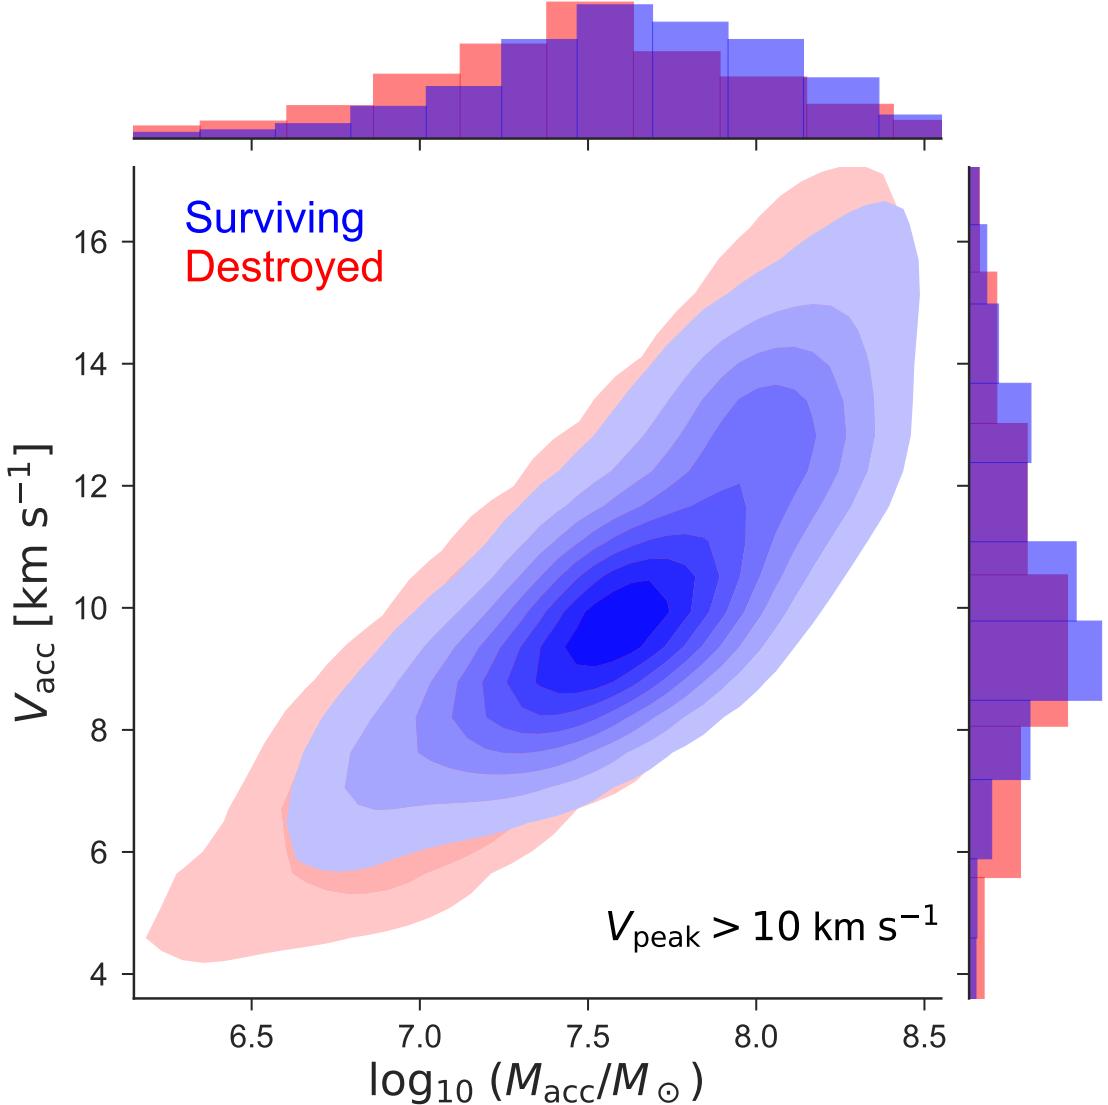


Figure 2.2: Normalized joint and marginal distributions of maximum circular velocity and virial mass at accretion for surviving (blue) and disrupted (red) subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ in the `m12i` and `m12f` FIRE simulations. We select disrupted subhalos starting at $a = 0.25$ ($z = 3$); note that $V_{\text{acc}} < V_{\text{peak}}$ for subhalos that are stripped prior to infall (e.g., see [46]).

with larger values of V_{acc} to host more massive galaxies and to experience more significant baryonic feedback, i.e., high-mass subhalos with larger values of V_{acc} are more likely to be disrupted. Thus, even though M_{acc} and V_{acc} are highly correlated, it is useful to train on both properties because subhalo concentration determines V_{acc} at fixed M_{acc} and provides physical information about whether a subhalo subject to given tidal forces is disrupted. The advantage of random forest classification is that it captures these complex relationships between subhalo properties and subhalo disruption.

2.4.3 Training and Validation

To train our classifier, we use the `GridSearchCV` function to search the space of random forest hyperparameters and select the ones that yield the highest out-of-bag (OOB) classification score averaged over ten cross-validation folds of the training data.⁴ These hyperparameters include the number of trees in the forest, the depth of each tree, the maximum number of features used by each tree, and the loss function. We train the classifier using a randomly selected 75% of the disrupted and surviving subhalos from our fiducial training set, with replacement. The number of folds and the ratio of the test-training split do not affect our results. The raw percentage of subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ from the hydrodynamic `m12i` and `m12f` simulations that are identified correctly by our classifier is 95%. We cross-validate this result by computing the OOB classification score, which is defined as the percentage of subhalos from the training data that the random forest classifies correctly when each tree does not vote on subhalos in its own training set. The optimal OOB score for our fiducial five-property classifier is 85%, and we find that at least 20 trees are needed to achieve this OOB score. Our classifier therefore identifies subhalos accurately, although the gap between the overall classification accuracy and the OOB scores suggests that we mildly overfit the training data. In particular, the raw accuracy is higher than the OOB score because decision trees are allowed to vote on subhalos within their respective training sets when classifying all subhalos. To illustrate the relative importance of each subhalo feature, Table 2.1 shows the OOB score along with the percentage of correct and incorrect predictions for subhalos in the test set, which is the set of all subhalos that are not included in the training set, for five different classifiers. We calculate these scores for each classifier by using the hyperparameters determined by `GridSearchCV` and averaging the results over 100 test-training splits. Each row of Table 2.1 lists the results for a classifier trained using an additional subhalo feature; as we add training features, the OOB score and the total classification accuracy generally improve. Note that there are more disrupted subhalos than surviving subhalos in our fiducial training set, so the raw classification accuracy for each set of features is higher than the mean classification accuracy inferred from Table 2.1. Thus, while the classification accuracy for surviving subhalos decreases when M_{acc} and V_{acc} are added, the increase in classification accuracy for disrupted subhalos outweighs this effect. We emphasize, however, that

⁴In n -fold cross-validation, the training set is divided into n subsets of equal size; $n - 1$ of these subsets are used for training, the remaining subset is used for cross-validation, and this procedure is repeated once for each possible cross-validation subset.

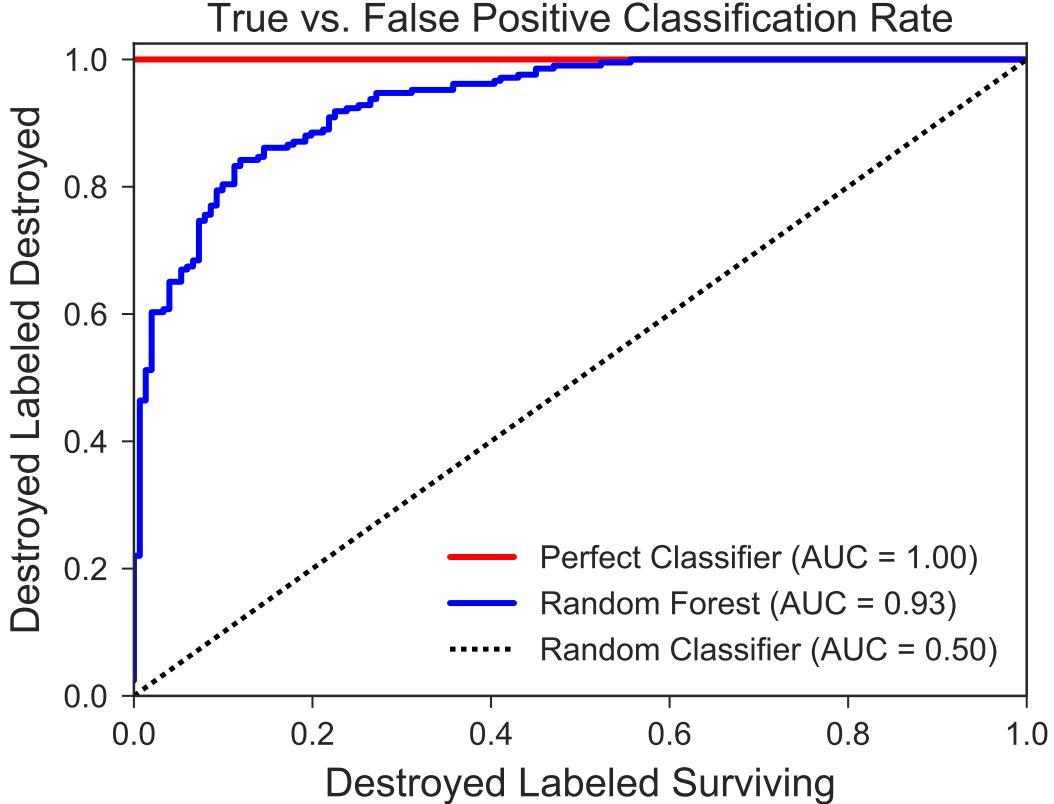


Figure 2.3: True- vs. false-positive classification rate for our fiducial five-property random forest classifier, which is trained on subhalos from the `m12i` and `m12f` FIRE simulations with $V_{\text{peak}} > 10 \text{ km s}^{-1}$. These classification rates apply to subhalos that are not included in the training set. The AUC is equal to 1 for a perfect classifier (red), 0.93 for our random forest classifier (blue), and 0.5 for a random classifier (black).

d_{peri} , a_{peri} , and a_{acc} contain most of the information about subhalo disruption in `m12i` and `m12f`.

Next, we examine our classifier’s receiver operating characteristic (ROC) curve, which illustrates the rate of true- versus false-positive classifications for subhalos in the test set. The ROC Curve for our five-property classifier is shown in Figure 2.3. The red and black lines illustrate perfect (100% true-positive rate) and random (true-positive rate equal to false-positive rate) classifiers. We quantify our classifier’s performance by calculating the area under the ROC curve (AUC), which confirms that the random forest classifies subhalos in the FIRE simulations accurately: its AUC is 0.93, while a random classifier has an AUC equal to 0.5 and a perfect classifier has an AUC equal to 1. Note that Figure 2.3 shows the ROC curve for a particular test-training split, but the scatter in the ROC curves for different test-training splits is small.

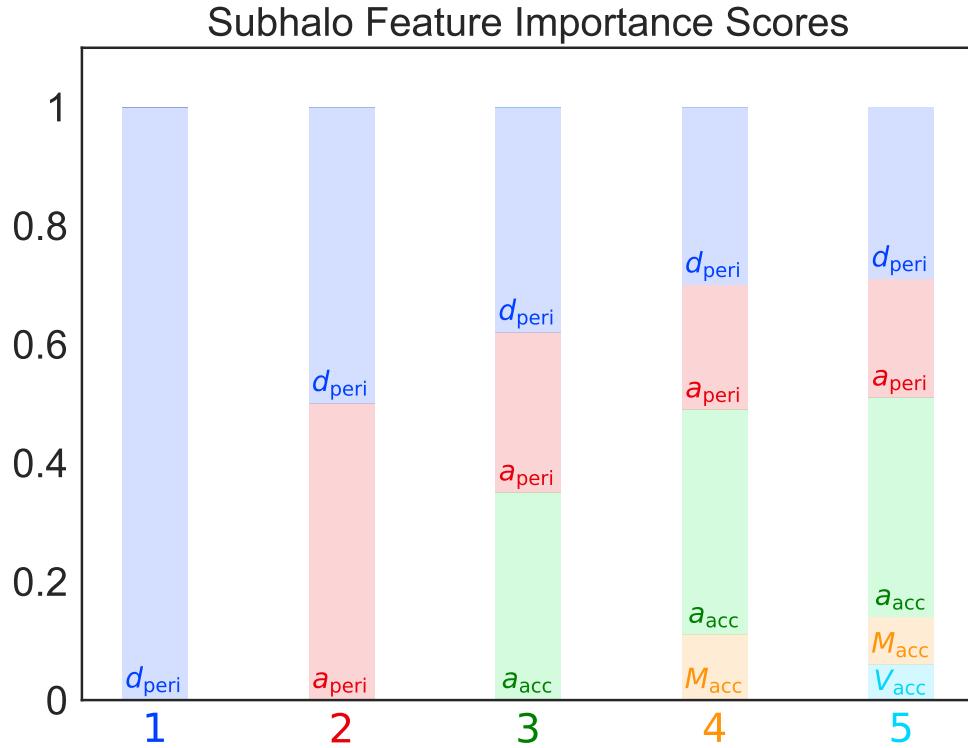


Figure 2.4: Feature importance scores for the five subhalo properties used to classify disrupted and surviving subhalos in the `m12i` and `m12f` FIRE simulations. The colored bars above each property indicate the feature importance scores averaged over 100 test-training splits when that property is added to the training features. Thus, the columns correspond to the five different classifiers in Table 2.1. For a given classifier, each property's score indicates its relative importance for classifying disrupted and surviving subhalos. Here d_{peri} and a_{peri} are the pericentric distance and scale factor at first pericentric passage after accretion, and a_{acc} , M_{acc} , and V_{acc} are the scale factor, virial mass, and maximum circular velocity at accretion.

2.4.4 Importance of Subhalo Features

The random forest algorithm determines the feature importance of the various subhalo properties included in the training process. The feature importance indicates the relative importance of each property for predicting whether a given subhalo is disrupted or whether it survives until $z = 0$. In particular, a property’s feature importance score is the suitably normalized change in the OOB classification score when the property is randomly shuffled among the subhalos in the training set. Thus, the property with the highest feature importance score is the most important for classifying disrupted and surviving subhalos in the `m12i` and `m12f` simulations. Figure 2.4 illustrates the mean feature importance scores for each classifier listed in Table 2.1; for a given classifier, the scores are averaged over 100 test-training splits, and the same hyperparameters are used for each realization. For our fiducial five-property classifier, which corresponds to the fifth column of Figure 2.4, we find mean feature importance scores of 0.28 for d_{peri} , 0.21 for a_{peri} , 0.35 for a_{acc} , 0.08 for M_{acc} , and 0.08 for V_{acc} . The variance in the feature importance scores for different test-training splits is small, and the scores depend very weakly on the random forest hyperparameters.

Figure 2.4 shows that pericentric distance is an important property for determining whether a given subhalo is disrupted; subhalos with close pericentric passages are more likely to be destroyed. The scale factors at accretion and at first pericentric passage after accretion are also important features. In particular, subhalos that accrete and reach pericenter earlier are preferentially disrupted. The fact that a_{acc} has the highest feature importance score suggests that the number of pericentric passages, rather than the distance and scale factor associated with each individual passage, is most strongly correlated with subhalo disruption. However, we note that interpreting the feature importance scores for d_{peri} and a_{peri} is complicated by the fact that we defined these properties as the instantaneous distance and scale factor at the final available snapshot for subhalos on infalling orbits that have not reached their true pericenter. The true pericenters for such subhalos occur at smaller values of d_{peri} and larger values of a_{peri} than we have assigned here; in a more detailed analysis, we would need to calculate these features by fitting individual subhalo orbits. However, the fraction of disrupted (surviving) subhalos in our fiducial training set that have not reached their true pericenter by the time of disruption ($z = 0$) is only 17% (20%), so the feature importance for d_{peri} and a_{peri} is reasonably accurate.

2.4.5 Model Limitations

Finally, we note that our classification method, like any other model, has its limitations. In particular,

1. our classifier is only trained on two zoom-in simulations of MW-mass host halos with a specific baryonic physics prescription, and thus it is not clear how well our algorithm will perform on subhalo populations associated with higher- or lower-mass host halos;

2. neither of the hosts that we train on experience a recent major merger, so our classifier might not apply to halos with significantly different formation histories;
3. both hosts form a central galactic disk that is responsible for most of the subhalo disruption, so our classifier mainly captures the dynamical effects of a central disk.

We discuss these limitations in more detail and comment on how they might affect our results in the following section.

2.5 Results

We now present our main results. In Section 2.5.1, we use our classifier to identify subhalos in DMO simulations of `m12i` and `m12f` that are likely to be disrupted in hydrodynamic resimulations. We analyze our predicted surviving subhalo populations by comparing the velocity functions and radial distributions to those from the FIRE, DISK, and DMO simulations in GK17. In Section 2.5.2, we predict surviving subhalo populations for the suite of DMO zoom-in simulations of MW-mass host halos from [315], and we study the resulting velocity functions, radial distributions, and implications for satellite searches.

2.5.1 Predictions for Dark Matter-only Simulations of the FIRE Halos

Subhalo Feature Distributions

There are about twice as many surviving subhalos at $z = 0$ in the DMO simulations of `m12i` and `m12f` as in the corresponding hydrodynamic simulations. As we have discussed, we expect many of these subhalos to be disrupted due to baryonic effects, including stellar feedback, enhanced tidal stripping, and disk shocking, and our random forest classifier can identify such subhalos based on their internal and orbital properties. In particular, to identify subhalos in the `m12i` and `m12f` DMO simulations that are likely to be disrupted by baryonic effects, we select subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ at $z = 0$, and we use our trained classifier to predict whether these subhalos should have been destroyed at some earlier time using their values of d_{peri} , a_{peri} , a_{acc} , M_{acc} , and V_{acc} . Note that this method does not require matching subhalos between DMO and hydrodynamic simulations.

Figure 2.5 shows the joint and marginal distributions of d_{peri} and a_{peri} for surviving subhalos from the `m12i` and `m12f` DMO simulations predicted by our random forest classifier. The random forest predicts a surviving subhalo population in $d_{\text{peri}} - a_{\text{peri}}$ space that agrees well with the hydrodynamic data; we also find good agreement in the spaces defined by the other subhalo features. Of course, since our classifier is trained on subhalos from the `m12i` and `m12f` FIRE simulations, we expect it to perform particularly well on the corresponding DMO simulations, which have identical initial conditions. Nevertheless, these results are encouraging: even though there is no galactic disk or stellar feedback in the DMO simulations, our classifier efficiently predicts subhalo populations that

are in good agreement with the hydrodynamic results. In particular, once the classifier has been trained on the hydrodynamic simulations, it can immediately predict surviving subhalo populations from DMO halo catalogs. Simulations that include baryonic effects by hand, such as the DISK simulations presented in GK17, are complementary to our approach, since they provide more direct physical modeling at the expense of increased computational costs.

In general, at least three mechanisms contribute to enhanced subhalo disruption in the `m12i` and `m12f` hydrodynamic simulations relative to the DMO simulations: tidal effects due to the central galactic disk, stellar feedback, and characteristic changes in subhalo orbits due to the presence of baryons. The results from the DISK simulations in GK17 indicate that the central disk is the main source of subhalo disruption in these simulations, but the frequency of disruption events might be enhanced by stellar feedback, which can soften central density cusps [201, 369, 502, 135]; as noted above, we multiply all circular velocities in the DMO and DISK simulations by a factor of $\sqrt{1 - f_b}$ to approximate this effect. Meanwhile, [500] analyzed the orbital properties of subhalos in hydrodynamic and DMO zoom-in simulations of an MW-mass host halo from the Aquarius Project and found that the distributions of subhalos in different orbital families change when baryons are included. It is difficult to assess the importance of the characteristic differences in internal and orbital subhalo properties between hydrodynamic and DMO simulations in general; however, the fact that we predict subhalo feature distributions starting from DMO halo catalogs that agree with hydrodynamic results suggests that these effects are relatively unimportant.

Subhalo Counts

Having shown that we can predict the feature distributions of surviving subhalos from DMO simulations of `m12i` and `m12f`, we turn to our predictions for the number of surviving subhalos as a function of various properties. In Figure 2.6, we present our predictions for the `m12i` and `m12f` velocity functions; the top panels show the velocity functions evaluated using V_{\max} , and the bottom panels show the velocity functions evaluated using V_{peak} . The blue lines show the most probable surviving subhalo populations predicted by our random forest algorithm for each host halo; we also plot the FIRE, DISK, and DMO results for comparison. We restrict the velocity functions to subhalos within 300 kpc of the center of their respective host at $z = 0$, since this roughly corresponds to the virial radii of `m12i` and `m12f`. Similarly, Figure 2.7 shows the distribution of tangential and radial orbital velocities for subhalos within 300 kpc of their respective host at $z = 0$, and Figure 2.8 shows the radial distribution of surviving subhalos at $z = 0$ within each host halo. In Figures 2.6–2.8, we only include subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ to match the cut used in our fiducial training set. The bottom panels in these figures show the number of surviving subhalos predicted by the most probable realization of our random forest classifier divided by the number of subhalos found in each hydrodynamic simulation. We also plot the Poisson error associated with the random forest predictions as shaded areas in each figure. In Appendix A, we show that the intrinsic scatter

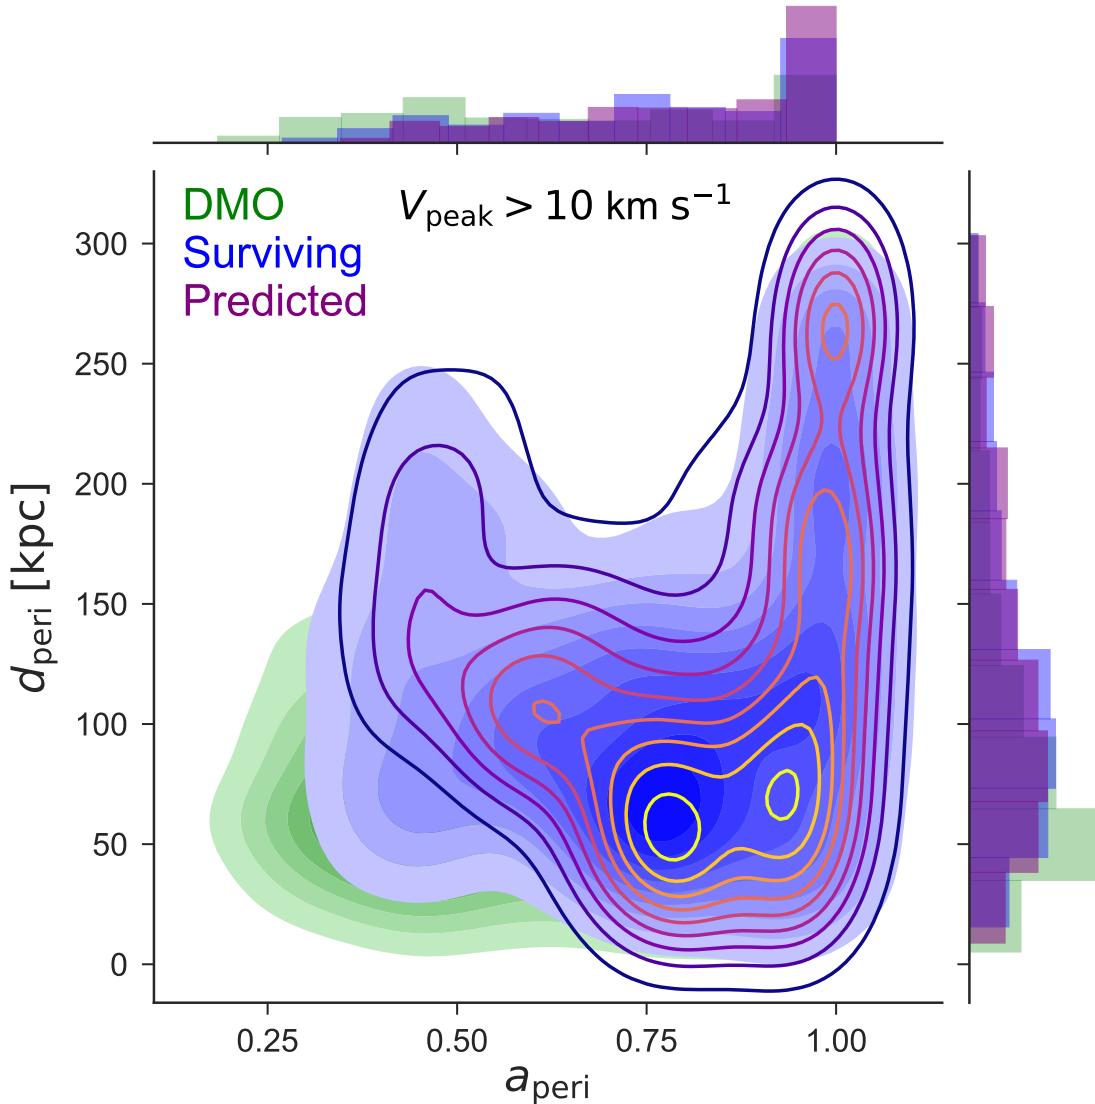


Figure 2.5: Normalized joint and marginal distributions of pericentric distance and scale factor at first pericentric passage after accretion for surviving subhalos in the `m12i` and `m12f` FIRE simulations (blue); surviving subhalos from the corresponding DMO simulations are shown in green. The unfilled contour and purple histograms show the most probable surviving subhalo population from the `m12i` and `m12f` DMO simulations predicted by our random forest classifier.

in the random forest predictions is small.

There are several interesting aspects of Figures 2.6–2.8 that are worth exploring. Our random forest algorithm predicts that the amount of substructure in each host is significantly reduced relative to the DMO simulations, bringing the velocity functions and radial distributions into good agreement with the FIRE results. The random forest predictions for the velocity functions are comparable to the DISK simulations at low velocities, which indicates that the effects of the disk are largely encoded in the subhalo properties that we use to train our classifier, at least for subhalos with low values of V_{\max} or V_{peak} . However, the random forest predicts more subhalo disruption than the DISK simulations for $V_{\max} \gtrsim 15 \text{ km s}^{-1}$ or $V_{\text{peak}} \gtrsim 20 \text{ km s}^{-1}$ and generally matches the FIRE results more closely in these regimes. The minor discrepancies for $V_{\max} \gtrsim 15 \text{ km s}^{-1}$ are likely caused by enhanced mass stripping due to baryonic effects, which would shift the predictions toward smaller velocities at high V_{\max} .

Our predicted radial distributions are also generally comparable to the DISK simulations; however, for $30 \text{ kpc} \lesssim r \lesssim 100 \text{ kpc}$, where the disk should be particularly effective at disrupting subhalos, our classifier predicts more subhalo disruption than the `m12i` DISK simulation and matches the FIRE results more closely for both hosts. Finally, Figure 2.7 shows that our classifier predicts a substantial reduction in the number of subhalos with low tangential velocities, even though it is not explicitly trained on orbital velocities. Our predicted tangential and radial velocity distributions are similar to the DISK results for `m12i`, while we slightly overpredict the number of high- V_{tan} and high- V_{rad} subhalos for `m12f`. Comparing our predictions to the DISK simulations is a particularly useful way to assess whether our classifier captures baryonic physics beyond the dynamical effects of a central galactic disk, since the DISK simulations do not modify internal subhalo properties. Thus, Figures 2.6–2.8 suggest that our classifier captures both the tidal effects of a disk and additional baryonic processes that contribute to subhalo disruption.

Our random forest classifier predicts that many subhalos with large values of V_{\max} and V_{peak} should be disrupted, while these subhalos are not necessarily destroyed in the DISK simulations (see Figures 2.6 and A.3). These subhalos either orbit at large radii, so that they are not significantly affected by the disk, or they are too tightly bound to be disrupted by the disk alone. We find that 45% (84%) of the disrupted subhalos from `m12i` and `m12f` with $V_{\text{peak}} > 20 \text{ km s}^{-1}$ have pericentric passages within 50 kpc (100 kpc) of their respective hosts. The disk does not seem to be the main factor that contributes to the destruction of the remaining subhalos, though a combination of stellar feedback and tidal forces could lead to their disruption. Interestingly, the region of the V_{\max} and V_{peak} functions where we predict enhanced subhalo disruption relative to the DISK simulations ($V_{\max} \gtrsim 15 \text{ km s}^{-1}$ and $V_{\text{peak}} \gtrsim 20 \text{ km s}^{-1}$) corresponds to the regime where baryonic physics can efficiently create cored subhalo density profiles [109, 442, 176]. It is also intriguing that our classifier predicts both the V_{\max} and V_{peak} functions accurately, even though it does not account for enhanced mass stripping beyond the $\sqrt{1 - f_b}$ circular velocity correction,

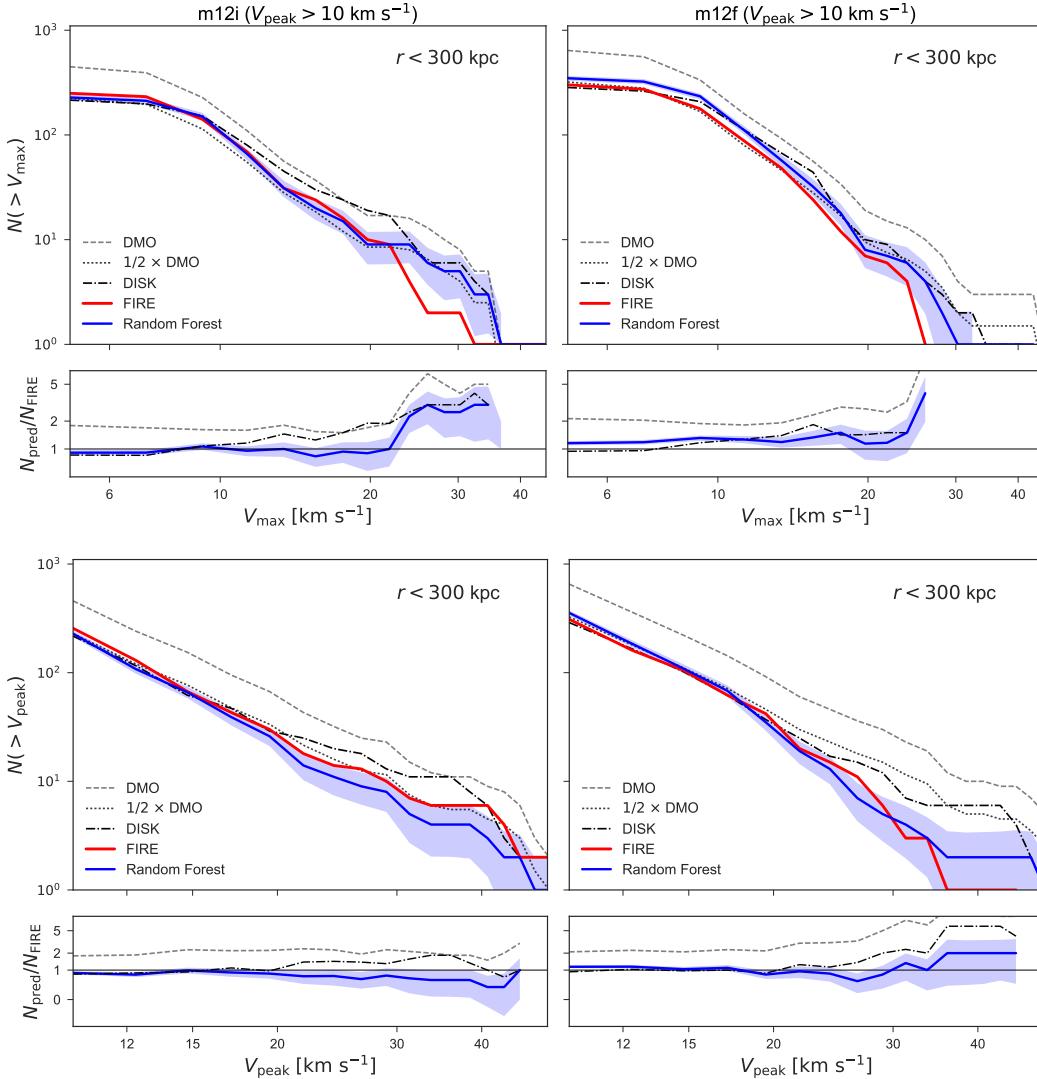


Figure 2.6: Velocity functions for subhalos hosted by m12i (left) and m12f (right), predicted from DMO simulations of these hosts by our random forest classifier (blue). The top panels show velocity functions evaluated using the maximum circular velocity at $z = 0$, and the bottom panels show velocity functions evaluated using the peak circular velocity V_{peak} . Our classifier is trained on subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ from both hydrodynamic simulations. The FIRE (red), DISK (dot-dashed), and DMO (dashed) results are shown for comparison; recall that we scale circular velocities in the DISK and DMO halo catalogs by a factor of $\sqrt{1 - f_b}$. Dotted lines show the DMO results scaled by a factor of 1/2 for comparison. We restrict these velocity functions to subhalos within 300 kpc of their respective host at $z = 0$. The bottom panels show the ratio $N_{\text{pred}}/N_{\text{FIRE}}$, where N_{pred} is the number of surviving subhalos predicted by the random forest and N_{FIRE} is the number of subhalos in each FIRE simulation. Shaded areas show the standard deviation about the most probable random forest prediction for 1000 draws from a Poisson distribution with a mean value of N_{pred} at each value of V_{max} or V_{peak} .

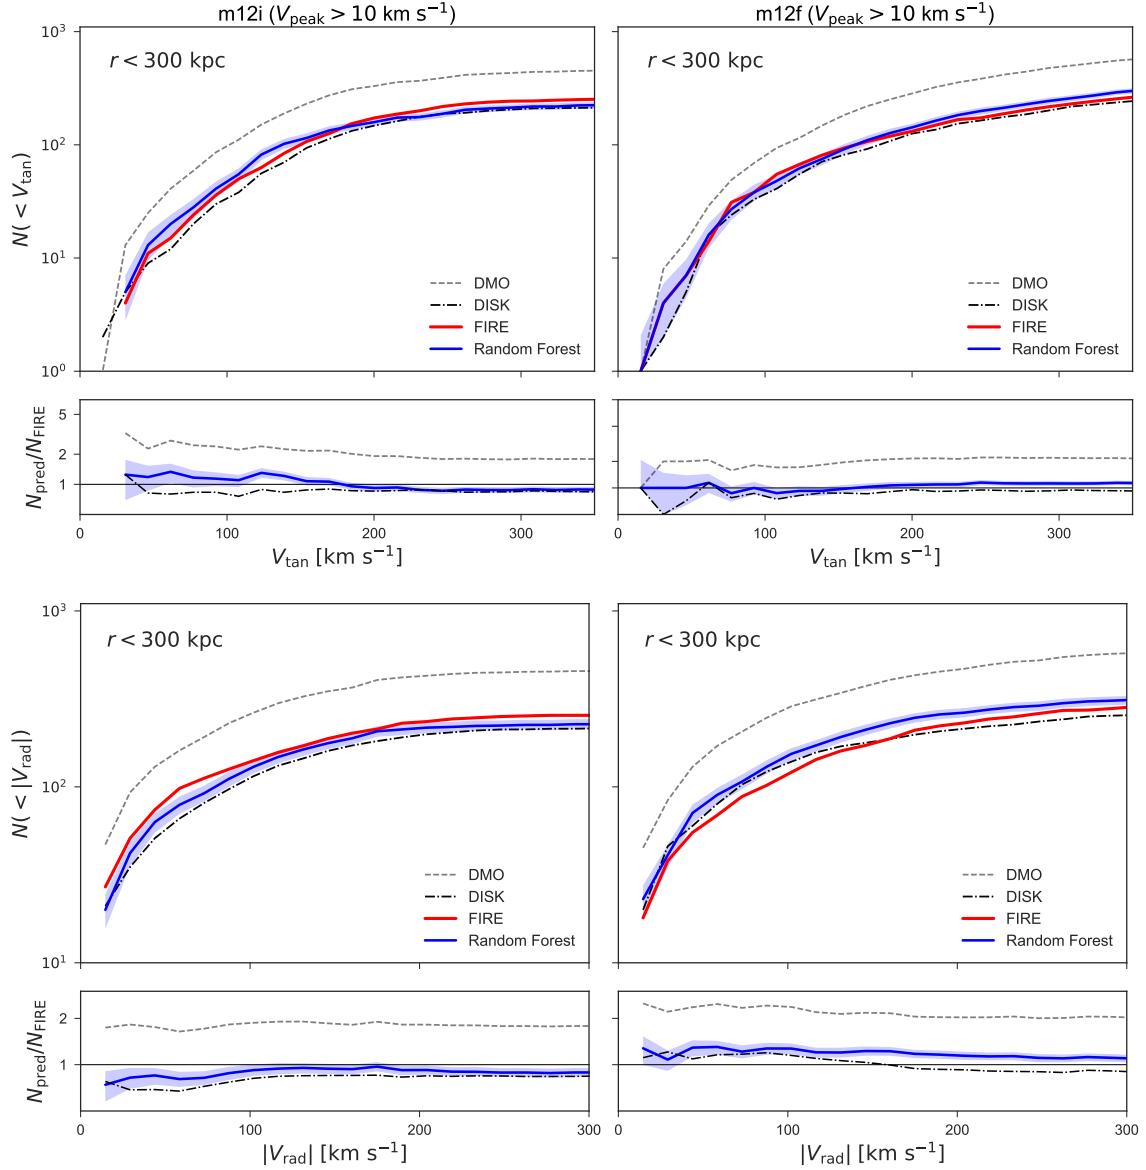


Figure 2.7: Distributions of tangential orbital velocities (top) and radial orbital velocities (bottom) for subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ hosted by m12i (left) and m12f (right) at $z = 0$, predicted from DMO simulations of these host by our random forest classifier (blue). The classifier is trained on subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ from both FIRE simulations. We restrict these distributions to subhalos within 300 kpc of their respective hosts at $z = 0$. The various curves and panels are described in Figure 2.6.

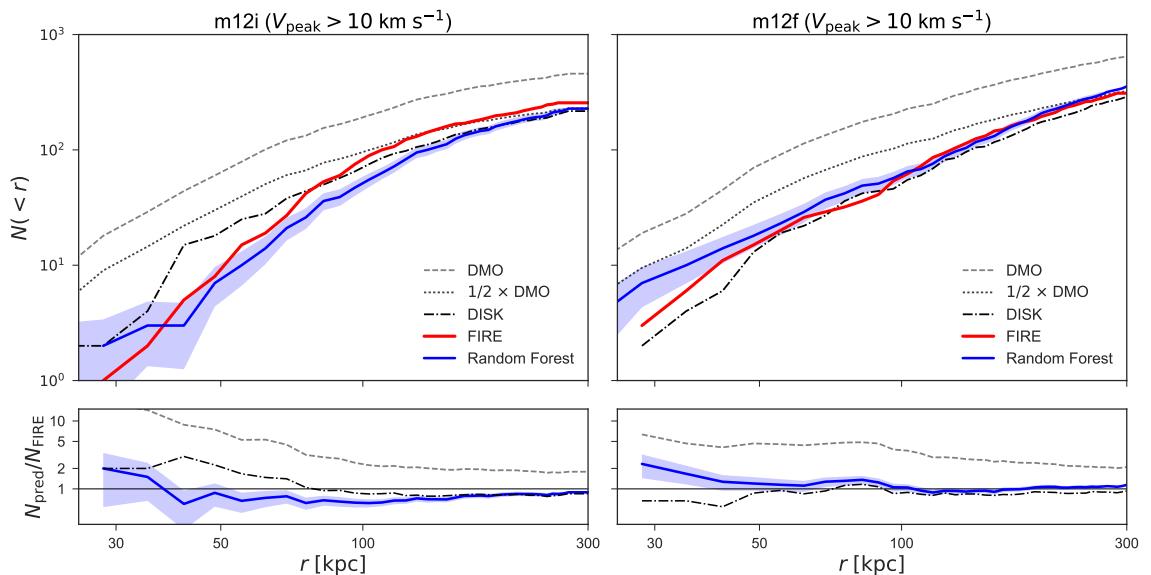


Figure 2.8: Radial distributions of subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ hosted by m12i (left) and m12f (right) at $z = 0$, predicted from DMO simulations of these hosts by our random forest classifier (blue). The classifier is trained on subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ from both FIRE simulations, and the various curves and panels are described in Figure 2.6. The scaled DMO curve overpredicts the number of surviving subhalos at small radii by an order of magnitude, highlighting the enhanced subhalo disruption in the inner regions of the hydrodynamic simulation due to the central disk.

which does not reproduce the hydrodynamic results on its own (for example, compare the ‘Raw DMO’ and ‘DMO’ curves in GK17). Since V_{\max} and V_{peak} are proxies for satellite luminosity, our method can therefore be extended to predict satellite galaxy populations associated with MW-mass host halos (see Figure 2.11); in addition, it can be used to constrain the cumulative mass functions of dark and luminous substructures relevant to gravitational-lensing analyses. Clearly, a more diverse training sample is required in order to make robust predictions regarding the populations of satellite galaxies around the MW and around the MW analogs from the SAGA survey. Nonetheless, Figures 2.6 and 2.8 show that classification algorithms can predict subhalo populations in good agreement with hydrodynamic simulations, providing an efficient way to explore the range of possible satellite galaxy populations associated with a particular host halo.

One could argue that the efficiency of our approach is outweighed by the fact that we must train our classifier on computationally expensive hydrodynamic simulations in order to predict surviving subhalo populations for corresponding DMO simulations. However, as we demonstrate in the following section, our method can be used to predict surviving subhalo populations when hydrodynamic simulations are unavailable. Of course, the surviving subhalo populations we predict in this paper are specific to the FIRE simulations that we use to train our classifier. Nonetheless, even though the generality of our results is limited by the small size of our training set, our work suggests that random forest classification can be used to predict realistic subhalo populations given a sufficiently diverse sample of hydrodynamic training simulations. In addition, we emphasize that our classifier is trained on simulations that yield satellite populations that are consistent with the observed mass functions and velocity dispersion functions for satellites of the MW and M31.

2.5.2 Predictions for a Suite of Dark Matter-only Zoom-in Simulations

Subhalo Counts

We now use our classifier to identify subhalos from a suite of independent DMO zoom-in simulations that are likely to be disrupted in hydrodynamic resimulations. In particular, we predict surviving subhalo populations for the 45 zoom-in simulations of MW-mass host halos from [315]. We refer the reader to [315] for a detailed description of the simulations, but we briefly highlight the most important aspects for this work. The host halos lie in the mass range $M_{\text{vir}} = 10^{12 \pm 0.03} M_{\odot}$ and have a variety of formation histories; we plot the mass accretion histories for these hosts in Figure 2.9. Note that **m12i** and **m12f** have formation histories that are consistent with these host halos, so we expect our model to perform well on this simulation suite. Of course, our model would not accurately predict subhalo disruption for hosts with significantly different formation histories due to the limited size of our training set. The zoom-in simulations were run at a lower resolution than the DMO simulations of **m12i** and **m12f**; the dark matter particle mass is $3 \times 10^5 M_{\odot}$, and [315] estimated that $V_{\max} \sim 9 \text{ km s}^{-1}$ is a conservative lower limit for the subhalo circular velocity resolution. Halo catalogs and merger trees were generated using the **ROCKSTAR** halo finder and the

`consistent-trees` merger code [47, 48]. Again, we scale all subhalo masses by $1 - f_b$ and all subhalo circular velocities by $\sqrt{1 - f_b}$ in our post-processing of the halo catalogs. As noted at the end of Section 2.2, the cosmological parameters for these simulations are slightly different than those used in the FIRE simulations, and we adjust the parameters in our analysis accordingly.

In Figure 2.10, we plot the maximum circular velocity functions and radial distributions for the subhalo populations from this simulation suite, along with those predicted by the most probable realization of our random forest classifier for each simulation. We also plot the results from the `m12i` and `m12f` FIRE simulations, along with the mean DMO curves scaled by a constant factor, for comparison. In particular, we scale the mean DMO curves by a factor of 2/3 so that the average number of subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ and $r < 300 \text{ kpc}$ matches the mean random forest prediction. The random forest predictions were generated using the method described above. We classify subhalos in each zoom-in simulation using the features d_{peri} , a_{peri} , a_{acc} , M_{acc} , and V_{acc} , and we restrict the velocity functions to subhalos within 300 kpc of their respective host at $z = 0$. We plot the most probable realization of the random forest prediction for each host. The intrinsic scatter in our random forest predictions is small.

Figure 2.10 shows that the reduction in the total number of subhalos predicted by our random forest classifier is larger than the host-to-host scatter for the subhalo populations in these zoom-in simulations. In particular, our classifier predicts that the total number of subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ and $r < 300 \text{ kpc}$ is reduced by a factor of 2/3, while the 1σ host-to-host scatter corresponds to an 87% reduction at most. This suggests that subhalo disruption due to baryonic effects, such as stellar feedback and the tidal influence of a central galactic disk, should not be neglected in semianalytic models that use the subhalo populations predicted by these DMO simulations as input. In particular, for MW-mass host halos that contain a central galactic disk similar to those found in the `m12i` and `m12f` FIRE simulations, the reduction in substructure due to the disk and other baryonic processes is larger than the scatter in subhalo abundance from host to host, so the impact of baryonic physics cannot be accounted for simply by marginalizing over the subhalo populations of host halos with a range of formation histories.

While the average amount of subhalo disruption is larger than the host-to-host scatter among the subhalo populations in these simulations, the impact of baryons on individual subhalo populations is largely consistent. In particular, our classifier predicts that the hosts with the most subhalos tend to have the largest number of surviving subhalos once baryonic effects are taken into account. Moreover, the number of DMO subhalos and the predicted number of surviving subhalos above different V_{max} thresholds and within various hostcentric radii are highly correlated for this simulation suite. For example, the Spearman rank correlation coefficient between the number of surviving subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ and $r < 300 \text{ kpc}$ predicted by the DMO simulations and by our classifier is 0.74. This implies that the shapes of the velocity functions and radial distributions are not strongly affected by baryonic physics; indeed, the scaled DMO curves in Figure 2.10 are very similar to the

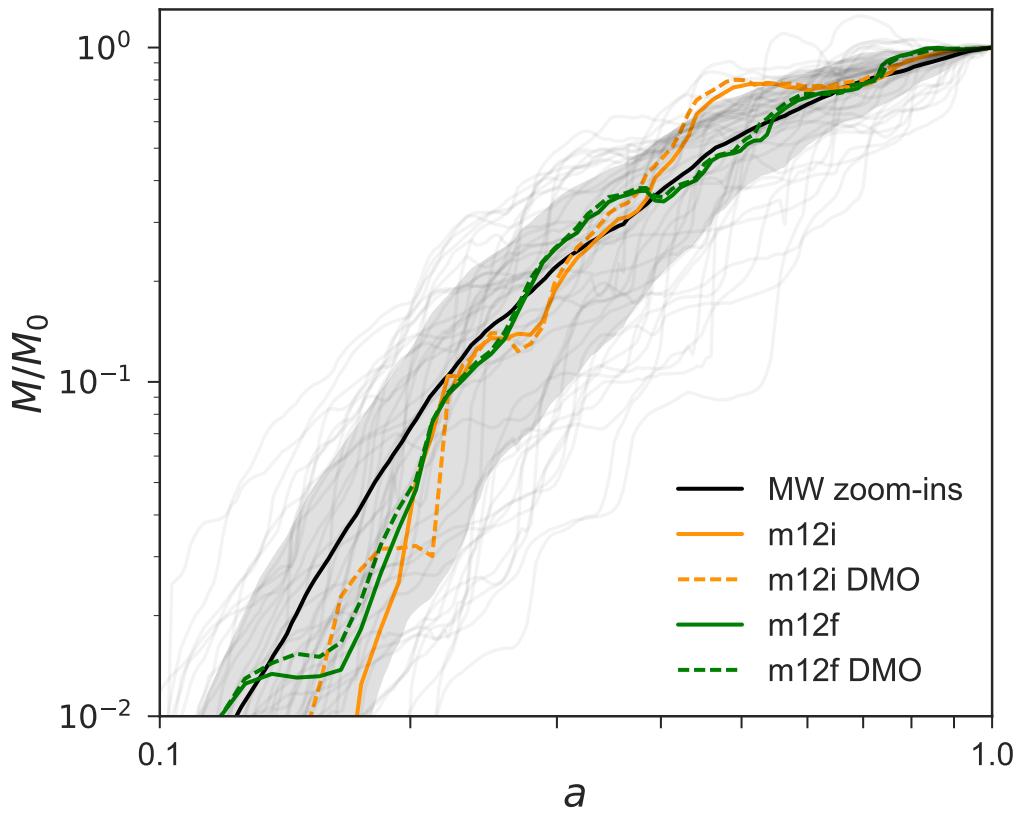


Figure 2.9: Mass accretion histories for the suite of DMO zoom-in simulations of MW-mass host halos presented in [315]. The black line shows the mean mass accretion history for the 45 hosts, and the shaded area shows the associated $\pm 1\sigma$ standard deviation. Mass accretion histories for the **m12i** and **m12f** FIRE simulations are shown in orange and green, respectively.

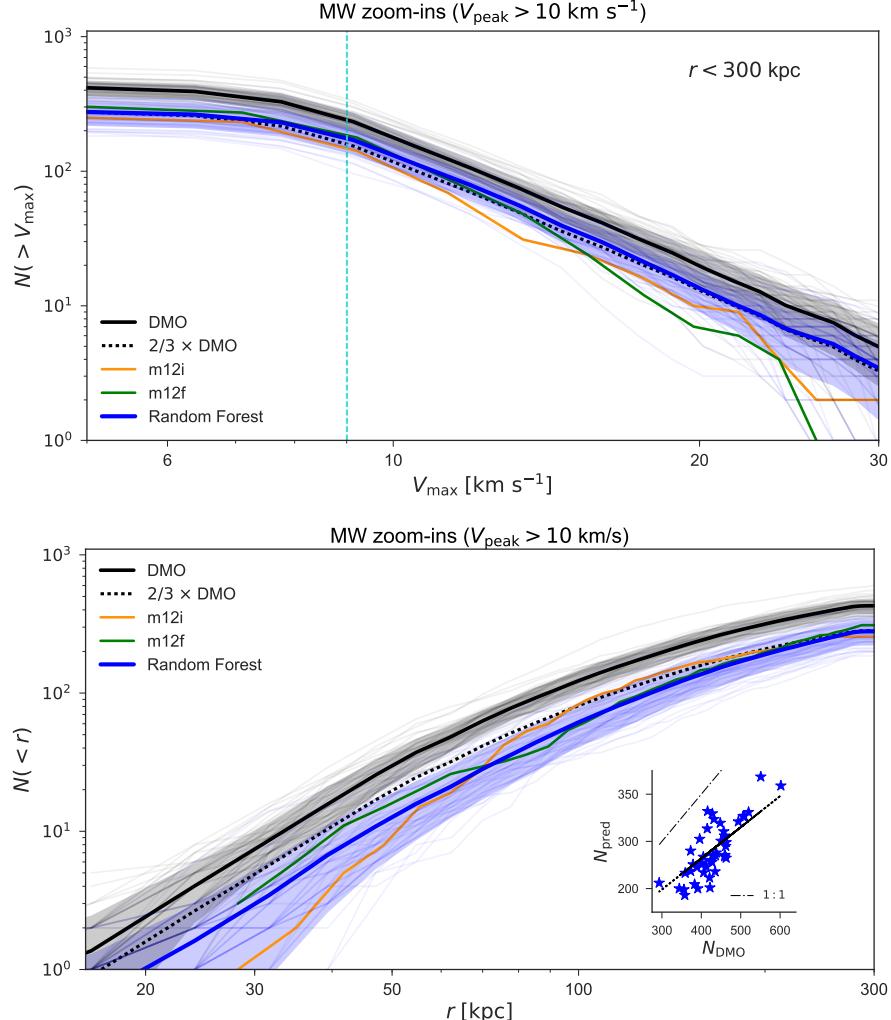


Figure 2.10: Velocity functions (top) and radial distributions (bottom) for the suite of 45 zoom-in simulations of MW-mass host halos presented in [315]. The thick lines show the mean number of subhalos predicted by the DMO simulations (black) and by our random forest classifier (blue), which is trained on the `m12i` and `m12f` FIRE simulations; the shaded areas show the $\pm 1\sigma$ standard deviation of these predictions. The thin lines show the DMO result and the most probable random forest prediction for each host. The thick dotted lines show the mean DMO velocity function and radial distribution scaled by a factor of 2/3 for visual comparison, and the orange and green lines show the results for `m12i` and `m12f`, respectively. Note that the scaled DMO line in the top panel is mostly obscured by the random forest prediction. The inset in the bottom panel shows the number of predicted surviving subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ and within 300 kpc of their respective host versus the number of such subhalos in the corresponding DMO simulations. The thick dotted line in the inset shows the constant fraction of surviving subhalos corresponding to the scaled DMO curves, and the thin dash-dotted line shows a 1:1 relationship for comparison. The vertical line at $V_{\text{max}} = 9 \text{ km s}^{-1}$ in the top panel represents a conservative resolution limit for these simulations.

random forest predictions, except at small radii, where subhalos are preferentially disrupted in the training data. The fractional amount of subhalo disruption is also consistent among the zoom-in simulations. In particular, the number of predicted surviving subhalos with $V_{\max} > 10 \text{ km s}^{-1}$ and $r < 300 \text{ kpc}$ for all 45 hosts is given by $N_{\text{pred}}/N_{\text{DMO}} = 0.65 \pm 0.09$. To illustrate these results, the inset in the bottom panel of Figure 2.10 shows the number of predicted surviving subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ and $r < 300 \text{ kpc}$ for each host versus the corresponding number of subhalos in each DMO simulation. The inset shows that the random forest predictions are consistent with an overall scaling of the DMO subhalo populations. Thus, subhalo disruption due to baryonic effects can be parameterized rather simply for these host halos in the context of our disruption model. We leave a detailed exploration of such a parameterization to future work informed by a wider range of hydrodynamic simulations, but we note that a simple one-parameter rescaling would not be sufficient to model subhalo disruption in detail; for example, Figure 2.10 shows that the shape of the mean radial subhalo distribution is somewhat altered by baryonic physics. Finally, we note that our random forest classifier predicts that these zoom-in simulations typically contain more high- V_{\max} subhalos than `m12i` or `m12f` and more subhalos at small radii than `m12i`. Determining whether these differences represent statistical fluctuations or systematic differences between the FIRE simulations and this simulation suite would require a larger sample of hydrodynamic results for comparison.

Implications for Satellite Searches

Our model, when applied to MW-size zoom-in simulations, suggests that MW-mass host halos are somewhat less likely to host bright satellite galaxies such as the Magellanic Clouds and that they have more extended radial satellite profiles than those inferred from DMO simulations. At face value, both of these predictions seem to be in tension with observations of MW satellites (e.g., see [264]). However, the MW itself could be an outlier, so here we also examine our model’s predictions for the satellite populations of MW analogs.

To estimate the impact of baryonic physics on the luminosity functions of MW, which can be compared to the results of the SAGA survey, we use the V_{peak} -luminosity abundance-matching relation tuned to the r -band luminosity function from the GAMA galaxy survey [295]; we refer the reader to [193] for details on the abundance matching procedure. Figure 2.11 shows the resulting luminosity functions for the MW zoom-in suite, along with the luminosity functions for these hosts predicted by our random forest classifier. We neglect the scatter in the V_{peak} -luminosity relation for this simple estimate because the host-to-host scatter among the zoom-in simulations is larger than the intrinsic scatter in the luminosity function introduced by abundance matching. Our classifier predicts a significant reduction in the number of bright satellites associated with MW analogs; the number of satellites with observed r -band magnitudes $M_{r,o} < -12.3$ inferred from the DMO simulations is 3.0 ± 1.6 , while our random forest predicts that only 1.5 ± 1.3 such satellites exist. Although these estimates of surviving satellite populations are simplistic, it will be interesting to

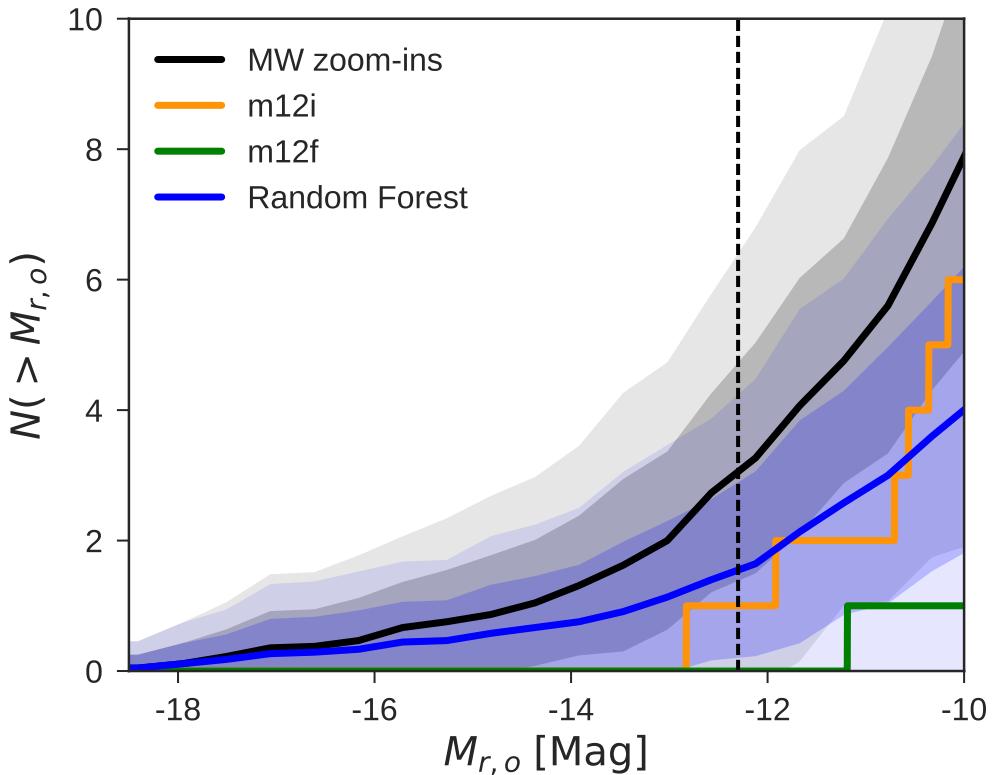


Figure 2.11: Luminosity functions for the DMO zoom-in simulations presented in [315] (black), inferred using the $V_{\text{peak}} - \text{luminosity}$ abundance-matching relation tuned to the r -band luminosity function from the GAMA survey [295]. We do not apply scatter in the $V_{\text{peak}} - M_{r,o}$ relation to highlight the host-to-host variability. The blue line shows our mean prediction for the luminosity function of surviving satellites, and shaded areas show $\pm 1\sigma$ and $\pm 2\sigma$ standard deviations. We also plot luminosity functions for m12i (orange) and m12f (green). Here $M_{r,o}$ is the observed r -band luminosity, and the vertical line corresponds to the completeness limit of the SAGA survey.

compare predictions informed by hydrodynamic simulations to observational results as the number of systems with high completeness limits improves.

2.6 Conclusions and Discussion

To conclude, we summarize our main results as follows.

1. We train a random forest classifier on disrupted and surviving subhalos in two hydrodynamic zoom-in simulations of MW-mass host halos from the FIRE project using five properties of each subhalo: d_{peri} , a_{peri} , a_{acc} , M_{acc} , and V_{acc} .
2. Our classifier identifies subhalos in the FIRE simulations with an 85% OOB classification score and predicts surviving subhalo populations from DMO simulations of these hosts that are in excellent agreement with the hydrodynamic results, often outperforming the DMO-plus-disk simulations presented in [190].
3. We argue that our classifier captures the effects of the central galactic disks that develop in the FIRE simulations, in addition to other baryonic disruption mechanisms such as stellar feedback.
4. We use our classifier to predict surviving subhalo populations for the suite of DMO zoom-in simulations of MW-mass host halos presented in [315], finding that the average amount of subhalo disruption is larger than the host-to-host scatter; however, the baryonic impact on each subhalo population is largely consistent, with $N_{\text{pred}}/N_{\text{DMO}} = 0.65 \pm 0.09$ for subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ and $r < 300 \text{ kpc}$.

We refer the reader to the end of Section 2.4 for a summary of the limitations of our classification method.

There are several interesting avenues for future work. For example, since we find that the average amount of subhalo disruption due to baryonic physics is larger than the host-to-host scatter among the suite of zoom-in simulations analyzed above, this characteristic reduction in the number of subhalos should be taken into account when marginalizing over the effects of baryonic physics for MW-mass host halos that contain a central galactic disk. Thus, it is plausible that the reduced number of surviving subhalos will change the conclusions drawn from semianalytic models that use the subhalo populations predicted by such simulations (e.g., [302, 303]).

Another potential application of our results concerns the radial segregation of dark matter subhalos with respect to various subhalo properties. Subhalo segregation, as studied by [452], directly depends on the subhalo populations predicted by DMO simulations. Since subhalo populations that have been altered by baryonic effects systematically differ from those predicted by DMO simulations, subhalo segregation could be affected by baryonic physics, and our classifier provides an

efficient method for predicting surviving subhalo populations in order to explore this possibility. Meanwhile, resolving the disruption of individual subhalos in detail is an important challenge for current simulations; for example, [451] estimated that 80% of all subhalo disruption in the Bolshoi simulation is numerical, rather than physical. The Latte simulations have $\sim 4000 \times$ smaller dark matter particle mass and $\sim 35 \times$ smaller dark matter force softening than Bolshoi, so these effects are likely much less severe, but it is nevertheless worth exploring whether artificial disruption persists in high-resolution hydrodynamic simulations and how these numerical artifacts might influence our results.

The algorithm presented in this paper is extremely simple, using only five subhalo properties as training features. Nevertheless, these properties encode the majority of the information about subhalo disruption in the `m12i` and `m12f` FIRE simulations, yielding a classifier that predicts surviving subhalo populations from DMO simulations that are in excellent agreement with hydrodynamic results. Of course, as the number of hydrodynamic training simulations grows, it will be worthwhile to explore more sophisticated classification algorithms and to study the feature selection in more detail. It will be interesting to assess how well a classifier can perform in principle, since there are characteristic differences between DMO and hydrodynamic simulations, including changes in subhalo orbits due to the presence of baryons [500], that our simple model cannot capture. Our results hint that these characteristic differences are relatively unimportant, but further tests should be performed using a larger sample of training simulations.

As more high-resolution zoom-in simulations become available, it will become feasible to train classifiers on increasingly diverse datasets, allowing for more robust predictions. Once a classifier has been trained on a wide variety of hydrodynamic simulations, it can predict a range of surviving subhalo populations associated with different central galaxy types and halo formation histories directly from DMO simulations. It is worth exploring whether these predictions can be used as input for neural networks in order to generate large samples of mock halo catalogs, perhaps eliminating the need for certain types of simulations entirely.

Machine-learning algorithms have the potential to identify large samples of realistic subhalo populations that can be used as input for models that populate subhalos with galaxies. Comparing the surviving subhalo populations predicted by such algorithms for host halos on different mass scales could provide insight into the original TBTF problem for MW-mass systems and into analogous problems for host halos of different masses. Moreover, comparing the results of classification algorithms that are trained on hydrodynamic simulations with different implementations of baryonic physics would be a promising step toward parameterizing the impact of baryons on the abundance and properties of dark matter subhalos.

Chapter 3

Modeling the Subhalo–Satellite Connection

Abstract

We develop a comprehensive and flexible model for the connection between satellite galaxies and dark matter subhalos in dark matter-only zoom-in simulations of Milky Way (MW)–mass host halos. We systematically identify the physical and numerical uncertainties in the galaxy–halo connection and simulations underlying our method, including (i) the influence of host halo properties; (ii) the relationship between satellite luminosities and subhalo properties, including the effects of reionization; (iii) the relationship between satellite and subhalo locations; (iv) the relationship between satellite sizes and subhalo properties, including the effects of tidal stripping; (v) satellite and subhalo disruption due to baryonic effects; and (vi) artificial subhalo disruption and orphan satellites. To illustrate our approach, we fit this model to the luminosity distribution of both classical MW satellites and those discovered in the Sloan Digital Sky Survey by performing realistic mock observations that depend on the luminosity, size, and distance of our predicted satellites, and we infer the total satellite population that will be probed by upcoming surveys. We argue that galaxy size and surface brightness modeling will play a key role in interpreting current and future observations, as the expected number of observable satellites depends sensitively on their surface brightness distribution.

3.1 Paper Status and External Contributions

This chapter is published in slightly modified form in the *Astrophysical Journal*, Volume 873, Issue 1, p.34–51 with the title, “Modeling the Connection between Subhalos and Satellites in Milky Way–like

Systems,” on which I am the corresponding author. It is the result of a collaboration with Yao-Yuan Mao, Gregory Green, and Risa Wechsler. Yao and Risa were involved in the development of the model and the interpretation of the results, and Greg contributed to the development of the statistical framework and constraints. In addition, Yao, Greg, and Risa made editorial contributions to the text. This work matured during the 2019 Kavli Institute for Theoretical Physics program “The Small-scale Structure of Cold(?) Dark Matter,” during which conference participants made helpful suggestions.

3.2 Introduction

Since the turn of the century, high-resolution N -body simulations have convincingly demonstrated that structure formation in a Lambda–Cold Dark Matter (Λ CDM) universe results in a significant number of self-gravitating DM subhalos that reside within the virial radius of Milky Way (MW)–mass host halos (e.g., [266, 325]). Relating these subhalos to observed satellite galaxies in MW-like systems requires either empirical modeling to statistically associate satellites with subhalos (e.g., [91, 185]), semi-analytic galaxy formation modeling (e.g., [290, 310, 211, 436, 209, 302]), or hydrodynamic simulations (e.g., [365, 408, 478]), all of which can yield satellite populations that are consistent with the luminosity function of the brightest MW satellites.

Modeling additional aspects of observed satellites such as their spatial, orbital, and size distributions will be necessary in order to interpret the results of current and future satellite searches in a cosmological context. Zoom-in simulations of MW-mass host halos are well-suited to this task because they provide high-fidelity realizations of the subhalo populations in these systems (e.g., [433, 185, 315, 206]). However, modeling the corresponding satellite populations using hydrodynamics is difficult because resolving MW-mass hosts along with their faint satellites requires exceptional resolution [481, 478]. In addition, sub-grid baryonic physics and star formation models, which remain uncertain, can have a significant impact on galaxy formation in this regime [280, 330]. While semi-analytic models offer flexible galaxy formation prescriptions that can be extended to ultra-faint systems, these approaches yield insights into the detailed astrophysical nature of the subhalo–satellite connection, rather than offering an easily interpretable coarse-grained description. The additional layer of modeling needed to track galaxy properties over time also increases the complexity of semi-analytic models relative to empirical approaches.

Predicting satellite populations directly from subhalo populations in dark matter-only (DMO) zoom-in simulations is therefore an attractive alternative. Several authors [440, 273, 277, 214, 264, 241, 340] have taken this approach to estimate the total number of MW satellites or to constrain the connection between subhalos and satellites. Many of these studies focus on specific aspects of MW satellite modeling (e.g., correcting for the completeness of observed satellite populations) and apply several distinct models to bracket the range of Λ CDM predictions.

Ingredient	Assumptions	Parameterization	Free?
3.4.1 Host Halo	Fixed by zoom-in simulations	None	No
3.4.2 Luminosities	Abundance match to GAMA survey Extrapolate luminosity function Lognormal ($M_V V_{\text{peak}}$) distribution No satellites below M_{peak} threshold	Non-parametric Faint-end slope α Constant scatter σ_M Cut on $M_{\text{peak}} < \mathcal{M}_{\text{min}}$	No Yes Yes Yes
3.4.3 Locations	On-sky positions set by subhalos Distances set by scaled subhalo radii	None $r_{\text{sat}} \equiv \chi r_{\text{sub}}$	No No
3.4.4 Sizes	[244] sizes at accretion Size reduction set by stripping Lognormal ($r'_{1/2} R_{\text{vir}}$) distribution	$r_{1/2} \equiv \mathcal{A} (c/10)^\gamma R_{\text{vir}}$ $r'_{1/2} \equiv r_{1/2} (V_{\text{max}}/V_{\text{acc}})^\beta$ Constant scatter σ_R	No No No
3.4.5 Baryons	[337] disruption model	$p_{\text{disrupt}} \rightarrow p_{\text{disrupt}}^{1/\beta}$	Yes
3.4.6 Orphans	Correspond to disrupted subhalos NFW host + dynamical friction Stripping after pericentric passages p_{disrupt} set by time since accretion	None $\ln \Lambda = -\ln(m_{\text{sub}}/M_{\text{host}})$ $\dot{m}_{\text{sub}} \sim -\frac{m_{\text{sub}}}{\tau_{\text{dyn}}} \left(\frac{m_{\text{sub}}}{M_{\text{host}}}\right)^{0.07}$ $p_{\text{disrupt}} \equiv (1 - a_{\text{acc}})^{\mathcal{O}}$	No No No No

Table 3.1: Summary of the physical ingredients, underlying assumptions, and parameterizations of various processes that enter our model for the subhalo–satellite connection. The final column indicates whether each component of the model is held fixed or allowed to vary for our fit to the luminosity distribution of classical and SDSS-identified satellites in Section 3.5. Bold values correspond to parameters that are varied in our fit to the observed luminosity function.

Herein we build upon these efforts to develop a comprehensive framework that simultaneously addresses the relevant observational and modeling uncertainties. In particular, we present a flexible model for mapping subhalos to satellites and we use observations to infer the connection between these systems. Our model addresses:

- (i) The influence of host halo properties on satellite populations;
- (ii) Satellite luminosities — including the impact of reionization on galaxy formation — by extrapolating an abundance matching relation to faint systems and imposing a galaxy formation threshold;
- (iii) The relationship between subhalo and satellite locations
- (iv) Satellites sizes — including the effects of tidal stripping — by extrapolating a galaxy size–halo size relation that accurately describes galaxy sizes in hydrodynamic simulations
- (v) Satellite and subhalo disruption due to baryonic effects using an algorithm calibrated on hydrodynamic zoom-in simulations of MW-mass host halos
- (vi) Orphan satellites using a semi-analytic model to track and reinsert disrupted subhalos

As an example application, we fit our model to the luminosity distribution of both classical MW satellites and those discovered in the Sloan Digital Sky Survey (SDSS; [21]), and we show that it predicts satellite populations that are qualitatively and quantitatively consistent with the

luminosity function, radial distribution, and size distribution of these systems. We then forward-model the total population of MW satellites and predict satellite abundance as a function of absolute magnitude and limiting observable surface brightness; these predictions are relevant to upcoming satellite searches with improved surface brightness limits that will be carried out by surveys like the Dark Energy Survey¹ (DES; [130]), the Hyper Suprime-cam Subaru Strategic Program (HSC-SSP; [220])², and the Vera C. Rubin Observatory Legacy Survey of Space and Time³ (LSST; [301]). We argue that satellite sizes, which have not consistently been included in subhalo-based models, are a key ingredient for interpreting current and future observations, as satellite detectability is highly dependent on surface brightness.

This paper is organized as follows. In Section 3.3, we describe the zoom-in simulations used in this work. We present our model for connecting subhalos to satellites in Section 3.4. In Section 3.5, we qualitatively and quantitatively compare our mock satellite populations to classical and SDSS-identified MW satellites. We describe an improved procedure for fitting the model to observed satellite populations, which we implement by performing mock observations and comparing our predictions to the luminosity distribution of classical and SDSS-identified systems. We discuss our results, predictions for ongoing and future surveys, implications for the low-mass subhalo–satellite connection, and caveats in Section 3.6. We summarize our model in Section 3.7. Throughout, we refer to bound DM systems as “subhalos” and to luminous galaxies as “satellites.” Furthermore, “log” refers to the base-10 logarithm.

3.3 Simulations

We primarily use six “MW-like” host halos (defined below) from the suite of forty-five MW-mass DMO zoom-in simulations presented in [315]. These forty-five host halos have virial masses⁴ in the range $M_{\text{vir}} = 10^{12.1 \pm 0.03} M_{\odot}$ and have a range of formation histories that are representative of $10^{12} M_{\odot}$ hosts. The highest-resolution particles in these simulations have a mass of $3.0 \times 10^5 M_{\odot} h^{-1}$, and the softening length in the highest-resolution regions is $170 \text{ pc } h^{-1}$. To test for convergence, we compare the subhalo maximum circular velocity function, radial distribution, and size distribution for one of these hosts (Halo 937) to a resimulation with a $4.0 \times 10^4 M_{\odot} h^{-1}$ high-resolution particle mass and an $85 \text{ pc } h^{-1}$ minimum softening length. These subhalo statistics are reasonably consistent among the fiducial- and high-resolution simulations (for example, see Figure 3.3), although we find a larger population of small-virial radius subhalos in the high-resolution run.

Halo catalogs and merger trees were generated using the ROCKSTAR halo finder and the CONSISTENT-TREES merger code [47, 48]. [315] estimate that subhalos in these simulations are well-resolved down

¹<https://www.darkenergysurvey.org/>

²<https://hsc.mtk.nao.ac.jp/ssp/>

³<https://www.lsst.org/>

⁴We define virial quantities according to the overdensity $\Delta_{\text{vir}} \simeq 99.2$ as appropriate for the cosmological parameters used in our zoom-in simulations: $h = 0.7$, $\Omega_m = 0.286$, $\Omega_b = 0.047$, and $\Omega_\Lambda = 0.714$.

to a maximum circular velocity of $V_{\max} \approx 9 \text{ km s}^{-1}$. To be conservative, we restrict our analysis to subhalos with both $V_{\max} > 9 \text{ km s}^{-1}$ and peak maximum circular velocity $V_{\text{peak}} > 10 \text{ km s}^{-1}$.

The MW might be an outlier compared to typical host halos and galaxies in the relevant mass ranges (e.g., [86, 100, 393, 293]), and its subhalo and satellite populations might be particularly influenced by the existence of the Large and Small Magellanic Clouds (LMC, SMC; e.g., [302, 145]). In this work, we therefore select hosts that have two Magellanic Cloud analogs, defined as subhalos with $V_{\max} > 55 \text{ km s}^{-1}$ following [302]; we find six such MW-like host halos in our simulation suite. Although secondary properties of the MW in addition to the existence of the Magellanic Clouds could also bias its subhalo population [173], we have checked that the results presented in Section 3.6 are not significantly affected if we randomly select host halos from our simulation suite.

3.4 Model

The following subsections describe the ingredients that enter our model for the subhalo–satellite connection, which we summarize in Table 3.1. Our model encompasses the influence of host halo properties on satellite populations (Section 3.4.1), the way in which satellites populate subhalos and the relationship between satellite and subhalo properties (Sections 3.4.2–3.4.4), and modifications to subhalo populations in DMO simulations (Sections 3.4.5, 3.4.6). We indicate the parameters associated with our method in the subsection headers; these parameters define our physically motivated empirical model. Note that unlike in semi-analytic approaches, we do not require our model components to represent specific astrophysical processes such as star formation or quenching. In addition, although we illustrate our model using the zoom-in simulations described previously, we stress that our framework does not depend on the specific host halos used in this paper.

3.4.1 Host Halo Properties

Although we primarily use the six MW-like host halos described in Section 3.3, we note that our model can be applied to different hosts by changing the underlying zoom-in simulations. While the masses of our host halos are consistent with several observational constraints for the MW (e.g., [100, 354]; also see the review in [63]), recent studies based on astrometric data suggest a more massive MW halo ([324, 370, 424, 471]; however, see [102]). Thus, studying how our predictions vary as a function of host halo mass is an important avenue for future work.

At fixed host halo mass, subhalo abundance depends on the host’s formation history and secondary properties (e.g., [496, 173]). For example, lower-concentration hosts accrete the majority of their subhalos later than higher-concentration hosts, leaving less time for subhalo disruption and resulting in a larger population of surviving subhalos at fixed host mass [315]. Although our MW-like host halos have lower concentration than average 10^{12} M_\odot hosts, marginalizing over concentration using MW-mass hosts with a range of formation histories does not affect the results of

the fit presented herein.

3.4.2 Satellite Luminosities

Next, we describe our procedure for assigning satellite luminosities to DM subhalos. We use abundance matching to link subhalos’ peak maximum circular velocities V_{peak} to their satellites’ absolute r -band magnitudes M_r down to a certain magnitude. In particular, we follow [193] by tuning our abundance matching relation to the GAMA galaxy survey [295] using the measured GAMA r -band luminosity function k -corrected to $z = 0$ and the halo V_{peak} function from a DMO simulation with 2560^3 particles and a side length of $250 \text{ Mpc } h^{-1}$, run using the same cosmological parameters and code as the Dark Sky Simulations [428]. We abundance match to the GAMA luminosity function down to $M_r = -13$ mag, and we extrapolate this relation to fainter galaxies, assuming that the faint-end satellite luminosity function follows a power law. We then convert our predicted r -band magnitudes to V -band magnitudes using the empirical relation $M_r \approx M_V - 0.2$ mag as in [193]. For a particular faint-end luminosity function slope, this procedure yields a mean $M_V - V_{\text{peak}}$ relation. Note that by assigning magnitudes based on V_{peak} , we have implicitly assumed that satellites’ absolute magnitudes are not altered after accretion (see [358] for a detailed discussion of this assumption; also see [491, 479, 376]). We now discuss our method for varying the faint-end slope and applying scatter to our luminosity relation.

Faint-end Slope (α)

The faint-end slope α of the luminosity function that enters our extrapolated abundance matching relation has been examined in previous MW satellite studies including [440]. Constraints on the faint-end slope for dwarf galaxies are limited, but can in principle be derived from the total luminosity function in the Local Volume [265] or from the satellite distributions around the MW, Andromeda (M31), or MW analogs, such as those observed by the Satellites around Galactic Analogs Survey (SAGA; [193]). For reference, [440] infer values of $-2 \lesssim \alpha \lesssim -1$ by applying completeness corrections to observed MW satellite populations. We find that a characteristic value of $\alpha = -1.3$ maps subhalos at our $V_{\text{peak}} \approx 10 \text{ km s}^{-1}$ resolution limit to satellites with $M_V \approx 5$ mag, which is significantly dimmer than the faintest spectroscopically confirmed MW satellite, Segue I ($M_V = -1.5$ mag; [321]).

Luminosity Scatter (σ_M)

Galaxy properties such as stellar mass and luminosity derived from abundance matching are usually assumed to follow lognormal distributions at fixed halo mass proxy, and the scatter in these properties is relatively well constrained to ~ 0.2 dex for host halos with masses above $\sim 10^{12} \text{ M}_\odot$ ([45, 379, 288];

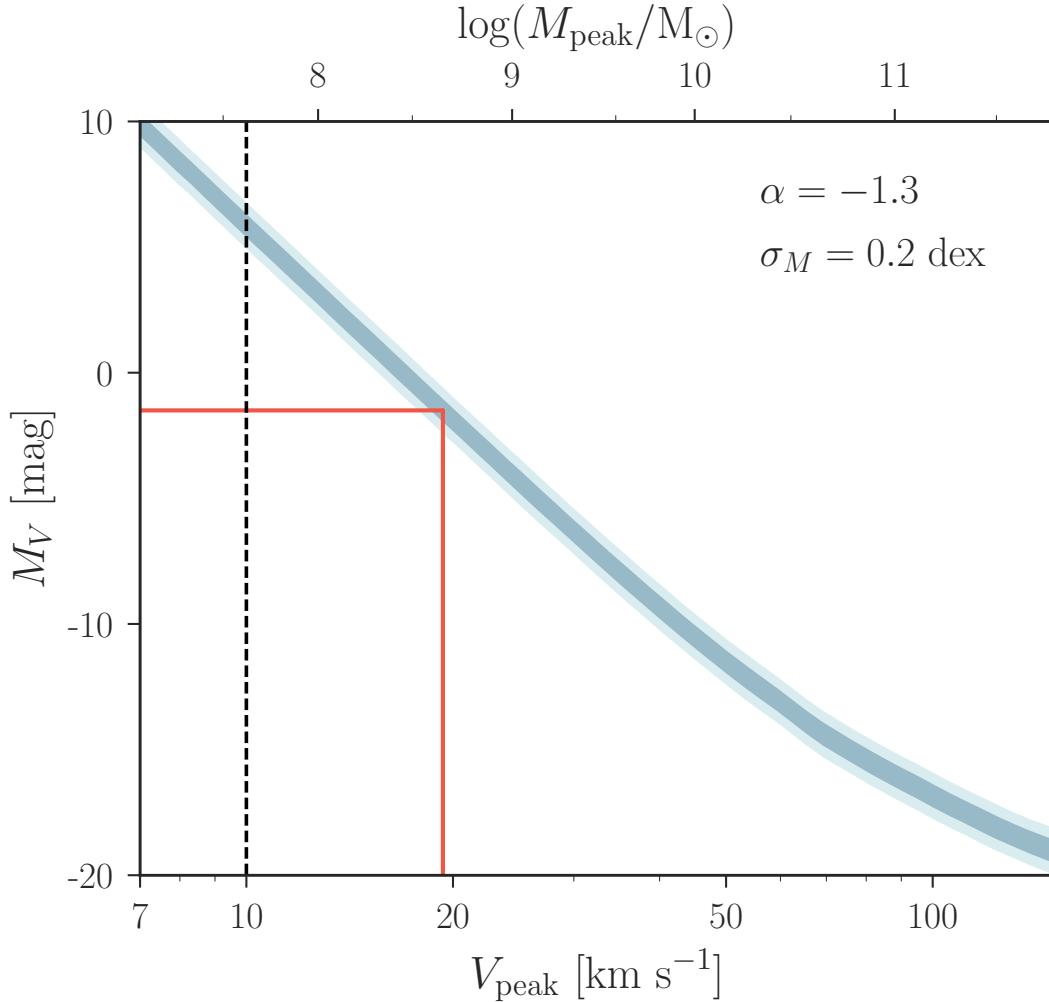


Figure 3.1: Example of the relation between satellite luminosity and subhalo peak circular velocity used in our model. We treat the faint-end power-law luminosity function slope α and the constant lognormal scatter σ_M as free parameters. Dark (light) shaded areas show $\pm 1\sigma_M$ ($\pm 2\sigma_M$) scatter. The dashed black line indicates a conservative upper bound on the resolution limit of our simulations, and the red lines indicate the absolute magnitude of the faintest spectroscopically confirmed MW satellite (Segue I; $M_V = -1.5$ mag) and the mean inferred V_{peak} of its subhalo for $\alpha = -1.3$ and $\sigma_M = 0.2$ dex.

also see the review in [474]).⁵ However, this scatter is not well constrained in the MW satellite regime [187, 328]. We therefore treat the scatter in our predicted absolute magnitudes at fixed V_{peak} , which we denote σ_M , as a free parameter with a lower bound of 0.2 dex. In particular, we apply scatter to each satellite’s absolute magnitude by drawing from a lognormal distribution with a mean set by our M_V – V_{peak} relation and a standard deviation of σ_M . Note that we do not deconvolve the abundance matching relation when we vary the scatter for computational efficiency; we have checked that this choice does not significantly affect the resulting M_V – V_{peak} relation. Although several semi-analytic models (e.g., [209]) and hydrodynamic simulations (e.g., [410, 176]) suggest that the scatter between galaxy properties and halo properties grows with decreasing halo mass, we treat σ_M as a constant for the fit presented in this paper. It might be necessary to relax this assumption in future work that addresses a more complete population of observed satellites.

Because we extrapolate our abundance matching relation to systems fainter than $M_r = -13$ mag, our procedure is not designed to match the global luminosity function in this regime; however, the global constraint is not our primary concern, and our power-law-plus-scatter parameterization results in a more flexible model. We show an example M_V – V_{peak} relation in Figure 3.1 for fiducial choices of α and σ_M , but we emphasize that these are free parameters in our model. Figure 3.1 illustrates that our simulations should resolve all subhalos that host satellites with $M_V < 0$ mag for this choice of α and σ_M (however, see the discussion on artificial subhalo disruption and orphan satellites in Section 3.4.6).

Galaxy Formation Threshold (\mathcal{M}_{min})

Many authors have studied the impact of reionization on galaxy formation in low-mass subhalos, finding that subhalos below a certain mass or V_{peak} threshold are likely to be dark (e.g., [439, 98, 429]). Because we simply assign satellite luminosities to subhalos, we must account for this effect; however, the details of the galaxy formation process for faint satellites are unclear. For example, [280], [409], and [176] respectively find that most isolated halos with $M_{\text{peak}} \approx 10^{10} \text{ M}_\odot$ ($V_{\text{peak}} \approx 45 \text{ km s}^{-1}$), $3 \times 10^9 \text{ M}_\odot$ ($V_{\text{peak}} \approx 33 \text{ km s}^{-1}$), and 10^9 M_\odot ($V_{\text{peak}} \approx 25 \text{ km s}^{-1}$) are dark. Meanwhile, [241] find that the peak virial mass of the subhalo hosting Segue I ($M_V = -1.5$ mag) is below $2.4 \times 10^8 \text{ M}_\odot$ ($V_{\text{peak}} \approx 16 \text{ km s}^{-1}$) at the 68% confidence level, and [64] find that star formation can proceed after supernova feedback in subhalos with peak virial masses down to $\sim 10^7 \text{ M}_\odot$ ($V_{\text{peak}} \approx 7 \text{ km s}^{-1}$).

These results are generally sensitive to the assumed redshift of reionization and ultraviolet background; to account for these uncertainties in a simple way, we treat the minimum peak virial mass necessary for galaxy formation, \mathcal{M}_{min} , as a free parameter. In particular, for a given value of \mathcal{M}_{min} , we discard satellites in all subhalos with $M_{\text{peak}} < \mathcal{M}_{\text{min}}$. This mass cut clearly does not capture

⁵The observational uncertainties associated with the conversion from stellar mass to luminosity are likely larger than the difference in scatter between these quantities [474], so we ignore this distinction here.

the complexities of galaxy formation in low-mass subhalos, which likely result in a smoothly varying galaxy occupation fraction rather than a sharp cutoff [409, 410, 175]. Although it is not *necessary* to introduce \mathcal{M}_{\min} because we fit to an incomplete sample of observed satellites, we will demonstrate that the classical-plus-SDSS luminosity distribution sets an interesting upper bound on this quantity, given our simple parameterization.

3.4.3 Satellite Locations

Next, we describe our procedure for assigning on-sky coordinates and radial distances to our mock satellites.

On-sky Positions

In general, we expect the positions of satellite galaxies to correspond reasonably well to the positions of their subhalos. Thus, we simply use the projected on-sky positions of subhalos in our zoom-in simulations to assign on-sky coordinates to our satellites. Since our DMO simulations do not contain galactic disks, we are free to perform arbitrary 3D rotations of our subhalo positions about host halo centers before projecting them onto the sky; these rotations can be fixed based on the positions of Magellanic Cloud–like subhalos in order to perform realistic mock MW satellite surveys. In addition, we convert satellites’ Galactocentric coordinates to heliocentric coordinates by placing mock observers 8 kpc from our host halo centers, as described in Section 3.5.

Radial Scaling (χ)

MW satellites seem to be unusually centrally concentrated compared to both the observed satellite population in M31 and typical subhalo populations in zoom-in simulations (e.g., [492, 202]; however, see [289]). While this apparent discrepancy might be caused in part by misestimates of observational incompleteness, several numerical effects could contribute to a mismatch between the radial distribution of simulated subhalos and observed satellites. For example, simulations might underestimate the amount of dynamical friction experienced by subhalos due to resolution effects (although we expect this to be a subdominant source of error in MW-mass systems), and halo finders might mis-track or fail to identify subhalos in dense, central regions [267]. We expect halo finder incompleteness to be mitigated in our analysis, given that we use the phase-space-based halo finder ROCKSTAR and because we reinsert disrupted subhalos using the orphan model described in Section 3.4.6.

To model these potential biases in our radial satellite distributions, we define the parameter $\chi \in (0, 1]$ as follows:

$$r_{\text{sat}} \equiv \chi r_{\text{sub}}, \quad (3.1)$$

where r_{sat} is a satellite’s distance from the center of its host halo (which we identify with its

Galactocentric distance) and r_{sub} is the Galactocentric distance of the corresponding subhalo. Thus, $\chi = 1$ corresponds to setting each satellite’s radial distance equal to that of its subhalo, while smaller values of χ shift our mock satellites inward relative to their subhalos. Although we fix $\chi = 1$ for the fit to the luminosity distribution of classical and SDSS-identified satellites presented below, we include it in our modeling framework for generality and as a useful phenomenological parameter for future work.

3.4.4 Satellite Sizes

To assign sizes to our mock satellites, we use a modified version of the galaxy size–halo virial radius relation from [244], which relates a galaxy’s 3D half-mass radius to its subhalo’s virial radius R_{vir} . Since we will compare the sizes of our mock satellites to measured half-light radii, we simply identify our satellites’ predicted 3D half-mass radii with their projected 2D half-light radii. This conversion neglects mass-to-light weighting and projection effects; these are both reasonable approximations, although the latter overestimates the sizes of highly elliptical dwarfs.

Mean Size Relation (\mathcal{A}, γ)

We use the following relation from [244] to set satellite half-light radii at the time of accretion:

$$r_{1/2} \equiv \mathcal{A} \left(\frac{c}{10} \right)^\gamma R_{\text{vir}}, \quad (3.2)$$

where $\mathcal{A} = 0.02$, $\gamma = -0.7$, and c denotes subhalo concentration measured at accretion.⁶ [244] find that this relation yields galaxy sizes that are consistent with those found in two hydrodynamic simulations with a residual scatter of ~ 0.15 dex about the hydrodynamic results. Although Equation 3.2 is essentially untested for ultra-faint dwarf galaxies, we hold \mathcal{A} and γ fixed at their fiducial values, and we account for this uncertainty by allowing the amount of size reduction that satellites undergo and the scatter in this relation to vary as described in the following subsections.

Size Reduction Due to Tidal Stripping (β)

One might expect the sizes of satellite galaxies to correlate tightly with their subhalos’ virial radii at accretion. However, dynamical effects such as tidal stripping can alter subhalo and satellite sizes after accretion, introducing scatter in this relationship (e.g., [276, 358]). To model these effects, we calculate each satellite’s $z = 0$ half-light radius $r'_{1/2}$ as follows:

$$r'_{1/2} \equiv r_{1/2} \left(\frac{V_{\text{max}}}{V_{\text{acc}}} \right)^\beta, \quad (3.3)$$

⁶We have also tested the size model from [275], but we find that the concentration dependence in the [244] relation leads to a more reasonable range of sizes compared to observed MW satellites.

where V_{\max} denotes the corresponding subhalo’s maximum circular velocity at $z = 0$, V_{acc} denotes its maximum circular velocity at accretion, and $\beta \geq 0$ is a model parameter. Thus, $\beta = 0$ corresponds to no change in satellite sizes after accretion, while larger values of β increase the amount by which satellite sizes are reduced based on the degree of tidal stripping that their subhalos undergo. This prescription does not capture the effects of tidal heating, which can *enlarge* satellites rather than reduce their sizes. In addition, although V_{\max}/V_{acc} depends on the orbital history of each subhalo, the amount of tidal stripping that their satellites undergo might also depend on orbital inclination with respect to the central disk, an effect that is not modeled in our DMO simulations. Thus, Equation 3.3 will likely need to be recalibrated on the size evolution of ultra-faint satellites in upcoming hydrodynamic simulations. Pending such results and a detailed model for stellar mass stripping, we fix $\beta = 0$ for our fiducial model, and we demonstrate in Appendix B.2 that our constraints are insensitive to this parameter.

Figure 3.2 compares the mean size distribution at accretion ($\beta = 0$) and at $z = 0$ using our tidal stripping model ($\beta = 1$) for mock satellites in our six MW-like host halos. Satellite sizes are reduced due to tidal stripping, as expected from Equation 3.3. Our extrapolation of the [244] relation yields a reasonable range of sizes compared to observed MW satellites, which roughly span $10 \text{ pc} \lesssim r_{1/2} \lesssim 3000 \text{ pc}$.

Size Scatter (σ_R)

The size relation that we have described has not been tested against hydrodynamic simulations for galaxies with half-light radii smaller than $\sim 400 \text{ pc}$. Thus, for small systems, the uncertainty in this relation might deviate from the ~ 0.15 dex residuals found in [244]. We therefore apply scatter by drawing each satellite’s size from a lognormal distribution with a mean given by Equation 3.3 and a standard deviation of σ_R , and we impose a minimum size of 20 pc as appropriate for the classical and SDSS-identified satellites studied herein. Although this scatter might be size-dependent, we treat σ_R as a constant for simplicity.

3.4.5 Subhalo Disruption due to Baryonic Effects (\mathcal{B})

To model the enhancement in subhalo disruption due to baryonic effects such as the presence of a central galactic disk, we apply the subhalo disruption algorithm from Chapter 2 [337], which uses the orbital and internal properties of subhalos in DMO simulations to predict the probability that they will be disrupted in hydrodynamic resimulations. This model was trained on two hydrodynamic simulations of MW-mass host halos from the Feedback In Realistic Environments project (FIRE; [223, 224]), both of which have classical satellite populations that are reasonably consistent with those in the MW and M31 [478, 190]. [337] found that this disruption model predicts surviving subhalo populations that are in better agreement with hydrodynamic results than DMO simulations that include the gravitational potential of the galactic disk found in corresponding hydrodynamic

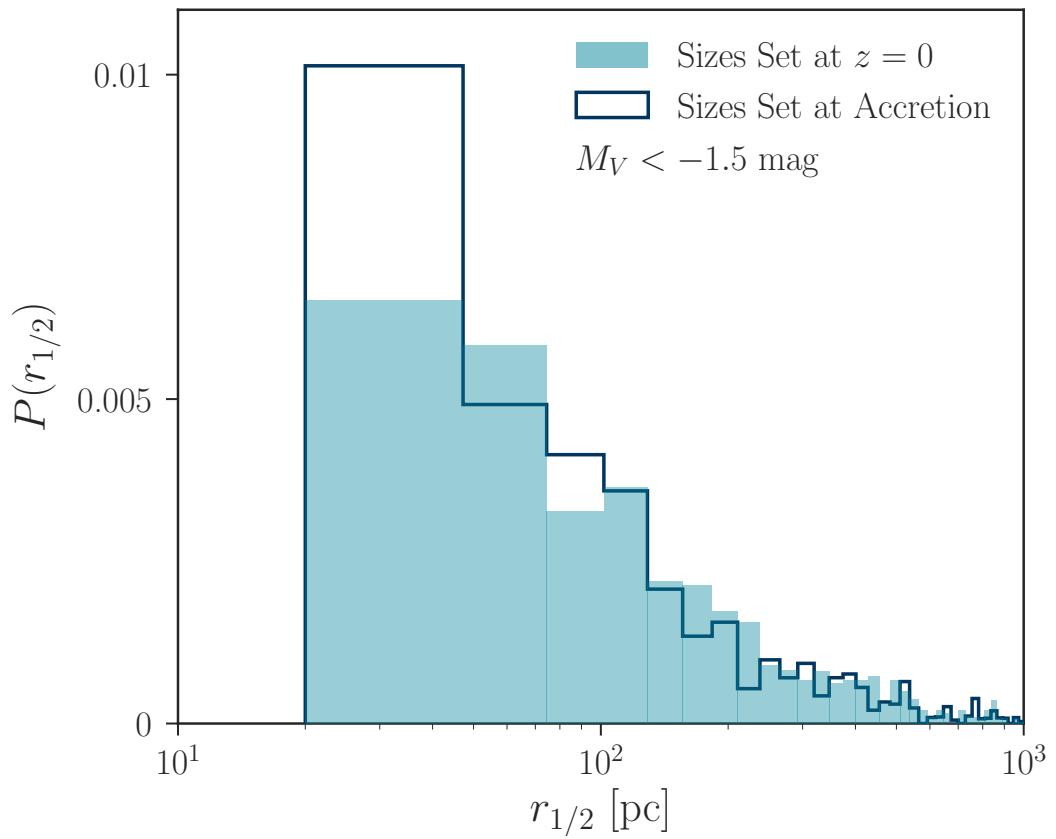


Figure 3.2: Size distributions for mock satellites in our six MW-like host halos, with satellite sizes set by subhalo sizes at accretion ($\beta = 0$) and at $z = 0$ using our prescription for satellite size reduction due to tidal stripping ($\beta = 1$). The distributions are weighted by survival probability using our subhalo disruption model with $\mathcal{B} = 1$.

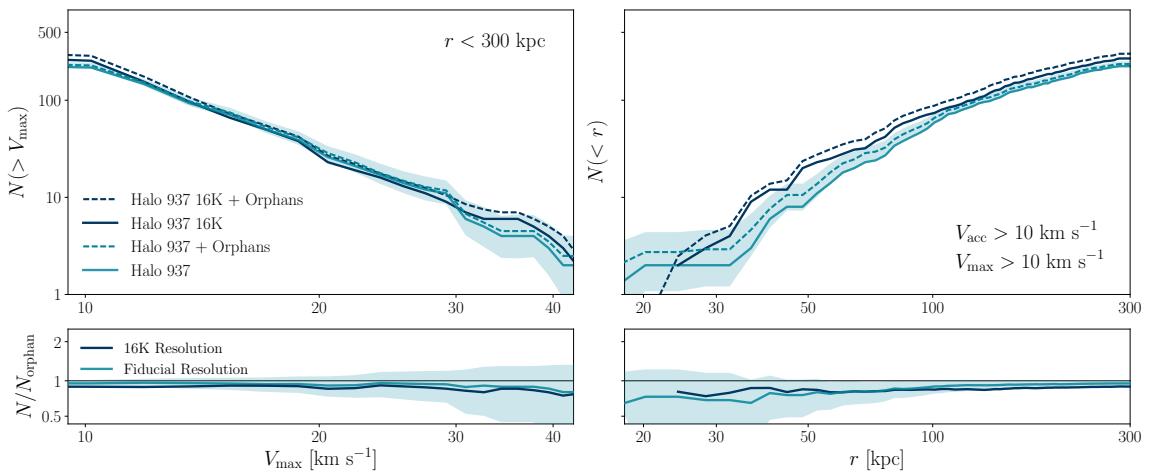


Figure 3.3: V_{\max} function (left panel) and radial distribution (right panel) of subhalos with $V_{\text{acc}} > 10 \text{ km s}^{-1}$ and $V_{\max} > 10 \text{ km s}^{-1}$ in one of our zoom-in simulations, shown with (dashed lines) and without (solid lines) our fiducial orphan model ($\mathcal{O} = 1$). The dark blue lines show results from a high-resolution simulation of this host halo, and the blue bands show the Poisson scatter for our prediction that includes all disrupted subhalos tracked to $z = 0$ in the fiducial-resolution simulation. The bottom panels show the ratio of the number of subhalos to the number of subhalos in the corresponding simulation including orphans for our fiducial-resolution (blue) and high-resolution (dark blue) runs. Including orphans brings the fiducial-resolution V_{\max} function and radial distribution into fairly good agreement with the high-resolution results.

simulations. Our method therefore captures both the tidal influence of a central disk and additional baryonic effects, so our work differs from studies that employ DM-plus-disk simulations to model the effects of baryons on subhalo populations (e.g., [241, 340]).

[337] showed that baryonic subhalo disruption approximately rescales the subhalo V_{\max} functions for the zoom-in simulations used in this work. However, the radial subhalo distributions are not scaled by a constant factor due to enhanced subhalo disruption in the inner ~ 50 kpc of host halos caused by the dynamical influence of the central galactic disks in the FIRE simulations. To account for the uncertainty associated with the limited training set in [337] — and particularly the limited set of host halo *accretion histories* that the model is trained on — we define the parameter \mathcal{B} to characterize subhalo disruption due to baryonic effects as follows. We assign each surviving subhalo in our DMO simulations a disruption probability by modifying the prediction from the [337] model according to $p_{\text{disrupt}} \rightarrow p_{\text{disrupt}}^{1/\mathcal{B}}$. Thus, $\mathcal{B} = 1$ corresponds to the unaltered disruption probabilities, $\mathcal{B} > 1$ ($\mathcal{B} < 1$) corresponds to increased (decreased) disruption probabilities relative to our fiducial model, and $\mathcal{B} = 0$ corresponds to no additional subhalo disruption due to baryonic effects. We assume that each satellite’s disruption probability is equal to that of its subhalo. Although this assumption is reasonable for disruption mechanisms like disk shocking, it warrants further investigation using controlled simulations.

3.4.6 Orphan Satellites (\mathcal{O})

Cosmological simulations often require a population of orphan galaxies (i.e., galaxies whose halos are not detected by the halo finder) to match observational galaxy clustering constraints (e.g., [468, 210, 372]). The details of orphan modeling depend on the simulation and target galaxy population in question; however, orphans should generally be included in analyses that are sensitive to systems near a resolution threshold. Thus, despite the relatively high resolution of our zoom-in simulations, orphans are potentially important in our study, since we aim to model faint MW satellites. Moreover, if *artificial* subhalo disruption is a significant effect (e.g., [453, 454]), in the sense that disrupted subhalos in our simulations should host observable satellite galaxies, then it becomes even more important to include orphans.

To model orphans, we therefore identify all disrupted subhalos in each simulation that contribute directly to the main host halo. We track the orbit of each disrupted subhalo until $z = 0$ using a softened gravitational force law and dynamical friction as follows:

$$\dot{\mathbf{v}} = -\frac{GM(< r)}{(r + \epsilon R_{\text{vir,host}})^2} \hat{\mathbf{r}} + \frac{\mathbf{F}_{\text{df}}}{m_{\text{sub}}}. \quad (3.4)$$

Here, r is a subhalo’s distance to the center of the host, m_{sub} is its virial mass, $M(< r)$ is the enclosed host halo mass⁷, $R_{\text{vir,host}}$ is the host halo’s virial radius, and the gravitational softening

⁷We calculate $M(< r)$ assuming an NFW host halo density profile—that is, $M(< r) = M_{\text{host}} f(cx)/f(c)$, where

$\epsilon = 0.01$ is chosen to avoid hard collisions with the host (we have checked that the resulting subhalo orbits are insensitive to this choice for reasonably small values of ϵ). To calculate \mathbf{F}_{df} , we use the [110] dynamical friction formula for an NFW host halo and a Maxwellian distribution of host particle velocities, which yields

$$\mathbf{F}_{\text{df}} = -4\pi \left(\frac{Gm_{\text{sub}}}{|\mathbf{v}_{\text{orb}}|} \right)^2 \ln \Lambda \rho(r) \left[\text{erf}(X) - \frac{2X}{\sqrt{\pi}} e^{-X^2} \right] \frac{\mathbf{v}_{\text{orb}}}{|\mathbf{v}_{\text{orb}}|}, \quad (3.5)$$

where \mathbf{v}_{orb} denotes subhalo orbital velocity, $\ln \Lambda$ is the Coulomb logarithm, $\rho(r)$ is the host halo's density profile, and $X \equiv \mathbf{v}_{\text{orb}}/[\sqrt{2}\sigma(r)]$ where $\sigma(r)$ is the local host halo velocity dispersion. We estimate $\sigma(r)$ using the fitting formula in [497] and we set $\ln \Lambda = -\ln(m_{\text{sub}}/M_{\text{host}})$ following [184].

To account for tidal stripping, we follow [44] by modeling mass loss for disrupted subhalos as follows:

$$\dot{m}_{\text{sub,infalling}} = 0 \quad (3.6)$$

$$\dot{m}_{\text{sub,outgoing}} = -1.18 \frac{m_{\text{sub}}}{\tau_{\text{dyn}}} \left(\frac{m_{\text{sub}}}{M_{\text{host}}} \right)^{0.07}, \quad (3.7)$$

where $\tau_{\text{dyn}} = (4\pi G \rho_{\text{vir}}/3)^{-1/2}$ is the dynamical timescale and derivatives are taken with respect to time. This mass stripping model is motivated by several synthetic and cosmological tests that have shown that the majority of subhalo mass loss occurs after pericentric passages [267, 46]. We use a modified version of the fitting formula from [243] to model the corresponding reduction in each disrupted subhalo's maximum circular velocity,

$$\frac{d \log V_{\text{max}}}{d \log m_{\text{sub}}} = 0.3 - 0.4 \frac{m_{\text{sub}}}{m_{\text{sub}} + m_{\text{sub,acc}}}, \quad (3.8)$$

where $m_{\text{sub,acc}}$ denotes subhalo virial mass at accretion. Finally, we calculate the sizes of our orphan satellites using Equations 3.2–3.3.

To vary the contribution from orphan satellites, we define the parameter \mathcal{O} by setting the disruption probability for each orphan equal to $(1 - a_{\text{acc}})^{\mathcal{O}}$, where a_{acc} is the final scale factor at which a subhalo enters the virial radius of its host halo. We find that this formula with $\mathcal{O} = 1$ describes the disruption probabilities predicted by the [337] model for *surviving* subhalos fairly well, and we use it because the uncertainties in the other features that enter the [337] model (e.g., pericentric distance) are potentially large for disrupted subhalos. Thus, $\mathcal{O} = 0$ corresponds to including zero orphan satellites, larger values of \mathcal{O} increase the contribution from orphans, and each orphan's disruption probability grows with the amount of time elapsed since accretion.

To test our orbit tracking and tidal stripping models, we calculate V_{max} functions and radial subhalo distributions for the fiducial- and high-resolution versions of Halo 937 described in Section

$x \equiv r/R_{\text{vir,host}}$, c is the host halo's concentration, and $f(\xi) \equiv \ln(1 + \xi) - \xi/(1 + \xi)$.

[3.3](#) by selecting subhalos with $V_{\text{acc}} > 10 \text{ km s}^{-1}$ and $V_{\text{max}} > 10 \text{ km s}^{-1}$ (recall that $V_{\text{max}} \approx 9 \text{ km s}^{-1}$ corresponds to the resolution threshold of our fiducial simulations). Figure [3.3](#) shows that our orphan model with $\mathcal{O} = 1$ brings both the velocity function and radial distribution from the fiducial-resolution simulation (light blue dashed lines) into fairly good agreement with the high-resolution results (dark blue solid lines). While the change in the velocity function due to orphans is consistent with a change in overall normalization (i.e., a constant scaling with respect to V_{max}) at both resolution levels, we find that orphans are preferentially added in the inner regions of our fiducial-resolution simulation, where they are most likely to be needed. On the other hand, our orphan model merely rescales the radial subhalo distribution in the high-resolution run, suggesting that there is less spurious disruption in this case.

We note that our orphan model can be interpreted in terms of the amount of artificial subhalo disruption that occurs in our simulations (e.g., [\[453, 454\]](#)). In particular, $\mathcal{O} = 0$ (no orphans included) corresponds to the assumption that all subhalo disruption in our simulations is both physical and coincides with the disruption of their satellites. Meanwhile, $\mathcal{O} \gg 1$ (all orphans included with zero disruption probability) implies that all subhalo disruption in our DMO simulations and all *additional* subhalo disruption due to baryonic effects in hydrodynamic simulations is artificial, in the sense that all disrupted subhalos should host satellites. We find that these extreme possibilities are respectively disfavored by the inner radial distribution and the total abundance of observed MW satellites, which effectively set lower and upper bounds on \mathcal{O} . We fix $\mathcal{O} = 1$ for the fit presented as follows, which roughly corresponds to the assumption that subhalo disruption in our DMO simulations is artificial while subhalo disruption in the hydrodynamic simulations that we train our disruption model on is physical. This assumption does not necessarily contradict the results of [\[454\]](#) and [\[453\]](#), since their subhalo disruption tests were performed without central disk potentials. Our hope is that future data will better constrain this parameter.

3.4.7 Comparison to Recent Models

Our approach differs from previous models for the subhalo–satellite connection in several regards. To illustrate these differences, we compare our model to those recently presented in [\[241\]](#) and [\[340\]](#):

1. Our procedure for assigning satellite luminosities to subhalos is tuned to match an observed luminosity function for systems brighter than $M_r = -13$ mag, unlike the empirical stellar mass–halo mass relations considered in [\[241\]](#); meanwhile, [\[340\]](#) estimate the total number of MW satellites by statistically comparing radial subhalo distributions in the AQUARIUS simulations to classical and SDSS-identified satellites without explicitly modeling the luminosity distribution of these systems.
2. We model the sizes of our mock satellites, while recent empirical studies, including [\[241\]](#) and [\[340\]](#), assume that all systems of a given luminosity have sufficient surface brightness to be

observed (though see [99]).

3. Our model for baryonic subhalo disruption is similar to the prescriptions in [241] and [340], which are based on DM-plus-disk simulations. However, our algorithm predicts surviving subhalo populations that are in better agreement with hydrodynamic results compared to DM-plus-disk simulations for the host halos that it was trained on [337].
4. We parameterize our baryonic disruption and orphan satellite models to allow for deviations from our fiducial prescriptions, unlike [241] and [340].

3.5 Comparison to Observed Satellites

We now demonstrate that our model can produce satellite populations that are both qualitatively and quantitatively consistent with classical and SDSS-identified MW satellites. We reiterate that other satellite populations — including those associated with M31 or with MW-mass hosts outside of the Local Volume — can be used to constrain our model, but we focus on the well-characterized population of classical and SDSS-identified MW satellites for clarity. Thus, we *do not* utilize all known MW satellites for the fit presented as follows.

3.5.1 Qualitative Comparison

Before fitting our model to observed MW satellites, we qualitatively compare its predictions to the abundance and properties of classical and SDSS-identified systems. In particular, Figure 3.4 shows projections of the predicted satellite population for one of our MW-like host halos in the observationally motivated parameter space of absolute magnitude, half-light radius, and heliocentric distance using fiducial values of our free parameters. We compare our predictions to the following classical and SDSS-identified MW satellites compiled in [321]: LMC, SMC, Sagittarius I, Fornax, Leo I, Sculptor, Leo II, Sextans I, Carina, Draco, and Ursa Minor (classical), and Canes Venatici I, Hercules, Boötes I, Leo IV, Ursa Major I, Leo V, Pisces II, Canes Venatici II, Ursa Major II, Coma Berenices, Willman I, Boötes II, Segue II, and Segue I (SDSS). We exclude Pisces I and Pegasus III because they were discovered using methods that do not adhere to our assumed SDSS detection criteria (described as follows), and we exclude Leo T because it lies outside of our fiducial 300 kpc reference radius. We refer the reader to [321] for references to the papers in which these systems were discovered.

Figure 3.4 illustrates all mock classical satellites in one of our host halos, along with all systems in a region corresponding to the area of the SDSS survey that points away from the LMC analog in this simulation. We define mock classical satellites as objects with $M_V \leq -8.8$ mag and we assume that observations of these systems are complete. We plot systems within 300 kpc that pass both the surface brightness limit of 30 mag arcsec $^{-2}$ estimated in [271] and the distance–magnitude SDSS

detection limit estimated in [273] as blue stars, while blue circles indicate systems in the mock SDSS footprint that *do not* pass both detection criteria. The [273] heliocentric completeness radius can be expressed as

$$R_{\text{eff}}(M_V) = 10^{-aM_V+b} \text{ kpc}, \quad (3.9)$$

where $a = 0.228$ and $b = 1.1$. Thus, a satellite with magnitude M_V passes this detection criterion if it falls within the effective radius $R_{\text{eff}}(M_V)$ given by Equation 3.9. Adopting the [465] version of the SDSS detection threshold, which corresponds to $a = 0.187$ and $b = 1.42$, yields similar constraints for the fit presented herein. For each mock satellite, we calculate absolute surface brightness using the relation

$$\mu_V = M_V + 36.57 + 2.5 \log[2\pi(r'_{1/2}/1 \text{ kpc})^2], \quad (3.10)$$

where we have left the units of mag arcsec $^{-2}$ implicit. Together, these detection criteria depend on the absolute magnitude, size, and radial distance of each mock satellite.

Figure 3.4 demonstrates that our predicted satellite populations agree fairly well with the abundance and properties of classical and SDSS-identified systems; we discuss the apparent deficit of satellites in the inner regions as follows. Although Figure 3.4 shows a particular realization of our model for a single MW-like host halo, we have checked that this mock satellite population is representative of our predictions for classical and SDSS systems. Our predicted satellite populations are similar to those inferred from semi-analytic models that account for satellite sizes, including [290], although we generally find fewer bright, compact systems than these works. Interestingly, the population of undiscovered mock satellites that do not pass the SDSS detection criteria depends sensitively on their size distribution. In particular, non-observations above the SDSS detection threshold in the M_V-r_\odot plane would be counted as detections in models that assume all dwarfs have sufficient surface brightness to be detected. This result suggests that size modeling will play an important role in interpreting current and future MW satellite observations. In addition, Figure 3.4 shows that the number of systems that do not pass the SDSS detection thresholds depends strongly on \mathcal{M}_{\min} , which implies that the observed abundance and properties of MW satellites can be used to place upper limits on the masses of the subhalos that host faint systems (e.g., [202, 241]).

Our fiducial model underpredicts the number of observed satellites at small radii ($r_\odot \lesssim 50$ kpc) relative to SDSS observations. Decreasing the strength of baryonic subhalo disruption does not directly resolve this issue, since we predict too few surviving subhalos in this regime regardless of their disruption probabilities, which can be seen from Figure 3.4. However, our simulations are likely subject to spurious subhalo disruption and halo finder issues for subhalos in central regions, so it is unclear whether the difference among the predicted and observed radial satellite distributions reflects a physical shortcoming of our model or a numerical shortcoming of our simulations. We find that increasing the contribution from orphan satellites (i.e., increasing \mathcal{O}) alleviates the discrepancy, which hints at the latter explanation, although scaling our satellite radii inward (i.e., decreasing χ)

also reduces the tension. We revisit the effects of these parameters on the radial distribution in Section 3.6.3.

3.5.2 Quantitative Comparison

We now describe our procedure for fitting the model described in Section 3.4 to observed satellite populations; we then specialize to the luminosity distribution of classical and SDSS-identified MW satellites. For each zoom-in simulation and each realization of our satellite model, we generate a predicted satellite population by performing a mock survey with an appropriate footprint and detection efficiency. We bin these satellites according to their physical properties \mathbf{s} (e.g., absolute magnitude, heliocentric distance, and half-light radius) by discretizing the space of satellite properties into bins of volume \mathcal{V} ; for example, choosing uniform bins for the parameter space plotted in Figure 3.4 would correspond to setting $\mathcal{V} = \Delta M_V \Delta r_{\odot} \Delta r_{1/2}$, where, for example, ΔM_V denotes the width of our absolute magnitude bins. To include the effects of observational incompleteness and satellite disruption, we count each satellite as $p_{\text{detect}} \times (1 - p_{\text{disrupt}})$ — rather than one — observed system, where p_{detect} is the probability of detecting a given satellite (determined by its properties and the survey sensitivity) and p_{disrupt} is its disruption probability.

We assume that observed satellites and mock satellites populate the parameter space in question according to a multidimensional Poisson point process with a rate parameter λ that is constant in each bin. Naively, λ can be calculated by averaging the number of mock satellites in each bin:

$$\lambda_i(\boldsymbol{\theta}) = \frac{1}{\mathcal{V}} \langle \hat{n}_i(\boldsymbol{\theta}) \rangle, \quad (3.11)$$

where $\hat{n}_i(\boldsymbol{\theta})$ indicates the number of mock observed satellites in bin i , $\boldsymbol{\theta}$ denotes the set of model parameters, and the average can be taken over different zoom-in simulations and realizations of the satellite population in each simulation (including different luminosity and size model realizations, observer locations, and survey orientations). In this formulation, the likelihood of observing N satellites $\mathbf{s}_1, \dots, \mathbf{s}_N$ in the survey corresponding to these mock observations is given by

$$P(\mathbf{s}_1, \dots, \mathbf{s}_N | \boldsymbol{\theta}) \approx \exp \left[- \sum_{\text{bins } j} \lambda_j(\boldsymbol{\theta}) \mathcal{V}_j \right] \prod_{\text{bins } i} \frac{\lambda_i(\boldsymbol{\theta})^{n_i}}{n_i!}, \quad (3.12)$$

where i and j index the bins and n_i is the number of observed satellites in bin i . This expression is approximate because the integral that appears in the normalization factor for a Poisson point process is replaced by a sum over discrete bins.

However, the estimate of λ obtained from Equation 3.11 is potentially noisy because we use a finite number of independent simulations. In addition, although the stochasticity in our predicted satellite populations is reduced because we restrict our analysis to host halos with two Magellanic Cloud analogs, it is necessary to impose a small lower bound on λ to avoid realizations with zero

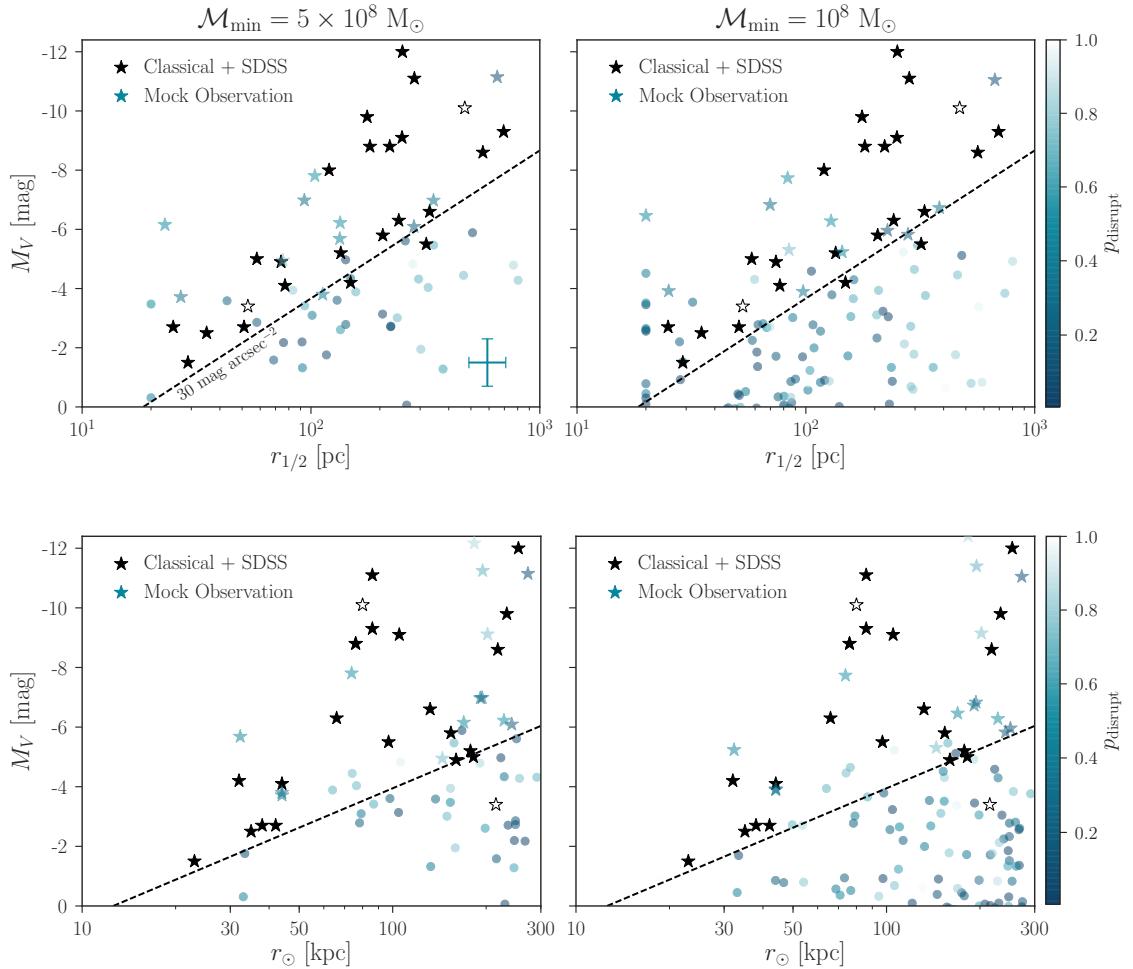


Figure 3.4: Visualizations of the predicted satellite population in one of our MW-like host halos (blue symbols) and the observed population of classical and SDSS-identified MW satellites (black stars) in the absolute magnitude–half-light radius plane (top panels) and the absolute magnitude–heliocentric distance plane (bottom panels) for $\mathcal{M}_{\min} = 5 \times 10^8 M_{\odot}$ (left panels) and $\mathcal{M}_{\min} = 10^8 M_{\odot}$ (right panels). All classical satellite analogs ($M_V \leq -8.8$ mag) and all systems in a mock SDSS footprint that pass both the SDSS surface brightness and completeness radius detection limits (dashed lines; [271, 273]) are plotted as blue stars, while blue circles show systems in the mock SDSS footprint that do not pass both detection criteria. The color bar indicates satellite disruption probability. This realization uses a faint-end slope of $\alpha = -1.3$, a luminosity scatter of $\sigma_M = 0.2$ dex, and a disruption parameter of $\mathcal{B} = 1$; we fix the remaining parameters according to Table 3.1. The error bars in the top-left panel show characteristic uncertainties for M_V and $r_{1/2}$. Pisces I and Pegasus III, which were discovered in SDSS using methods that do not adhere to our assumed detection criteria, are plotted as unfilled stars.

likelihood if we adopt Equations 3.11–3.12. This choice of lower bound is necessarily arbitrary.

We therefore approach the problem in a different manner. Rather than calculating a single estimate of the rate parameter in each bin, $\lambda_i(\boldsymbol{\theta})$, from the mock observed satellites, we marginalize over an unknown rate parameter in each bin. In particular, if we observe n_i real satellites and $\hat{n}_{i,j}$ mock satellites in bin i , where $j = 1, \dots, \hat{N}$ runs over all simulations and model realizations, we have

$$\begin{aligned} P(n_i | \hat{n}_{i,1}, \dots, \hat{n}_{i,\hat{N}}) &= \int P(n_i | \lambda_i) P(\lambda_i | \hat{n}_{i,1}, \dots, \hat{n}_{i,\hat{N}}) d\lambda_i \\ &= \frac{1}{P(\hat{n}_{i,1}, \dots, \hat{n}_{i,\hat{N}})} \int P(n_i | \lambda_i) P(\hat{n}_{i,1} | \lambda_i) \dots P(\hat{n}_{i,\hat{N}} | \lambda_i) P(\lambda_i) d\lambda_i \\ &= \left(\frac{\hat{N}+1}{\hat{N}} \right)^{-(\hat{n}_{i,1} + \dots + \hat{n}_{i,\hat{N}} + n_i)} \times (\hat{N}+1)^{-n_i} \frac{(\hat{n}_{i,1} + \dots + \hat{n}_{i,\hat{N}} + n_i)!}{n_i! (\hat{n}_{i,1} + \dots + \hat{n}_{i,\hat{N}})!}, \end{aligned} \quad (3.13)$$

where we have left the dependence on the model parameters $\boldsymbol{\theta}$ implicit, and we have assumed (i) a flat prior on λ_i for $\lambda_i \geq 0$ and (ii) that n_i and all $\hat{n}_{i,j}$ are drawn from the same Poisson distribution with rate parameter λ_i .

Because we produce non-integer numbers of mock satellites by counting each system as $p_{\text{detect}} \times (1 - p_{\text{disrupt}})$ object, we replace the factorials in Equation 3.13 with the appropriate Gamma functions to obtain the final form of the likelihood. Our results are unaffected if we enforce integer counts by performing multiple mock observations of each predicted satellite population.

We note that Equation 3.13 is similar to the likelihood derived in [241]; however, these authors marginalize over λ_i , given a single predicted satellite population from a particular simulation and model realization, and then average the resulting probabilities, while our likelihood treats all simulations and model realizations simultaneously. In Appendix B.1, we demonstrate that our likelihood converges to the underlying Poisson distribution in the limit of many mock observations, while the likelihood used in [241] does not.

Finally, we use Bayes' theorem to compute the resulting posterior distribution over our free parameters:

$$P(\boldsymbol{\theta} | \mathbf{s}_1, \dots, \mathbf{s}_N) = \frac{P(\mathbf{s}_1, \dots, \mathbf{s}_N | \boldsymbol{\theta}) P(\boldsymbol{\theta})}{P(\mathbf{s}_1, \dots, \mathbf{s}_N)}, \quad (3.14)$$

where $P(\boldsymbol{\theta})$ is our prior distribution, $P(\mathbf{s}_1, \dots, \mathbf{s}_N)$ is the Bayesian evidence, and

$$P(\mathbf{s}_1, \dots, \mathbf{s}_N | \boldsymbol{\theta}) = \prod_{\text{bins } i} P(n_i | \hat{n}_{i,1}, \dots, \hat{n}_{i,\hat{N}}) \quad (3.15)$$

is the likelihood.

For our fit to classical and SDSS-identified satellites, we use the six MW-like host halos described previously. For each host, we generate five mock satellite populations by simultaneously drawing (i) satellite luminosities and sizes from our luminosity and size relations, (ii) random observer locations

Parameter	Prior	Prior Set By
α	$\arctan \alpha \sim \text{unif}(-1.1, -0.9)$	Uninformative prior for $-2 \lesssim \alpha \lesssim -1.25$.
σ_M	$\sigma_M \sim \text{unif}(0.2 \text{ dex}, 1.0 \text{ dex})$	$\sigma_M \approx 0.2 \text{ dex}$ for higher-mass halos.
\mathcal{M}_{\min}	$\log(\mathcal{M}_{\min}/\text{M}_\odot) \sim \text{unif}(7.5, 10)$	Conservative upper bound based on [241].
\mathcal{B}	$\mathcal{B} \sim \text{Lognormal}(\mu = 1, \sigma = 0.5)$	$\mathcal{B} = 1$ corresponds to hydrodynamic results.

Table 3.2: Prior distributions for the parameters varied in our fit to the luminosity distribution of classical and SDSS-identified MW satellites.

8 kpc from the host halo center from the vertices of an octahedron, and (iii) random survey orientations. We assume that mock classical satellites ($M_V \leq -8.8$ mag) over the entire sky are always detected, and that mock satellites in the SDSS footprint that pass both of the detection criteria described in Section 3.5.1 are always detected, and we fit our model to the luminosity distribution of the classical and SDSS-identified systems listed in Section 3.5.1. We choose uniform absolute magnitude bins, and we assume that $P(\theta)$ factorizes into a product of independent prior distributions; we list our choices for these priors in Table 3.2. We have checked that our results are not significantly affected by our choice of magnitude bins and priors.

Due to the limited constraining power of the classical-plus-SDSS luminosity distribution, we only vary the following parameters: α , σ_M , \mathcal{M}_{\min} , and \mathcal{B} . We fix $\chi = 1$, $\mathcal{A} = 0.02$ and $\gamma = -0.7$ (i.e., the fiducial [244] size relation), $\beta = 1$ (i.e., our fiducial satellite size reduction model), $\sigma_R = 0.01$ dex, and $\mathcal{O} = 1$ (i.e., our fiducial orphan model) for the fit presented here. Our choice of σ_R is motivated by the fact that larger values of this scatter produce an overabundance of small observed mock satellites, since systems that scatter to small sizes at fixed distance and absolute magnitude are more likely to be observed. Thus, implementing an appreciable amount of size scatter likely requires modifying our assumptions that satellite size follows a lognormal distribution at fixed subhalo properties and that the scatter in our size relation is size-independent. In Appendix B.2, we test whether our radial scaling, size reduction, and orphan models are preferred by the classical-plus-SDSS luminosity distribution by computing Bayes factors for fits with $\chi = 1$, $\beta = 1$, and $\mathcal{O} = 1$ versus fits with $\chi = 0.8$, $\beta = 0$, and $\mathcal{O} = 0$. We find very weak evidence in favor of $\chi = 0.8$, $\beta = 1$, and $\mathcal{O} = 1$, implying that our fit to classical-plus-SDSS satellites is not sensitive to these effects and justifying our choice to fix these parameters.

To sample the posterior distribution for our fit to the classical-plus-SDSS luminosity distribution, we run 5×10^4 iterations of the Markov Chain Monte Carlo (MCMC) sampler `emcee` [180] using 20 walkers, and we discard the first 5000 burn-in steps. We have verified that our results are stable to changes in the number of walkers and that we have sampled a reasonable number of autocorrelation lengths for each chain.

3.6 Results and Discussion

The posterior distribution from our MCMC run is shown in Figure 3.5; we now summarize our main results.

3.6.1 Derived Constraints

1. Our fit favors a faint-end luminosity function slope of $\alpha = -1.34^{+0.04}_{-0.03}$ (68% confidence interval), which is shallower than most constraints from [440] and fairly consistent with the global constraint from GAMA ($\alpha = -1.26 \pm 0.07$; [295]).⁸ The GAMA constraint is derived from the luminosity function of galaxies with $-24 \text{ mag} \lesssim M_r \lesssim -13 \text{ mag}$, while we constrain the power-law slope of the luminosity function for systems dimmer than $M_r = -13 \text{ mag}$ by fitting to satellites with $-18 \text{ mag} \lesssim M_V \lesssim -1 \text{ mag}$.
2. Our fit is consistent with the luminosity scatter inferred for higher-mass galaxies ($\sigma_M \approx 0.2 \text{ dex}$), but it allows for significantly larger values, with $\sigma_M = 0.21^{+0.36}_{-0.00} \text{ dex}$ (68% confidence interval).
3. Our fit strongly favors a galaxy formation threshold of $\mathcal{M}_{\min} < 5 \times 10^8 \text{ M}_\odot$, with $\log(\mathcal{M}_{\min}/\text{M}_\odot) = 7.54^{+0.60}_{-0.04}$ (68% confidence interval); our results are consistent with the upper bound of $2.4 \times 10^8 \text{ M}_\odot$ from [241]. Decreasing \mathcal{M}_{\min} below 10^8 M_\odot rarely results in additional mock classical or SDSS observations, so the marginal likelihood for \mathcal{M}_{\min} is roughly flat in this regime.
4. Our fit is consistent with $\mathcal{B} = 1$: we find $\mathcal{B} = 1.04^{+0.56}_{-0.33}$ (68% confidence interval), which implies that our fiducial baryonic disruption model is compatible with the observed classical-plus-SDSS luminosity distribution.

To test whether our model provides an adequate fit to the data, we draw samples from the posterior distribution and plot the resulting 68% and 95% confidence intervals for the luminosity distribution, radial distribution, and size distribution of classical and SDSS satellites in Figures 3.6 and 3.7. Our predictions are largely consistent with both the observed luminosity function and the observed radial and size distributions of these systems, despite the fact that we have only fit to their luminosities. As noted previously, our model slightly underpredicts the observed population of satellites close to the center of the MW ($r_\odot \lesssim 50 \text{ kpc}$). However, since we have only fit to an observed luminosity distribution using fixed radial scaling and orphan prescriptions, this discrepancy might not persist for a joint fit to observed satellite luminosities, radii, and sizes that varies χ and \mathcal{O} . The dashed red lines in Figure 3.7 illustrate that decreasing χ reduces the tension among the predicted and observed inner radial distributions.

⁸Note that [295] constrain the Schechter function faint-end slope while we measure the power-law luminosity function slope itself.

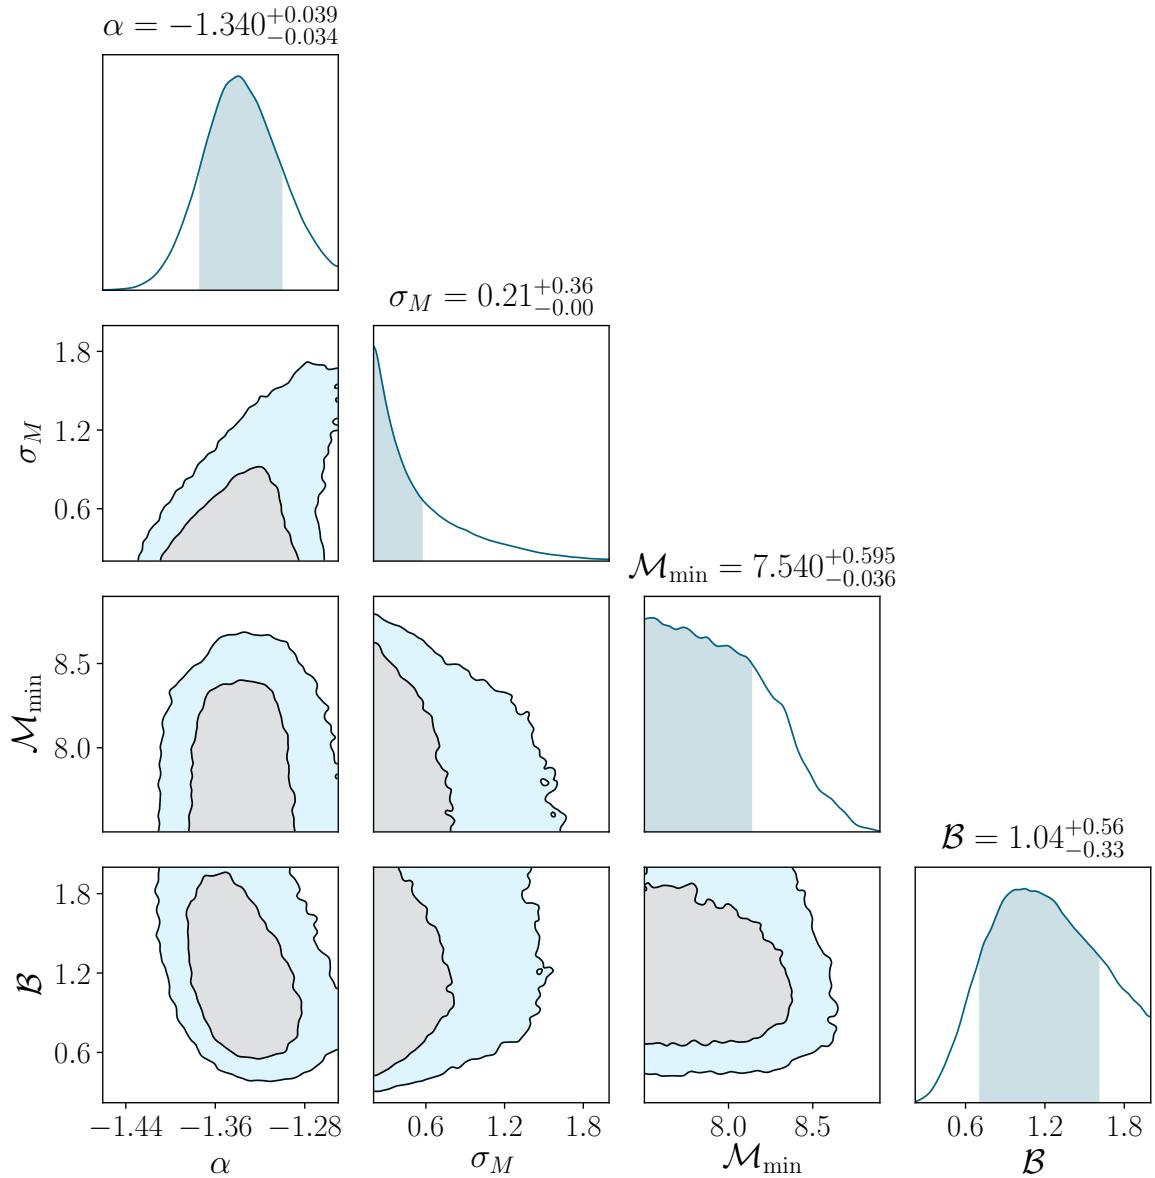


Figure 3.5: Posterior distribution from our fit to the luminosity distribution of classical and SDSS-identified satellites. Dark (light) shaded contours show 68% (95%) confidence intervals, and shaded areas in the marginal distributions show 68% confidence intervals. Note that σ_M is reported in dex and \mathcal{M}_{\min} is reported as $\log(\mathcal{M}_{\min}/M_{\odot})$ in this plot.

3.6.2 Predictions for Future Surveys

Given the sky coverage and detection efficiency of future MW satellite searches, we can use our model to predict the abundance and properties of the satellites that we expect to be discovered. To place our results in context, we first study the total number of satellites within 300 kpc of the MW — independent of surface brightness — inferred from our fit to the classical-plus-SDSS luminosity distribution. The left-hand panel of Figure 3.8 compares our prediction for the total number of MW satellites as a function of absolute magnitude to the results in [440], [271], and [340], and to the predictions derived in [241] based on three different stellar mass–halo mass relations. We predict 95 ± 29 (134 ± 44) total satellites with $M_V < -1.5$ mag ($M_V < 0$ mag) within 300 kpc of the MW at the 68% confidence level. Our estimate for the total number of MW satellites is consistent with but more conservative than most previous results, likely due to the fact that our subhalo disruption model captures both the effects of a central disk and additional baryonic physics. We refer the reader to [340] for a discussion of the discrepancy between the [440] prediction and other estimates. Finally, we note that the high- \mathcal{M}_{\min} tail of our posterior results in a small number of realizations with fewer total satellites than currently observed (including both spectroscopically confirmed systems and candidate satellites); we choose not to incorporate this constraint in our fit to restrict our analysis to classical and SDSS data alone.

Next, we make predictions for satellite searches with improved surface brightness limits by calculating the total number of MW satellites as a function of limiting observable surface brightness. In particular, the right-hand panel of Figure 3.8 shows the total number of MW satellites within 300 kpc inferred from our fit to the classical-plus-SDSS luminosity distribution, assuming that satellites over the entire sky are observed to a limiting surface brightness μ_{\lim} . We predict that 83 ± 26 (92 ± 29) satellites with $M_V < -1.5$ mag would be observed if the entire sky were covered to a limiting surface brightness of 32 (34) mag arcsec $^{-2}$. Similarly, we predict that $\sim 95\%$ of all MW satellites with $M_V < -1.5$ mag would be observed if satellites down to 33 mag arcsec $^{-2}$ were observed over the entire sky. We emphasize that these estimates depend on the details of our size model and surface brightness calculations.

To connect these predictions to ongoing and future surveys, we indicate approximate surface brightness detection thresholds for SDSS, DES, and LSST in Figure 3.8. We estimate the sensitivity of LSST satellite searches by comparing the 5σ limiting point-source magnitudes of a recent HSC-SSP satellite search [220] to those expected for LSST. LSST will likely achieve comparable sensitivity to this HSC-SSP survey in its first year of operation [237]; Equation 3.10 implies that the two satellite candidates recently detected in HSC-SSP data have $\mu_V = 31.6$ mag arcsec $^{-2}$ (Virgo I) and $\mu_V = 30.9$ mag arcsec $^{-2}$ (Cetus III), so we adopt an approximate surface brightness threshold of 32 mag arcsec $^{-2}$ for LSST Y1 (K. Bechtol 2019, private communication). We note that satellite searches with surveys such as HSC-SSP and LSST face the unique challenge of distinguishing faint stars associated with dwarf satellites in the Galactic halo from distant unresolved galaxies (e.g.,

[484]), adding to the uncertainty in our limiting surface brightness estimate.

3.6.3 Implications for the Connection between Low-mass Subhalos and Faint Satellites

The properties of the subhalos that host faint MW satellites can be used to constrain DM models that produce a cutoff in the subhalo mass function (e.g., [309, 260, 241]), along with the impact of reionization on galaxy formation (e.g., [327, 202]). Although we have not explicitly imposed a relationship between the luminosity of our mock satellites and the present-day mass of their subhalos, our model can be used to predict the joint distribution of subhalo mass and satellite luminosity, which we illustrate in Figure 3.9 by populating our MW-like host halos using the best-fit model derived previously with fixed $\mathcal{M}_{\min} = 10^8 M_{\odot}$. The tight correspondence between peak circular velocity and absolute magnitude enforced by our $M_V - V_{\text{peak}}$ relation is broadened by the mass–concentration relation (which relates V_{peak} to M_{peak}) and further broadened by tidal stripping (which relates M_{peak} to M_{vir}).

By drawing from our full posterior distribution (i.e., allowing \mathcal{M}_{\min} to vary) and repeatedly populating our six MW-like host halos, we find that satellites with $-6 \text{ mag} \lesssim M_V \lesssim 0 \text{ mag}$ typically reside in subhalos with present-day virial masses $2 \times 10^7 M_{\odot} \lesssim M_{\text{vir}} \lesssim 2 \times 10^8 M_{\odot}$ and peak virial masses $10^8 M_{\odot} \lesssim M_{\text{peak}} \lesssim 5 \times 10^8 M_{\odot}$ (68% confidence intervals). We caution that baryonic effects can systematically reduce subhalo masses [349, 329, 409], so these results should be interpreted as subhalo masses *in DMO simulations*; an additional layer of modeling calibrated on hydrodynamic simulations would be necessary to modify these values to account for the presence of baryons. The quoted lower bounds depend mildly on the low-mass cutoff in our prior for \mathcal{M}_{\min} , since mock satellites in low-mass subhalos can scatter into the absolute magnitude range of interest; however, we caution that our adopted low-mass cutoff is near the resolution limit of our simulations. Our upper bounds are consistent with the results in [241], who showed that mass thresholds in this range can be used to place competitive constraints on warm dark matter models.

We find that the faint satellites considered previously typically inhabit subhalos with $14 \text{ km s}^{-1} \lesssim V_{\text{peak}} \lesssim 22 \text{ km s}^{-1}$ (68% confidence interval). This lower bound on V_{peak} can potentially decrease for a joint fit to the luminosities and radial distances of observed satellites; for example, [202] showed that it might be necessary to populate subhalos down to $V_{\text{peak}} \approx 10 \text{ km s}^{-1}$ to match the observed radial distribution of MW satellites. However, these bounds are also dependent on our radial scaling and orphan satellite models, since decreasing χ or increasing \mathcal{O} raises the predicted abundance of mock satellites, particularly in the inner regions ($r \lesssim 50 \text{ kpc}$). For example, by drawing from our posterior and repeatedly populating our MW-like host halos with fixed $\mathcal{O} = 1$, we predict 9 ± 3 satellites with $r < 50 \text{ kpc}$ and $M_V < 0 \text{ mag}$ if satellite radii are set equal to subhalo radii ($\chi = 1$), and 14 ± 5 (37 ± 12) such satellites if $\chi = 0.8$ ($\chi = 0.5$). On the other hand, for fixed $\chi = 1$, we

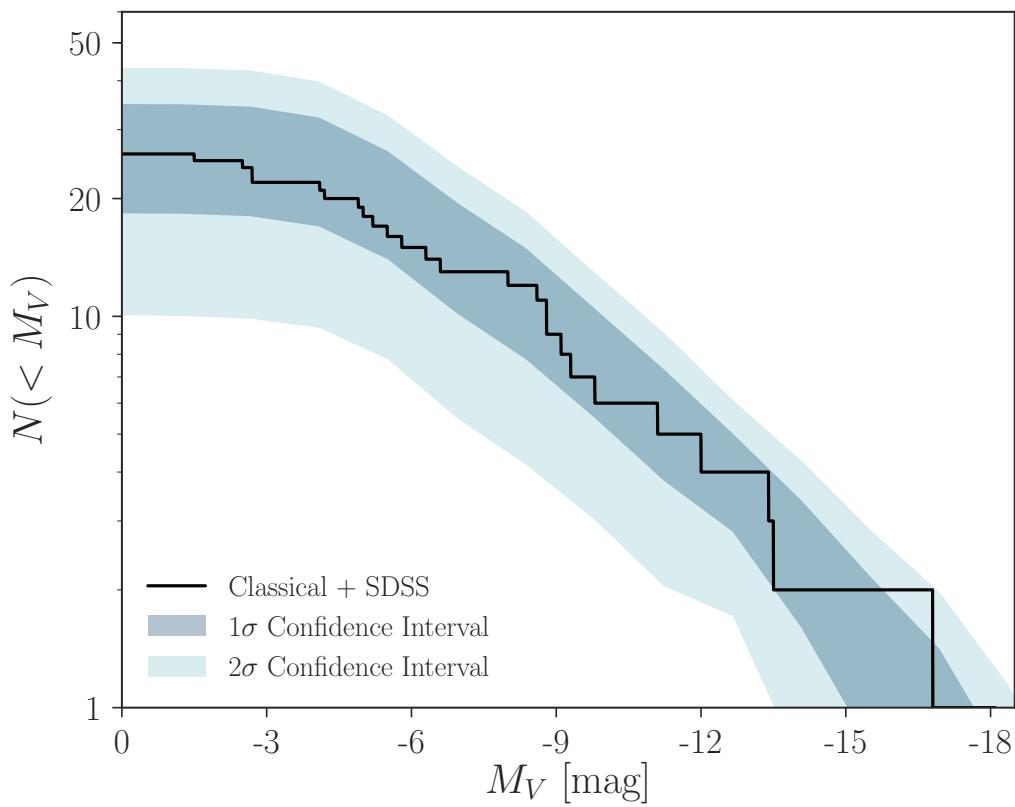


Figure 3.6: Our prediction for the observed luminosity function of classical and SDSS satellites inferred from our fit to the absolute magnitude distribution of these systems. Dark (light) shaded areas show 68% (95%) confidence intervals.

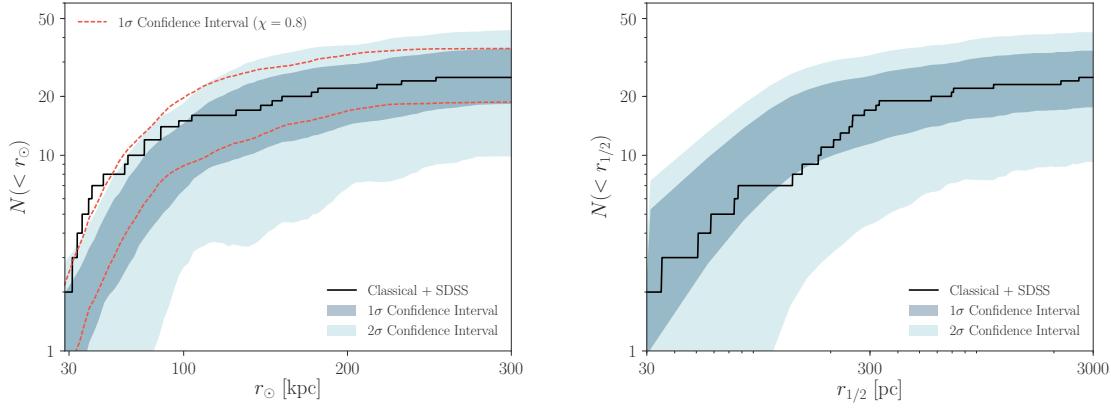


Figure 3.7: Left panel: the radial distribution of classical and SDSS satellites inferred from our fit to the observed luminosity distribution of these systems. Dashed red lines show 68% confidence intervals for a fit with satellite radii scaled inward relative to subhalo radii by a factor of 0.8. Right panel: the corresponding size distribution, calculated by setting satellite sizes based on subhalo properties at accretion ($\beta = 0$) with a constant lognormal scatter of $\sigma_R = 0.01$ dex. In both panels, dark (light) shaded areas show 68% (95%) confidence intervals.

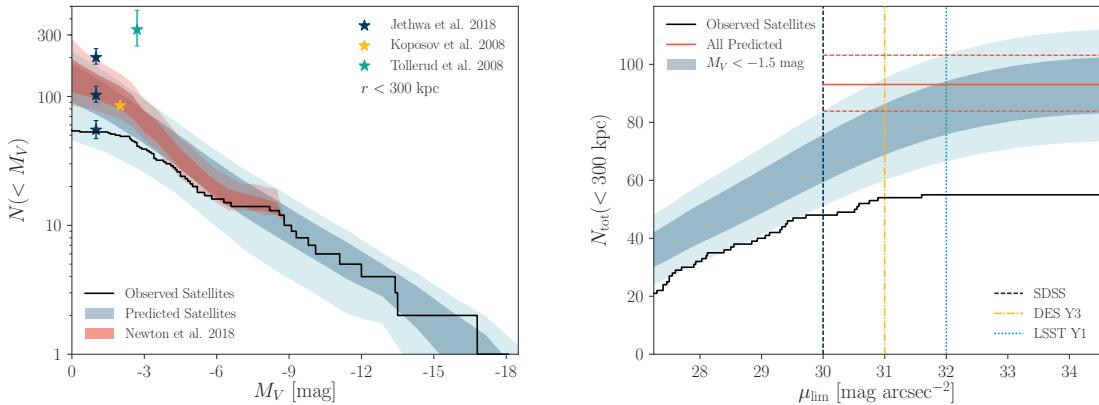


Figure 3.8: Left panel: the total number of satellites within 300 kpc of the MW inferred from our fit to the classical-plus-SDSS luminosity distribution (blue) compared to results from previous studies and to all observed MW satellites including candidate systems (black). Right panel: the total number of satellites with $M_V < -1.5$ mag within 300 kpc of the MW as a function of limiting observable surface brightness. The red line shows our prediction for the total number of satellites with $M_V < -1.5$ mag independent of surface brightness; dashed red lines show 68% Poisson confidence intervals. In both panels, dark (light) shaded areas indicate 68% (95%) confidence intervals.

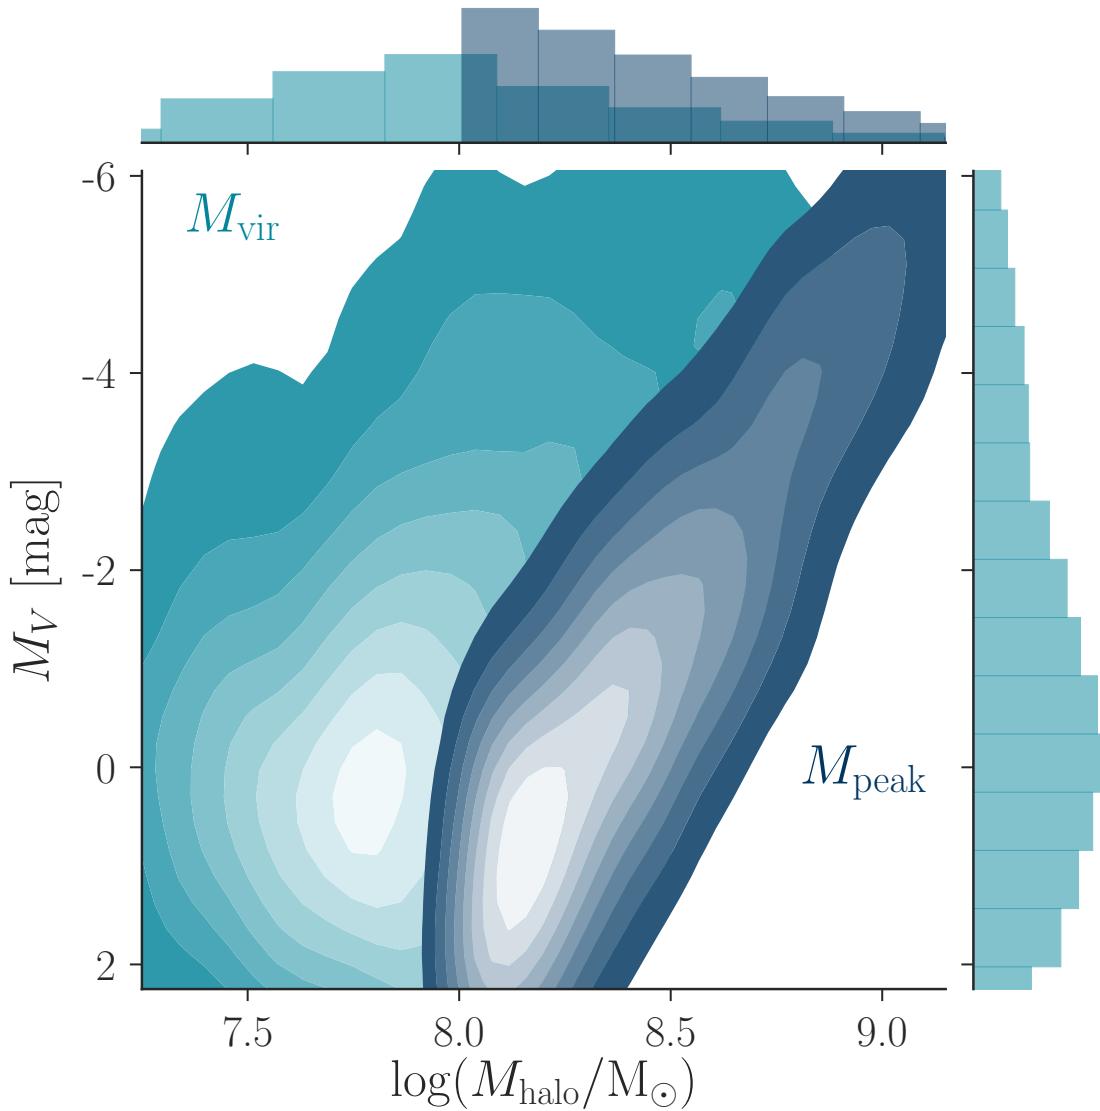


Figure 3.9: Joint distribution of subhalo mass and satellite absolute magnitude predicted by our best-fit model with $\mathcal{M}_{\min} = 10^8 \text{ M}_\odot$. The dark (light) contour corresponds to peak (present-day) subhalo virial mass, and the contours indicate satellite number density.

predict 5 ± 2 satellites with $r < 50$ kpc and $M_V < 0$ mag if no orphans are included ($\mathcal{O} = 0$), and 13 ± 4 such systems if all orphans are included with zero disruption probability ($\mathcal{O} \gg 1$; 68% confidence intervals). Our radial scaling and orphan parameters are therefore somewhat degenerate, but the fact that 18 satellites with $r < 50$ kpc have been observed (including candidate satellites) implies that it is difficult to reconcile models with fixed $\chi = 1$ and $\mathcal{O} = 0$ with the observed radial distribution.

Thus, in the context of our model, it is unclear whether the observed radial distribution requires a lower galaxy formation threshold than expected in standard reionization quenching scenarios. Future work that addresses a larger population of faint MW satellites will shed light on this issue; for example, including the ultra-faint dwarfs discovered by DES [42, 150] in a joint fit to satellite luminosities and radii can potentially decrease our upper bound on \mathcal{M}_{\min} and break the degeneracies necessary to constrain our radial scaling and orphan models. Such a study will require a careful treatment of LMC satellites, which is beyond the scope of this paper.

3.6.4 Caveats and Future Work

There are several caveats and possible extensions to the model presented in this paper. Most notably, the fact that we underpredict the number of satellites in the inner regions can potentially bias our parameter constraints and predictions for future surveys. However, we find that refitting the classical-plus-SDSS luminosity distribution with $\chi = 0.8$ to alleviate this tension (see Figure 3.7) does not significantly affect our results.

Although we have focused on modeling the observed absolute magnitudes, radial distances, and physical sizes of MW satellites, stellar velocity dispersion measurements provide an additional constraint on the subhalos that host these systems. The V_{\max} distribution of the subhalos that host our mock classical satellites extends to significantly higher values than those inferred from stellar velocity dispersion measurements using the [487] V_{\max} estimator; thus, our model suffers from the canonical “too big to fail” problem [84, 85], despite the fact that we include subhalo disruption due to baryonic effects. This shortcoming can likely be mitigated by incorporating the effects of stellar feedback on the inner density profiles of subhalos, along with a more sophisticated conversion between V_{\max} values and observed stellar velocity dispersions ([502, 311, 92, 103, 457]; see [242] and [97] for comprehensive discussions of the “too big to fail” problem).

By construction, our approach relies on zoom-in simulations; however, given an analytic method for generating subhalo V_{peak} functions, radial distributions, and size distributions, along with parameterizations of tidal stripping, subhalo disruption due to baryonic effects, and the contribution of orphan satellites, mock satellite populations could potentially be generated without using a particular set of simulations (see [213] for work along these lines). Such an approach would effectively combine our empirical framework with semi-analytic models to increase the level of detail at which we model satellite properties.

Exploring how our results depend on host halo mass would require either a different set of simulations or an analytic model, and characterizing this dependence might be particularly important in light of recent results that favor a relatively massive MW halo ([324, 370, 424, 471]; however, see [102]). In addition, although we have focused on satellite luminosities, radii, and sizes, comparing the orbital properties of observed and simulated systems will likely be fruitful in the era of precision astrometric measurements.

3.7 Summary

We have described a flexible model for populating subhalos in DMO zoom-in simulations of MW-mass hosts with satellite galaxies. We demonstrated that this model produces reasonable satellite populations in the regime of classical and SDSS-identified MW satellites, and we presented an improved method for fitting the model to observed satellite populations. Our fit to the classical-plus-SDSS luminosity distribution produces satellite populations that are qualitatively and quantitatively consistent with the luminosity function, radial distribution, and size distribution of observed systems, modulo modest tension in the inner radial distribution.

We briefly summarize the key aspects of our approach and highlight several open questions:

- (i) *Host halo properties*: We fix host halo properties based on our zoom-in simulation suite. Our host halos lie in the mass range $10^{12.1 \pm 0.03} M_{\odot}$ and have a variety of accretion histories and secondary properties. For our fit to observed MW satellites, we select host halos with two Magellanic Cloud-like systems. How do our predicted satellite populations vary as a function of host halo mass?
- (ii) *Satellite luminosities*: We assign satellite luminosities to DMO subhalos by abundance matching to the GAMA survey down to $M_r = -13$ mag and extrapolating this relation to fainter systems assuming a power-law luminosity function. We set the threshold for galaxy formation due to reionization using a cut on peak subhalo virial mass. Are constraints on the subhalo–satellite connection derived from MW satellites consistent with those from M31, the Local Volume, and SAGA hosts? What galaxy formation threshold is consistent with observations of ultra-faint satellites discovered since SDSS? What is the lowest subhalo mass that future observations of MW satellites can probe?
- (iii) *Satellite locations*: We model the locations of our mock satellites by scaling the radial distances of subhalos in our zoom-in simulations and projecting these systems onto the sky. Is the seemingly centrally concentrated radial distribution of MW satellites rare, or is it an artifact of misestimated observational incompleteness?
- (iv) *Satellite sizes*: We assign satellite sizes to DMO subhalos using the size relation from [244]. Is there evidence (e.g., from hydrodynamic simulations) that this size relation holds for ultra-faint dwarf galaxies?

- (v) *Baryonic effects*: We model subhalo disruption due to baryonic physics, such as the tidal influence of a galactic disk, using a model calibrated on hydrodynamic simulations. Were MW subhalos and satellites tidally disrupted in a manner that is consistent with hydrodynamic results?
- (vi) *Orphan satellites*: We include orphan satellites by tracking the orbits and modeling the tidal stripping of disrupted subhalos in our simulations. If subhalo disruption is a numerical artifact (e.g., [453, 454]), such that a significant population of disrupted subhalos should host observable satellite galaxies, can it be accounted for by modeling orphans? What is the relationship between subhalo disruption and satellite disruption?

Chapter 4

Galaxy–Halo Connection Constraints Including the LMC

Abstract

The population of Milky Way (MW) satellites contains the faintest known galaxies and thus provides essential insight into galaxy formation and dark matter microphysics. Here we combine a model of the galaxy–halo connection with newly derived observational selection functions based on searches for satellites in photometric surveys over nearly the entire high Galactic latitude sky. In particular, we use cosmological zoom-in simulations of MW-like halos that include realistic Large Magellanic Cloud (LMC) analogs to fit the position-dependent MW satellite luminosity function. We report decisive evidence for the statistical impact of the LMC on the MW satellite population due to an estimated 6 ± 2 observed LMC-associated satellites, consistent with the number of LMC satellites inferred from Gaia proper-motion measurements, confirming the predictions of cold dark matter models for the existence of satellites within satellite halos. Moreover, we infer that the LMC fell into the MW within the last 2 Gyr at high confidence. Based on our detailed full-sky modeling, we find that the faintest observed satellites inhabit halos with peak virial masses below $3.2 \times 10^8 M_\odot$ at 95% confidence, and we place the first robust constraints on the fraction of halos that host galaxies in this regime. We predict that the faintest detectable satellites occupy halos with peak virial masses above $10^6 M_\odot$, highlighting the potential for powerful galaxy formation and dark matter constraints from future dwarf galaxy searches.

4.1 Paper Status and External Contributions

This chapter is published in slightly modified form in the *Astrophysical Journal*, Volume 893, Issue 1, p.48-70 with the title, “Milky Way Satellite Census. II. Galaxy–Halo Connection Constraints

Including the Impact of the Large Magellanic Cloud,” on which I am the corresponding author. It is the result of a collaboration with the Dark Energy Survey (DES) Milky Way Working Group (MWWG) and particularly Risa Wechsler, Keith Bechtol, Yao-Yuan Mao, Gregory Green, and Alex Drlica-Wagner. Along with other members of the MWWG including Mitch McNanna and Sidney Mau, Alex and Keith led the observational effort to characterize the census of Milky Way satellite galaxies in DES and Pan-STARRS1 data along with their observational selection function, providing key input for this work. Risa, Keith, Yao, Greg, and Alex all provided helpful feedback on interfacing and interpreting the satellite model with this data, and they all made editorial contributions to the text. As with other DES papers, this work would not have been possible without significant collaboration, and I am particularly grateful to MWWG conveners and members (both past and present) for facilitating these efforts.

4.2 Introduction

The sample of confirmed and candidate Milky Way (MW) satellite galaxies has more than doubled in the last 5 years. Modern imaging surveys have driven these discoveries; in particular, following the successes of the Sloan Digital Sky Survey (SDSS) in the early 2000s [485, 486, 54, 53, 50, 52, 51, 207, 208, 399, 503, 234, 464], the Dark Energy Survey (DES) and the Panoramic Survey Telescope and Rapid Response System Pan-STARRS1 (PS1) have discovered 17 and three new satellite galaxy candidates, respectively [42, 150, 261, 272, 282, 283, 305]. These systems are identified as arcminute-scale overdensities of individually resolved stars, and many have already been spectroscopically confirmed. Meanwhile, other surveys with the Dark Energy Camera and VST ATLAS have recently discovered several additional satellites [318, 148, 445, 446, 443, 274, 320].

Nonetheless, the current census of MW satellites is likely highly incomplete, particularly for faint systems in the outer regions of the MW halo. This is evidenced by the detection of three new ultrafaint satellites in the first $\sim 676 \text{ deg}^2$ of Hyper Suprime-Cam Strategic Survey Program (HSC-SSP) imaging data [222, 220, 221] and by the discovery of Antlia II, the lowest surface brightness galaxy currently known, using RR Lyrae member stars identified in Gaia DR2 [444]. In the near future, the Legacy Survey of Space and Time (LSST) conducted from the Vera C. Rubin Observatory will be able to detect satellites over the entire southern sky down to a surface brightness of $\mu_V \sim 32 \text{ mag arcsec}^{-2}$ [237, 440, 214, 336].

Interpreting the cosmological and astrophysical implications of these discoveries requires a detailed understanding of the observational selection effects for each survey under consideration. In a companion paper ([149], hereafter Paper I), we derive observational selection functions for DES and PS1 based on searches for simulated satellites in each dataset. These selection functions encode the probability that satellites in either survey are detectable as a function of their absolute magnitude, heliocentric distance, physical size, and position on the sky. They incorporate realistic

photometric error models, selection masks that exclude highly reddened regions near the Galactic disk, and the influence of local stellar density on satellite detectability. Detection sensitivity is linked to sky position because various surveys have imaged different parts of the sky at varying depths, and accurately modeling this effect is crucial in order to disentangle anisotropy in the underlying MW satellite system from selection effects.

In this paper, we combine the observational selection functions derived in Paper I with a detailed model of the galaxy–halo connection and high-resolution cosmological zoom-in simulations of MW-mass host halos to infer the position-dependent MW satellite luminosity function. Although several empirical models have recently been used to study subsets of the MW satellite population [241, 264, 340, 336], this is the first analysis that is directly based on imaging data over more than $\sim 15,000 \text{ deg}^2$; indeed, our analysis covers 75% of the high Galactic latitude sky. Moreover, our galaxy–halo connection model allows us to marginalize over astrophysical uncertainties in our fit to the observed DES and PS1 satellite populations. We quantify the impact of the largest MW satellite, the Large Magellanic Cloud (LMC), and its associated satellites on the observed DES and PS1 satellite populations. We find that the satellites accreted with a realistic LMC analog—defined in terms of its mass, heliocentric distance, and infall time—are essential to fit the DES and PS1 luminosity functions simultaneously; this finding constitutes a remarkable confirmation of hierarchical structure formation. We predict that 4.8 ± 1.7 (1.1 ± 0.9) of the known satellites observed by DES (PS1) fell into the MW with the LMC, consistent with the number of LMC-associated satellites inferred from Gaia proper-motion measurements [249, 356].

Our analysis constrains the properties of the lowest-mass halos that host observed satellites, which we infer to have peak virial masses below $3.2 \times 10^8 M_\odot$ at 95% confidence. This finding, along with constraints on the faint-end slope of the luminosity function, can be used to inform feedback prescriptions in hydrodynamic simulations [410, 175, 427, 330, 482]. Constraints on the minimum halo mass also hold broad implications for the microphysical properties of dark matter (e.g., [151, 335]). Crucially, our model can be extended to explore the degeneracies between baryonic physics and deviations from the cold dark matter (CDM) paradigm.

This paper is organized as follows. We first provide an overview of our framework in Section 4.3. We then describe the simulations (Section 4.4), galaxy–halo connection model (Section 4.5), observational selection functions (Section 4.6), and statistical framework (Section 4.7) used in our analysis. We present our results in Section 4.8, focusing on the observed DES and PS1 satellite populations (§4.8.1), the impact of the LMC system (§4.8.2), the total MW satellite population (§4.8.3), galaxy–halo connection model constraints (§4.8.4), the properties of halos that host the faintest observed satellites (§4.8.5), and the implications of our findings for dark matter micro-physics (§4.8.6). We discuss the main theoretical uncertainties in our analysis in Section 4.9, and we conclude in Section 4.10. Appendices provide additional details on our galaxy–halo connection model (Appendix C.1) and statistical framework (Appendix C.2), the robustness of our results to

observational systematics (Appendix C.3) and resolution effects (Appendix C.4), and the observed DES and PS1 satellite populations (Appendix C.5).

Throughout, we use the term “galaxy–halo connection model” to refer to a model that describes how the properties of galaxies, including luminosity and size, are related to the properties of halos, such as peak virial mass. Furthermore, “log” refers to the base-10 logarithm.

4.3 Analysis Overview

Using the observed population of MW satellites to constrain our galaxy–halo connection model requires the following ingredients (see Figure 4.1 for a visualization of each step).

1. A model that predicts the *underlying* MW satellite population.
2. An observational selection function that, convolved with the prediction in the previous step, yields a prediction for the *observed* satellite population.
3. A model for the likelihood of producing the true MW satellite population given the prediction from the previous step.

The first step above can either be performed using either a hydrodynamic simulation, in which galaxy formation is modeled at the simulation level, or an empirical prescription for painting galaxies onto halos. We take the latter approach in this paper, which allows for a more flexible modeling framework, as well as the use of simulations that are closer approximations to the MW system. Indeed, our results may help to constrain feedback prescriptions in hydrodynamic simulations. Note that the assumed dark matter model (e.g., cold vs. warm dark matter or collisionless vs. interacting dark matter) affects the underlying satellite population and often manifests as a cutoff in the abundance of halos—and thus faint galaxies—below a halo mass threshold determined by the microphysical properties of dark matter.

The steps outlined above each rely on tools that have been developed in previous studies and in Paper I. Here we simply provide a brief description of each step, and we defer additional details to the Appendices.

4.4 Simulations

4.4.1 General Description

Our model of the underlying MW satellite population is built on high-resolution dark matter–only zoom-in simulations of MW-mass host halos selected from the suite of 45 hosts presented in [315],

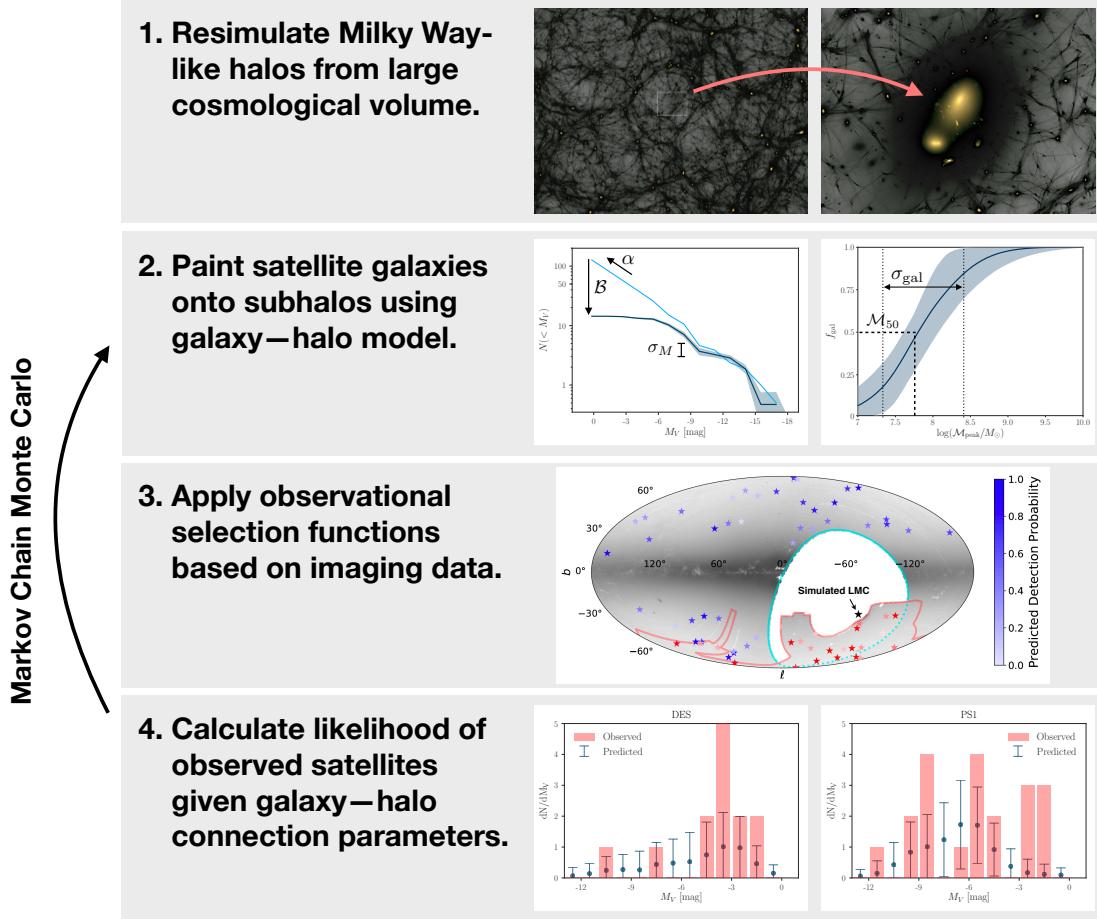


Figure 4.1: Visualization of our MW satellite modeling framework. In the first step, we perform high-resolution zoom-in simulations of MW-like halos selected from a larger cosmological volume (Section 4.4); in the second step, we paint galaxies onto subhalos using a parametric model for the galaxy–halo connection (Section 4.5); in the third step, we use the observational selection functions derived in Paper I to compute the probability that these satellites would be observed in DES or PS1 imaging data (Section 4.6); and in the final step, we calculate the likelihood of producing the true DES and PS1 satellite populations given many mock satellite population realizations at fixed galaxy–halo connection model parameters (Section 4.7). We then iterate this process to constrain our galaxy–halo connection model.

which have virial masses between 1.2 and $1.6 \times 10^{12} M_{\odot}$.¹ The highest-resolution particles in these simulations have a mass of $3 \times 10^5 M_{\odot} h^{-1}$, and the softening length in the highest-resolution regions is $170 \text{ pc } h^{-1}$.

Halo catalogs and merger trees were generated using the ROCKSTAR halo finder and the CONSISTENT-TREES merger code [47, 48]. Subhalos in these simulations are well resolved down to a present-day maximum circular velocity of $V_{\max} \approx 9 \text{ km s}^{-1}$ [315]. To be conservative, we only use subhalos with both $V_{\max} > 9 \text{ km s}^{-1}$ and peak maximum circular velocity $V_{\text{peak}} > 10 \text{ km s}^{-1}$. In Appendix C.4, we show that these resolution thresholds are sufficient for modeling the satellite populations of interest here.

4.4.2 Host Halo and LMC Analog Selection

The MW might be atypical compared to the average halos of a similar mass (e.g., [86, 100, 393, 173]); in particular, its satellite population is likely affected by the existence of the LMC system and the “satellites of satellites” that accreted with the LMC into the virial radius of the MW [306, 143, 302, 145]. In addition, the detailed merger history of the MW—such as the early accretion of an LMC-mass galaxy inferred from Gaia data—might affect its faint satellite population [78].

Thus, we select MW-like host halos that each have an LMC analog with realistic internal and orbital properties; both of these hosts experience an early major merger similar to the Gaia-Enceladus accretion event (see Appendix C.1.3 for details). We define realistic LMC analogs as subhalos with

1. present-day maximum circular velocity $V_{\max} \geq 55 \text{ km s}^{-1}$,
2. present-day heliocentric distance $40 \text{ kpc} < D < 60 \text{ kpc}$, and
3. time of accretion onto the MW less than 2 Gyr ago.

These criteria yield two MW-like host halos with virial masses of 1.57 and $1.26 \times 10^{12} M_{\odot}$, respectively. Both of these hosts were used in the less restrictive host halo set defined in Chapter 3 [336], and both have a Navarro–Frenk–White (NFW) concentration parameter that is consistent with constraints set using the combination of satellite and globular cluster dynamics measured by Gaia [102, 471]. The LMC analogs in these two simulations have present-day virial masses of 1.6 and $2.5 \times 10^{11} M_{\odot}$ respectively, and both have peak virial masses of $3.0 \times 10^{11} M_{\odot}$. These LMC analogs accreted onto their host halos 1 and 1.5 Gyr ago, respectively, and their orbital dynamics are consistent with LMC proper-motion measurements (e.g., [250]).

Our fiducial LMC analogs have masses that are consistent with LMC mass estimates based on stellar stream dynamics, satellite dynamics, and the orbital histories of both Magellanic Clouds

¹We define virial quantities according to the Bryan–Norman virial overdensity [93], with $\Delta_{\text{vir}} \simeq 99.2$ as appropriate for the cosmological parameters adopted in our zoom-in simulations: $h = 0.7$, $\Omega_m = 0.286$, $\Omega_b = 0.047$, and $\Omega_{\Lambda} = 0.714$ [315].

[60, 360, 164, 163].² However, different studies have adopted various definitions of “LMC mass,” and precision in the LMC mass definition (and particularly in the distinction between peak and present-day halo mass) is crucial going forward. We expect that our inference is most sensitive to the peak mass rather than the present-day mass of the LMC because peak mass correlates more directly with the expected abundance of LMC satellites, particularly for recent infall scenarios. Other probes of LMC mass are likely sensitive to these quantities in different ways, and some—including timing arguments (e.g., [360]) and orbit-rewinding to infer LMC satellite abundances (e.g., [356])—might be most sensitive to the ratio of the LMC and MW halo masses.

Although the masses of our host halos are consistent with observational constraints for the MW (e.g., [100, 63, 355]), our simulations span a narrower range of host mass relative to the uncertainty on this quantity inferred from Gaia measurements. For example, [102] found that the MW host virial mass lies between 1.0 and $1.8 \times 10^{12} M_\odot$ at the 95% confidence level (also see [106, 292, 291]). Since subhalo abundance is proportional to host halo mass, predicted satellite abundances scale linearly with MW mass, modulo second-order changes in subhalo disruption due to variations in the mass accretion history of the central galaxy [259, 404]. Ideally, our analysis would be performed using MW-like host halos—all of which include realistic LMC analogs—that bracket the current range of allowed MW host mass; however, the availability of such MW-like systems is limited by the statistics of our simulations. Thus, we do not marginalize over the full range of allowed MW host masses in this work. We estimate the potential impact of this uncertainty in Appendix C.1.1.

4.5 Galaxy–Halo Connection Model

To associate satellite galaxies with subhalos in the simulations described above, we use a modified version of the model developed in Chapter 3 [336]. This model parameterizes the relationship between satellite and subhalo properties and the effects of baryonic physics on subhalo populations in flexible ways, which allows us to marginalize over the relevant theoretical uncertainties. Additional model details and tests are presented in Appendix C.1.

4.5.1 Satellite Luminosities

To associate satellite luminosities with subhalos, we follow [336] by employing an abundance-matching procedure that monotonically relates the absolute V -band magnitude of satellites, M_V , to the peak circular velocity of subhalos, V_{peak} .³ This relation is constrained by the GAMA survey [295, 193] for bright systems ($M_V < -13$ mag) and is extrapolated into the regime of dim satellites

²Although detailed exploration of Magellanic Cloud binary systems is beyond the scope of this work, we note that [419] and [107] found that LMC analogs with SMC-like companions are typically more massive than isolated LMC analogs.

³We perform abundance matching using V_{peak} to incorporate the effects of halo assembly bias and to mitigate the impact of subhalo tidal stripping [379, 288].

by treating the faint-end slope of the satellite luminosity function, α , and the lognormal scatter in luminosity at fixed V_{peak} , σ_M as free parameters. We assume that this scatter is lognormal and constant as a function of halo properties in our fiducial analysis; we explore a mass-dependent scatter model in Appendix C.1.2.

Our abundance-matching model is a simple, empirical prescription for assigning satellite luminosities that is not designed to capture the complexities of star formation in ultrafaint dwarf galaxies. For example, [79] argued that star formation in systems dimmer than $M_V \sim -5$ mag is effectively shut down by reionization, resulting in two distinct galaxy populations today. While our abundance-matching model is consistent with the current data, which are fit fairly well by a single power-law luminosity function (see Paper I), it will be valuable to investigate more detailed models of stellar mass growth and to compare against a wider range of observables, including the inferred star formation histories of MW satellites, in future work.

4.5.2 Satellite Sizes

We assign physical sizes to satellites by extrapolating a modified version of the size–virial radius relation from [275], which links a galaxy’s stellar 3D half-mass radius to its halo’s virial radius, into the faint satellite regime. In particular, we set the mean predicted size of each satellite at accretion according to

$$r_{1/2} \equiv \mathcal{A} \left(\frac{R_{\text{vir}}}{R_0} \right)^n, \quad (4.1)$$

where \mathcal{A} and n are free parameters, R_{vir} denotes the subhalo virial radius measured at accretion, and $R_0 = 10$ kpc is a normalization constant. Following [336], we equate the 3D half-mass radii predicted by Equation 4.1 to azimuthally averaged projected half-light radii; this conversion neglects mass-to-light weighting and projection effects. Nonetheless, this size relation yields reasonable mean sizes when compared to the observed population of classical and SDSS-discovered satellites [336].

We draw the size of each satellite from a lognormal distribution with a mean given by Equation 4.1 and a standard deviation of $\sigma_{\log R}$, which is a free parameter in our model. When fitting the observed satellite populations, we only compare predicted and mock satellites with $r_{1/2} > 10$ pc in order to exclude likely star clusters from the analysis. We explore a more conservative cut of $r_{1/2} > 20$ pc in Appendix C.3.2.

The size prescription described above assumes that satellite sizes are fixed after accretion onto the MW. However, post-infall effects such as tidal stripping and tidal heating can shrink or enlarge satellites depending on their orbital histories [359, 165, 169]. In Appendix C.1.4, we show that our key results are not sensitive to these effects using a simple model for satellite size evolution due to tidal stripping.

4.5.3 Subhalo Disruption Due to Baryonic Effects

To incorporate the effects of baryonic physics—and particularly the tidal influence of the Galactic disk—on our simulated subhalo populations, we apply a random forest algorithm trained on hydrodynamic simulations to predict the probability that each subhalo will be disrupted in a hydrodynamic resimulation based on its orbital and internal properties [190, 337]. We model the strength of this disruption effect using the free parameter \mathcal{B} , which is defined such that $\mathcal{B} = 1$ corresponds to fiducial hydrodynamic predictions [337] and larger (smaller) values of \mathcal{B} correspond to more effective (less effective) subhalo disruption. For each subhalo, we set

$$p_{\text{disrupt}} \equiv (p_{\text{disrupt},0})^{1/\mathcal{B}}, \quad (4.2)$$

where $p_{\text{disrupt},0}$ is the fiducial disruption probability returned by the machine-learning algorithm in [337].

4.5.4 Galaxy Formation Efficiency

The stochastic, nonlinear nature of galaxy formation in low-mass halos likely leads to a smoothly varying fraction of occupied halos, rather than a sharp cutoff in the efficiency of galaxy formation [410, 175, 330, 482]. Thus, in our fiducial model, we parameterize the fraction of halos that host galaxies of any mass, referred to as the *galaxy occupation fraction*, following [202],

$$f_{\text{gal}}(\mathcal{M}_{\text{peak}}) \equiv \frac{1}{2} \left[1 + \text{erf} \left(\frac{\mathcal{M}_{\text{peak}} - \mathcal{M}_{50}}{\sqrt{2}\sigma_{\text{gal}}} \right) \right], \quad (4.3)$$

where $\mathcal{M}_{\text{peak}}$ is the largest virial mass a subhalo ever attains, which typically occurs before infall into the MW; \mathcal{M}_{50} is the peak halo mass at which 50% of halos host galaxies of any mass; and σ_{gal} is the width of the galaxy occupation fraction. In our fiducial model, \mathcal{M}_{50} and σ_{gal} are free parameters. Note that in the limit $\sigma_{\text{gal}} \rightarrow 0$, this reduces to a model in which all halos with $\mathcal{M}_{\text{peak}} > \mathcal{M}_{50}$ host a galaxy.

Although our analysis does not constrain σ_{gal} , Equation 4.3 is a simple, physically motivated form of the occupation fraction that will be interesting to explore in future work. Note that we parameterize the occupation fraction in terms of peak halo mass (rather than, e.g., V_{peak}) because $\mathcal{M}_{\text{peak}}$ is more easily interpretable and connects directly to constraints on alternative dark matter models (e.g., [335]).

4.5.5 Orphan Satellites

Because we model faint galaxies that can potentially inhabit subhalos near the resolution limit of our simulations, it is important to account for artificially disrupted subhalos that might host “orphan”

satellites ([211], and see [78] for a recent example of the importance of orphans in satellite modeling). To model orphans, we follow the prescription in Chapter 3 [336], which identifies disrupted subhalos in each simulation, interpolates their orbits to $z = 0$ using a softened gravitational force law and a dynamical friction model, and accounts for tidal stripping with a mass-loss model calibrated on high-resolution test simulations. We parameterize the effective abundance of orphans by setting their disruption probabilities equal to

$$p_{\text{disrupt}} \equiv (1 - a_{\text{acc}})^{\mathcal{O}}, \quad (4.4)$$

where a_{acc} is the final scale factor at which each subhalo enters the virial radius of the MW, and \mathcal{O} is a parameter that captures deviations from disruption probabilities in hydrodynamic simulations, which are well fit by $\mathcal{O} = 1$ [336]. Thus, larger (smaller) values of \mathcal{O} correspond to a greater (smaller) contribution from orphan satellites.

Following [336], we include orphan satellites by fixing $\mathcal{O} = 1$ in our fiducial model. Thus, we effectively assume that subhalo disruption in dark matter–only simulations is a numerical artifact [453, 454] but that subhalo disruption in hydrodynamic simulations is a physical effect. We show that our results are largely insensitive to the value of \mathcal{O} in Appendix C.1.6.

4.6 Observational Selection Functions

We employ the DES and PS1 survey selection functions derived in Paper I, which have been publicly distributed as machine-learning classifiers that predict satellite detection probability given absolute magnitude, M_V ; heliocentric distance, D ; azimuthally averaged projected half-light radius, $r_{1/2}$; and sky position.⁴ The predicted detection probabilities are derived from searches for simulated satellites in catalog-level DES and PS1 data, and they employ geometric cuts that restrict observable satellites to lie within the respective survey footprint and that mask regions where satellite detection is challenging due to interstellar extinction, bright nearby stars, and bright extragalactic objects.

We self-consistently apply these position-dependent detection criteria to our predicted satellite populations by matching the on-sky position of our LMC analogs to the true on-sky position of the LMC. In particular, we choose random observer locations 8 kpc from the halo center, and we perform appropriate rotations to our subhalo populations for each observer location to match the true LMC position. We apply the DES selection function for satellites within the overlap region of the two surveys, and we only count satellites within a fiducial heliocentric distance of 300 kpc.

⁴The DES Y3A2 and PS1 DR1 selection functions are publicly available at <https://github.com/des-science/mw-sats>.

4.7 Statistical Framework

To fit our galaxy–halo model to the DES and PS1 luminosity functions derived in Paper I, we generate predicted satellite populations given a set of galaxy–halo connection model parameters, $\boldsymbol{\theta}$, by performing mock DES-plus-PS1 surveys using the selection functions described above. For each host halo and each realization of our satellite model, we bin mock-observed satellites according to their absolute magnitude. We further split satellites in each absolute magnitude bin into high ($\mu_V < 28 \text{ mag arcsec}^{-2}$) and low ($\mu_V \geq 28 \text{ mag arcsec}^{-2}$) surface brightness samples to incorporate satellite size information in our fit.⁵ We list the DES and PS1 satellites used in this analysis in Table C.2.

Next, we calculate the number of predicted satellites in each bin i via

$$n_i = \sum_{s_i} p_{\text{detect},s_i} \times (1 - p_{\text{disrupt},s_i}) \times f_{\text{gal},s_i}, \quad (4.5)$$

where s_i indexes the satellites in bin i , p_{detect} is the detection probability returned by the appropriate observational selection function, p_{disrupt} is the disruption probability due to baryonic effects (Equation 4.2), and f_{gal} is the galaxy occupation fraction (Equation 4.3). For objects that lie in the overlap region of the DES and PS1 footprints, we calculate p_{detect} using the DES selection function.

We note that detection probability mainly depends on surface brightness and present-day heliocentric distance (see Paper I), disruption probability mainly depends on orbital properties [337], and galaxy occupation depends on $\mathcal{M}_{\text{peak}}$ according to Equation 4.3. Thus, our model for satellite detectability is coupled to our galaxy occupation fraction model, since surface brightness is directly linked to $\mathcal{M}_{\text{peak}}$ due to our abundance-matching assumption. Nonetheless, our results are largely unaffected if we exclude the galaxy occupation fraction from our model, confirming that f_{gal} can be interpreted as the probability that a halo hosts a satellite brighter than $M_V = 0 \text{ mag}$, corresponding to the faintest satellite in our observational sample.

We assume that satellites populate each bin in absolute magnitude–surface brightness parameter space according to an independent Poisson point process with a rate parameter λ that depends on absolute magnitude, surface brightness, and our galaxy–halo connection model parameters. Because our model yields noisy estimates of λ , we marginalize over its range of possible values in each bin, following [336]. The likelihood of observing the set of DES and PS1 satellites, \mathbf{s}_{DES} and \mathbf{s}_{PS1} (specified by their absolute magnitudes and surface brightnesses), given a set of model parameters $\boldsymbol{\theta}$ is then

$$P(\mathbf{s}_{\text{DES}}, \mathbf{s}_{\text{PS1}} | \boldsymbol{\theta}) = \prod_{\text{bins } i} P(n_{\text{DES},i} | \hat{\mathbf{n}}_{\text{DES},i}) \times P(n_{\text{PS1},i} | \hat{\mathbf{n}}_{\text{PS1},i}), \quad (4.6)$$

where $n_{\text{DES},i}$ ($n_{\text{PS1},i}$) is the observed number of DES (PS1) satellites in bin i , and $\hat{\mathbf{n}}_{\text{DES},i}$ ($\hat{\mathbf{n}}_{\text{PS1},i}$) is

⁵In particular, we calculate the effective surface brightness averaged within the half-light radius as $\mu_V = M_V + 36.57 + 2.5 \log(2\pi r_{1/2}^2)$, where $r_{1/2}$ is measured in units of kpc.

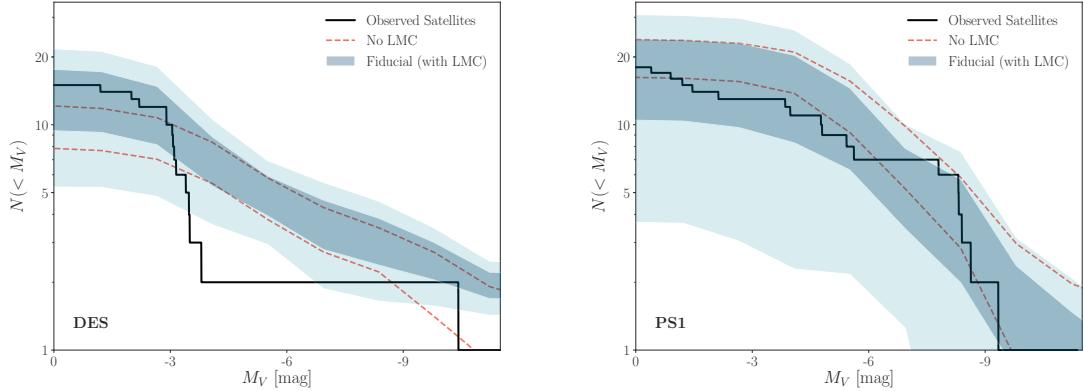


Figure 4.2: Predicted DES and PS1 satellite luminosity functions resulting from a joint fit to these satellite populations. Dark (light) blue bands correspond to 68% (95%) confidence intervals from our fiducial eight-parameter galaxy–halo connection model, dashed red lines show the 68% confidence interval for a model using host halos without LMC analogs (“No LMC”), and black lines show the observed luminosity functions within each survey footprint. Our fiducial model, which includes realistic LMC analogs, is decisively favored over the No LMC scenario, with a Bayes factor of $\sim 10^4$.

a vector of the number of mock DES (PS1) satellites in bin i from several realizations of our model at fixed $\boldsymbol{\theta}$. These realizations include draws over host halos, observer locations, and our galaxy–halo connection model, which is stochastic at fixed $\boldsymbol{\theta}$. Note that steps 1–3 in Figure 4.1 generate mock satellite populations $\hat{\mathbf{n}}_{\text{DES}}$ and $\hat{\mathbf{n}}_{\text{PS1}}$, and step 4 compares these to the observed populations n_{DES} and n_{PS1} . The explicit forms of $P(n_{\text{DES},i}|\hat{\mathbf{n}}_{\text{DES},i})$ and $P(n_{\text{PS1},i}|\hat{\mathbf{n}}_{\text{PS1},i})$ are given in Equation C.7.

Finally, we use Bayes’s theorem to compute the posterior distribution over galaxy–halo connection model parameters,

$$P(\boldsymbol{\theta}|\mathbf{s}_{\text{DES}}, \mathbf{s}_{\text{PS1}}) = \frac{P(\mathbf{s}_{\text{DES}}, \mathbf{s}_{\text{PS1}}|\boldsymbol{\theta})P(\boldsymbol{\theta})}{P(\mathbf{s}_{\text{DES}}, \mathbf{s}_{\text{PS1}})}, \quad (4.7)$$

where $P(\boldsymbol{\theta})$ is our prior on the galaxy–halo connection model parameters (given in Appendix C.2.2), $P(\mathbf{s}_{\text{DES}}, \mathbf{s}_{\text{PS1}})$ is the Bayesian evidence, and $P(\mathbf{s}_{\text{DES}}, \mathbf{s}_{\text{PS1}}|\boldsymbol{\theta})$ is given by Equation 4.6. To sample from this posterior, we run 10^5 iterations of the Markov Chain Monte Carlo (MCMC) sampler `emcee` [180] to sample the eight free parameters $\boldsymbol{\theta} = (\alpha, \sigma_M, \mathcal{M}_{50}, \mathcal{B}, \sigma_{\text{gal}}, \mathcal{A}, \sigma_{\log R}, n)$ using 32 walkers. We discard a burn-in period of 20 autocorrelation lengths, corresponding to $\sim 10^4$ steps, which leaves more than 100 independent samples.

4.8 Results

We now present our results, focusing on the observed DES and PS1 satellite populations (§4.8.1), the impact of the LMC system (§4.8.2), the total MW satellite population (§4.8.3), the galaxy–halo

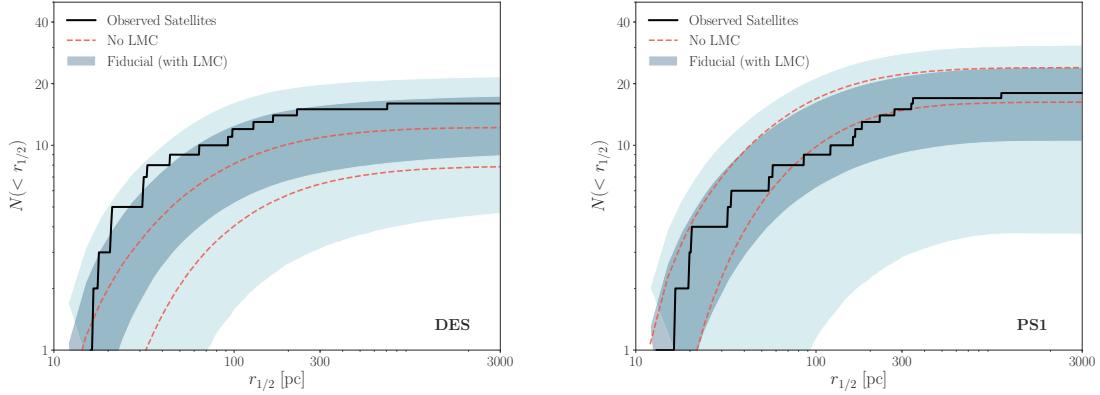


Figure 4.3: Size distributions derived by fitting our galaxy–halo connection model to the DES and PS1 satellite populations. Dark (light) blue bands correspond to 68% (95%) confidence intervals from our fiducial eight-parameter model, dashed red lines show the 68% confidence interval for a model using host halos without LMC analogs (“No LMC”), and black lines show the observed size distributions.

connection model constraints (§4.8.4), the properties of the halos that host faint satellites (§4.8.5), and the implications for dark matter microphysics (§4.8.6).

4.8.1 Observed Satellite Populations

Figure 4.2 shows the 68% and 95% confidence intervals for the observed DES and PS1 satellite luminosity functions given by draws from the posterior of our fiducial model, which is consistent with both datasets. We note that the DES and PS1 likelihoods individually yield consistent results.

Figure 4.3 shows the corresponding satellite size distributions drawn from our fiducial posterior. Our model is consistent with the sizes of both the observed DES and PS1 satellites. It very slightly overpredicts the sizes of observed DES systems; however, we reiterate that our size model does not allow for size reduction due to tidal stripping or size enlargement due to tidal heating, which affect satellites with close pericentric passages to the Galactic disk (e.g., [24]). Our findings in Appendix C.1.4 suggest that the post-infall size evolution of satellites in subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ and $V_{\text{max}} > 9 \text{ km s}^{-1}$ does not significantly affect our inference.

Our fiducial model is consistent with the outer radial distributions of both DES and PS1 satellites, but it slightly underpredicts the number of satellites near the center of the MW ($D \lesssim 100 \text{ kpc}$), particularly in PS1. We explore this minor discrepancy in Appendix C.1.3, where we show that our galaxy–halo connection model constraints and inferred total MW satellite population are largely unaffected if the radial distribution is forced to match the data.

4.8.2 The Impact of the LMC

To assess the impact of the LMC and its satellites on the MW satellite population, we test the following models in addition to our fiducial model, which includes a realistic, recently accreted LMC system by construction.

1. No LMC: a model with four host halos that have the same mean concentration as our fiducial hosts but no LMC analog.
2. Misplaced LMC: a model with our fiducial host halos in which subhalo positions are reflected, effectively placing the DES footprint in the northern hemisphere.
3. Early LMC Infall: a model with two host halos that have the same mean concentration as our fiducial hosts with LMC analogs that pass our LMC V_{\max} and heliocentric distance cuts but fall into the MW 2 and 6 Gyr ago, respectively.

For each alternative LMC scenario listed above, we refit the observed DES and PS1 satellite populations, sampling over the same eight parameters used in our fiducial analysis.

Our fiducial model is favored over the No LMC, Misplaced LMC, and Early LMC Infall scenarios with Bayes factors of $\sim 10^4$, 10^4 , and 10^3 , respectively. In addition, both host halos in the Early LMC Infall case are individually disfavored with Bayes factors of $\sim 10^3$. Thus, we find decisive statistical evidence for the impact of the LMC on the MW satellite population, particularly within and near the DES footprint. Moreover, we infer that the LMC system fell into the MW within the last 2 Gyr at high confidence. We also note that, in our fiducial host with more massive MW and LMC halos, the LMC reaches pericenter near the second-to-last simulation snapshot (i.e., ~ 150 Myr ago). Performing our analysis using the final snapshot for this host noticeably degrades the fit due to the dispersal and disruption of LMC satellites during the LMC’s pericentric passage. Thus, we use the second-to-last snapshot for this host in our fiducial analysis, and we remark that satellite abundances can potentially constrain the number of allowed pericentric passages for the LMC.

The alternative LMC scenarios defined above are strongly disfavored relative to our fiducial model because they cannot produce a sufficient number of dim satellites in the DES footprint without overpredicting the number of observed PS1 satellites. This is a direct consequence of the spatial overdensity of subhalos near the LMC analogs in our fiducial simulations; in particular, the projected density of resolved subhalos within 50° of the LMC on the sky is enhanced by $\sim 50\%$ relative to the density on a random patch of sky.

To quantify the number of satellites in our fiducial model that are associated with the LMC, we explore the following definitions of LMC-associated subhalos.

1. Fiducial definition. A subhalo is associated with the LMC if it is within the virial radius of the LMC halo at the time of LMC infall into the MW.

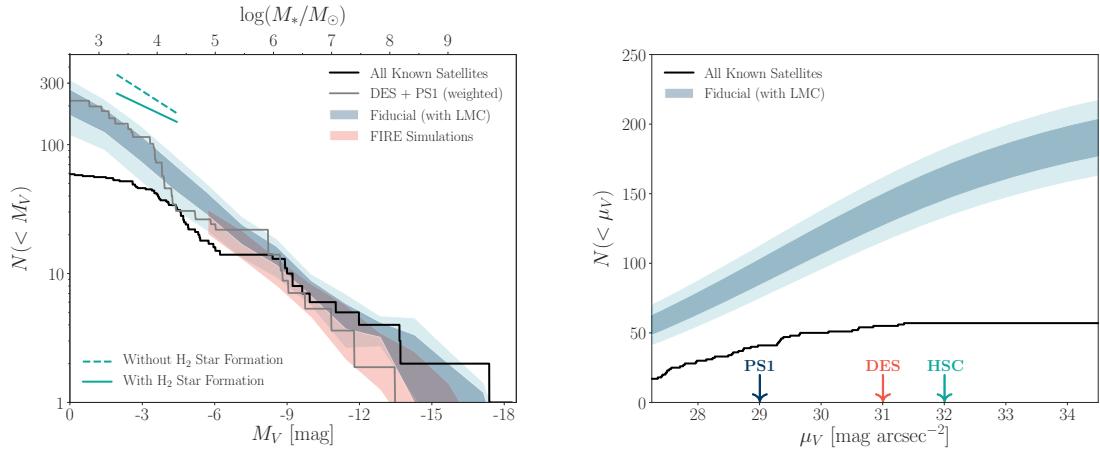


Figure 4.4: Left panel: total MW satellite luminosity function inferred from our joint fit to the DES and PS1 satellite populations (blue) compared to the current census of confirmed and candidate MW satellites (black) and the empirical estimate derived in Paper I (gray), which assumes an isotropic satellite distribution and a cored NFW radial satellite distribution. The 68% confidence intervals from hydrodynamic simulations of the Local Group using the FIRE feedback prescription are shown in red [188]. Luminosity function slopes predicted from hydrodynamic simulations with (solid green line) and without (dashed green line) H₂-based star formation are shown for comparison [330]; these predictions do not account for subhalo disruption due to the Galactic disk. Note that the Paper I prediction (gray) differs from the “All Known Satellites” curve (black) at the bright end because it does not include the LMC, SMC, or Sagittarius. Right panel: The surface brightness distribution of MW satellites with $M_V < 0$ mag and $r_{1/2} > 10$ pc as a function of the limiting observable surface brightness of an all-sky survey. Arrows indicate approximate detection limits for current surveys. Note that LSST Y1 is expected to have similar detection sensitivity to HSC [237, 440, 214, 336].

2. Gravitationally influenced definition. A subhalo is associated with the LMC if it has ever passed within the virial radius of the LMC halo.

Here LMC infall is defined as the snapshot at which the center of the LMC halo crosses within the MW virial radius. Note that nearly all systems that satisfy our strict definition are bound to the LMC at the time of LMC infall.

Under the fiducial (gravitationally influenced) definitions above, we predict that 52 ± 8 (181 ± 25) total LMC-associated subhalos (above our cuts of $V_{\text{peak}} > 10 \text{ km s}^{-1}$ and $V_{\text{max}} > 9 \text{ km s}^{-1}$) exist within the virial radius of the MW today, where the 95% confidence interval is estimated by drawing from our fiducial posterior. We predict that 48 ± 8 (164 ± 25) of these subhalos form galaxies with $M_V < 0$ mag and $r_{1/2} > 10$ pc (in agreement with an earlier estimate by [241]), and that 41 ± 7 (118 ± 21) of these satellites survive tidal disruption due to the Galactic disk. Of these surviving LMC-associated satellites, we predict that 4.8 ± 1.7 (11 ± 3.6) are currently observed by DES and that 1.1 ± 0.9 (6.1 ± 2.1) are currently observed by PS1.

Our statistical probe of LMC satellite association is remarkably consistent with the number of observed LMC satellites inferred from Gaia proper-motion measurements, which indicate that four galaxies in or near the DES footprint—excluding the Small Magellanic Cloud (SMC)—are associated with the LMC, and that two satellites in or near the PS1 footprint are potentially associated with the LMC ([249, 356]).⁶. In addition, the orbital dynamics of our predicted LMC satellites are consistent with Gaia proper-motion measurements for these likely LMC-associated systems. These predictions are also consistent with other empirical models [127, 240, 145, 400, 249, 164, 498] and with hydrodynamic simulations of isolated LMC analogs [238].

In Appendix C.4, we show that the properties of our LMC-like systems are not significantly affected by the realizations of small-scale power in our fiducial simulations. However, we caution that the number of predicted LMC satellites observed by DES and PS1 depends on the particular properties of our two LMC analogs. Thus, exploring the robustness of these results using a suite of zoom-in simulations selected to contain realistic LMC systems with a range of internal and orbital properties is an important avenue for future work.

4.8.3 The Total MW Satellite Population

Figure 4.4 shows the total MW satellite luminosity function and surface brightness distribution resulting from our fit to the DES and PS1 satellite populations. We infer that a total of 220 ± 50 satellites with $M_V < 0$ mag and $r_{1/2} > 10$ pc exist within the virial radius of the MW, where uncertainties correspond to 68% confidence intervals calculated by sampling from our fiducial posterior. Thus, we predict that ~ 150 satellites remain undiscovered in a standard CDM scenario, roughly one-fourth of which are associated with the LMC. This is larger than the fraction of satellites that

⁶A recent analysis based on Gaia proper-motion measurements and hydrodynamic simulations suggests that two bright satellites in or near DES, Fornax and Carina, may also be LMC-associated ([353]; however, see [356])

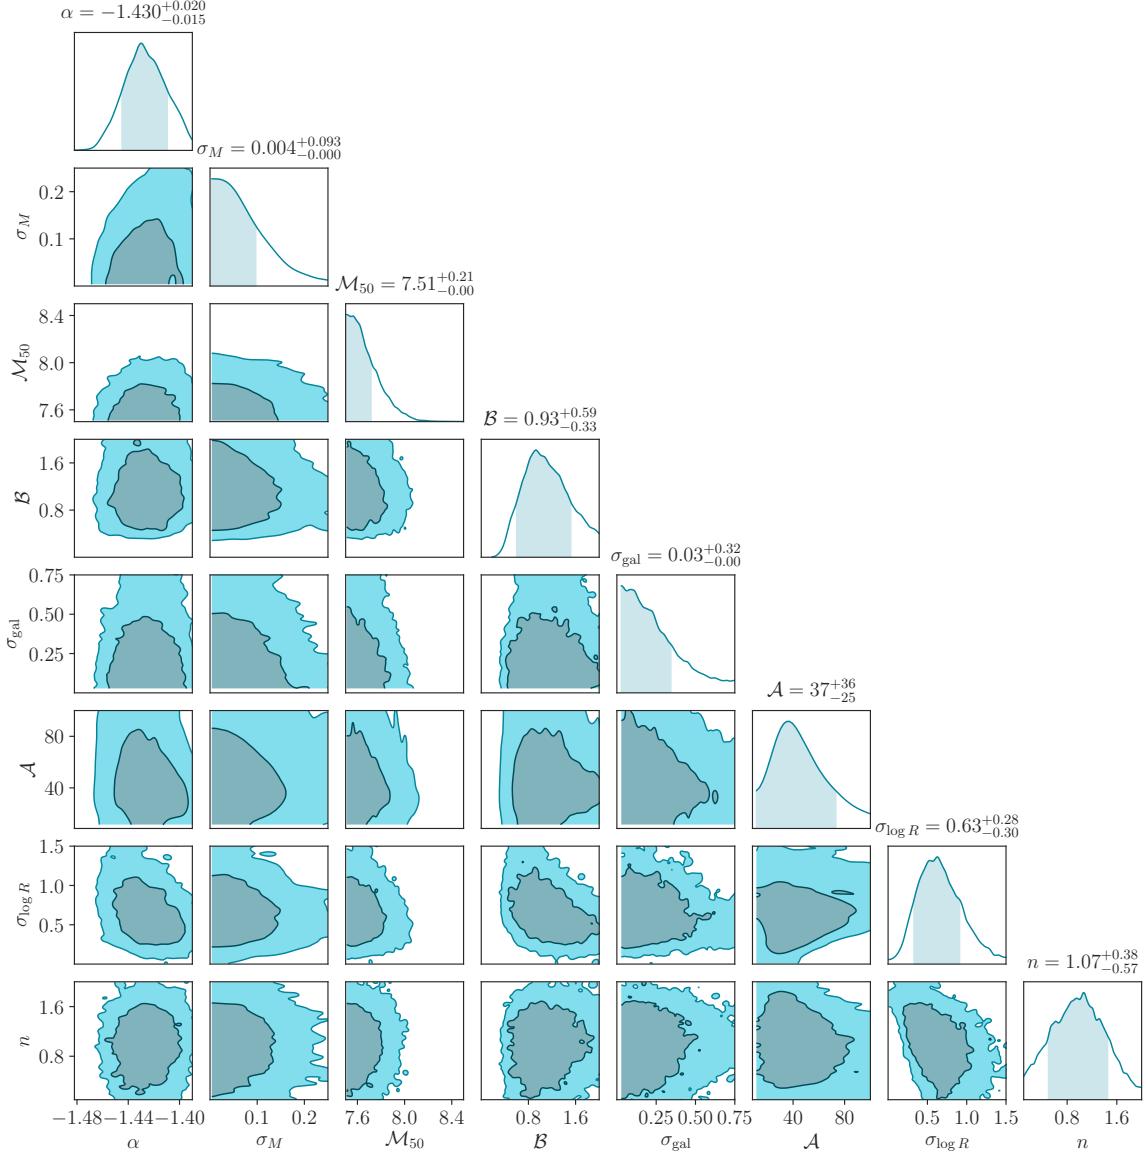


Figure 4.5: Posterior distribution from our fit to the DES and PS1 satellite populations. Dark (light) shaded contours represent 68% (95%) confidence intervals. Shaded areas in the marginal distributions and parameter summaries correspond to 68% confidence intervals. Note that σ_M , σ_{gal} , and $\sigma_{\log R}$ are reported in dex, \mathcal{M}_{50} is reported as $\log(\mathcal{M}_{50}/M_\odot)$, and \mathcal{A} is reported in pc.

Parameter	Physical Interpretation	Confidence Interval
α	Power-law slope of satellite luminosity function	$-1.46 < \alpha < -1.39$
σ_M	Scatter in luminosity at fixed V_{peak}	$0 \text{ dex}^* < \sigma_M < 0.19 \text{ dex}$
\mathcal{M}_{50}	Mass at which 50% of halos host galaxies	$7.5^* < \log(\mathcal{M}_{50}/M_{\odot}) < 7.93$
\mathcal{B}	Strength of subhalo disruption due to baryons	$0.3 < \mathcal{B} < 2.1$
σ_{gal}	Scatter in galaxy occupation fraction	$0 \text{ dex}^* < \sigma_{\text{gal}} < 0.67 \text{ dex}$
\mathcal{A}	Amplitude of galaxy–halo size relation	$0^* \text{ pc} < \mathcal{A} < 110 \text{ pc}$
$\sigma_{\log R}$	Scatter in half-light radius at fixed halo size	$0 \text{ dex}^* < \sigma_{\log R} < 1.2 \text{ dex}$
n	Power-law index of galaxy–halo size relation	$0^* < n < 1.8$

Table 4.1: Galaxy–halo connection model constraints derived from our fit to the DES and PS1 satellite populations. Asterisks mark prior-driven constraints.

have ever fallen into the MW that are associated with the LMC because our fiducial LMC analogs accreted recently, making their satellites less likely to be disrupted. Our prediction for the total number of MW satellites is consistent with several recent studies [241, 264, 340, 336], and it is lower than the empirical estimate in Paper I, which was recognized to be inflated due to the assumption of an isotropic satellite distribution. This prediction will be tested by upcoming deep imaging surveys; indeed, HSC-SSP has already started to probe this population of distant, low surface brightness MW satellites by discovering three new ultrafaint satellite candidates in $\sim 676 \text{ deg}^2$ of imaging data [222, 220, 221].

To estimate whether our predictions are consistent with HSC-SSP observations, we draw realizations of the MW satellite population from our fiducial posterior and calculate the number of systems within the DES or PS1 footprints that would not be observed by the appropriate survey. We then estimate the number of these systems currently observed by an HSC-like survey covering 676 deg^2 that detects all satellites (i.e., systems with $M_V < 0 \text{ mag}$ and $r_{1/2} > 10 \text{ pc}$) down to a surface brightness of $\mu_V = 32 \text{ mag arcsec}^{-2}$ and out to a heliocentric distance of 300 kpc, assuming an isotropic satellite distribution at high Galactic latitudes and accounting for subhalo disruption. There are six known satellites in the HSC footprint, but two of the six (Sextans and Leo IV) are detected at high significance in PS1 by at least one of the search algorithms in Paper I. We find that our mock HSC survey detects 1.75 ± 0.6 satellites, which is in slight tension with the number of systems detected by HSC (four, after discounting Sextans and Leo IV).

Figure 4.4 illustrates several predictions from hydrodynamic simulations of isolated and satellite dwarf galaxies. Our results are largely consistent with the luminosity function of bright MW satellites in hydrodynamic simulations of the Local Group using the Feedback In Realistic Environments (FIRE) feedback prescription, down to the FIRE resolution limit of $\sim -6 \text{ mag}$ [188]. Note that these FIRE simulations do not include LMC or SMC analogs, which accounts for the discrepancy with both our predictions and the observed luminosity function at $M_V < -16 \text{ mag}$. Interestingly, other recent hydrodynamic simulations indicate that different star formation prescriptions significantly

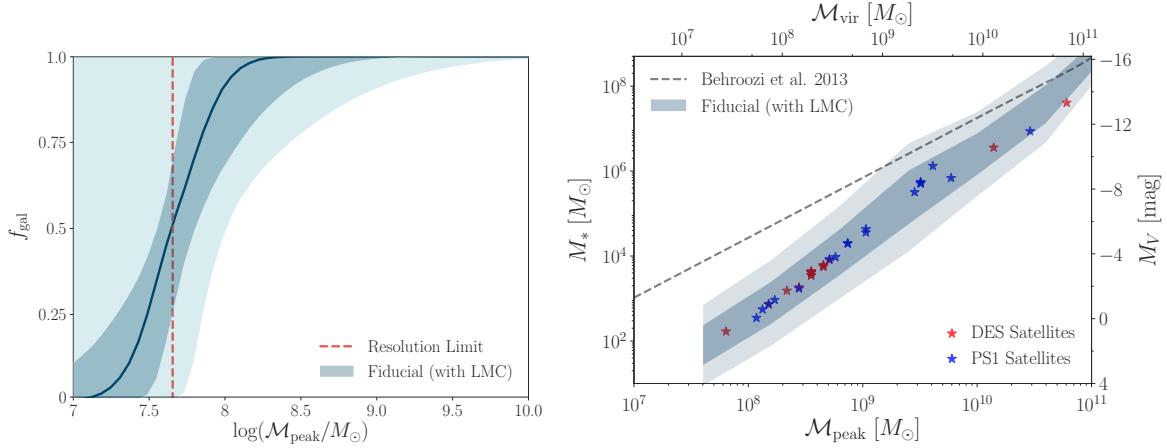


Figure 4.6: Left panel: fraction of halos that host galaxies, inferred from our fit to the DES and PS1 satellite populations. The solid line shows the median inferred galaxy occupation fraction, and dark (light) shaded contours represent 68% (95%) confidence intervals. The resolution limit of our simulations is indicated by the dashed vertical line. Right panel: SMHM relation inferred from our fit to the DES and PS1 satellite populations. An extrapolation of the mean SMHM relation derived from more luminous field galaxies is shown in gray [45]. Stars illustrate the mean of the predicted M_{peak} range corresponding each observed DES and PS1 satellite, and top ticks indicate the corresponding present-day virial masses of the halos that host these systems.

impact the amplitude and faint-end slope of the luminosity function for satellites of isolated LMC-like halos [330]. Thus, our constraints on the faint-end slope, which are driven by satellites with $M_V > \sim -6$ mag (corresponding to stellar mass $M_* < \sim 10^5 M_{\odot}$), can be used to inform subgrid star formation prescriptions.

4.8.4 Galaxy–Halo Connection Model Constraints

The posterior distribution for our fiducial model is shown in Figure 4.5, and the corresponding galaxy–halo connection model constraints are listed in Table 4.1. Note that we obtain statistically consistent results when fitting the DES and PS1 satellite populations with either of our two fiducial simulations individually in terms of both the Bayesian evidence and the galaxy–halo connection model constraints. In particular, no parameter constraints shift by more than 1σ relative to the fiducial values reported below when the fit is performed using either simulation individually. We now discuss each constraint in detail.

1. The inferred faint-end slope of the satellite luminosity function is steeper than that reported in our study based on classical and SDSS satellites in Chapter 3 [336]. Our constraint is consistent with the faint-end slope derived from higher-luminosity field galaxies in the GAMA survey [295, 488], even though it is based on a sample that extends nearly 10 mag fainter

than that used in GAMA. We note that α is the most sensitive parameter in our analysis to modeling assumptions and details of the observed satellite population.

2. The scatter in luminosity at fixed V_{peak} is constrained to $\sigma_M < 0.19$ dex at 95% confidence, which may inform hydrodynamic feedback prescriptions that predict a steep increase in luminosity scatter at low masses (e.g., see [474]). Our lack of a lower limit on σ_M is consistent with previous studies of faint galaxy samples (e.g., [288]). Meanwhile, large values of σ_M are not allowed because too many low- V_{peak} satellites upscatter to observable luminosities, resulting in overpredicted luminosity functions. To confirm that this upper limit is robust, we calculate Bayes factors by drawing samples from the posterior in bins of σ_M , finding that $\sigma_M = 0.15$ dex ($\sigma_M = 0.2$ dex) is disfavored relative to $\sigma_M = 0$ dex with a Bayes factor of 30 (100).⁷ These upper limits are comparable to the scatter typically inferred from abundance-matching analyses of brighter systems ($\sigma_M \sim 0.2$ dex) and smaller than that from hydrodynamic simulations of dwarf galaxies (e.g., [381]); however, we caution that our constraint might be impacted by the use of only two independent realizations of the MW satellite population. In addition, it is potentially misleading to compare global constraints on scatter to those derived from the MW alone. Both of these caveats are important to explore in future work.
3. The peak mass at which 50% of halos host galaxies is inferred to be less than $8.5 \times 10^7 M_{\odot}$ at 95% confidence. Note that this summary statistic depends on the lower limit of our prior on \mathcal{M}_{50} , since the \mathcal{M}_{50} posterior flattens near its lower limit, which is chosen based on the resolution of our simulations. Thus, we also calculate Bayes factors by drawing from the posterior in bins of \mathcal{M}_{50} to confirm that this summary statistic is robust. We find that $\mathcal{M}_{50} = 8.5 \times 10^7 M_{\odot}$ ($\mathcal{M}_{50} = 1.5 \times 10^8 M_{\odot}$) is disfavored relative to arbitrarily low values of \mathcal{M}_{50} with a Bayes factor of 50 (100). The current data are not able to place a lower limit on \mathcal{M}_{50} , which would correspond to the detection of a cutoff in galaxy formation.
4. Our posterior is consistent with $\mathcal{B} = 1$, corresponding to our fiducial baryonic disruption model. Although a large spread in disruption strength is allowed by the data, extremely efficient ($\mathcal{B} > 2.1$) and inefficient ($\mathcal{B} < 0.3$) subhalo disruption relative to hydrodynamic simulations is strongly disfavored. These constraints widen when our lognormal prior on \mathcal{B} is relaxed; however, zero subhalo disruption (corresponding to $\mathcal{B} = 0$) is robustly ruled out.
5. The scatter in the galaxy occupation fraction is consistent with zero, which makes sense given our lack of a lower limit on \mathcal{M}_{50} . Models with large scatter in the occupation fraction ($\sigma_{\text{gal}} < 0.67$ dex), corresponding to extremely stochastic galaxy formation, are disfavored relative to a step function occupation fraction at 95% confidence. Note that the slope of the σ_{gal} posterior is driven by the lower limit of our \mathcal{M}_{50} prior; as this limit decreases, the σ_{gal} posterior flattens.

⁷We provide details on our Bayes factor calculations in Appendix C.2.3.

6. The amplitude of the galaxy–halo size relation, defined as the typical size of a satellite in a halo with $R_{\text{vir}} = 10$ kpc at accretion, is constrained to lie between 12 and 73 pc at 68% confidence. For larger values of \mathcal{A} , satellites are too large to be detected with high probability, and the DES and PS1 luminosity functions are underpredicted; for smaller values of \mathcal{A} , many predicted satellites do not pass our $r_{1/2} > 10$ pc cut, and the luminosity functions are underpredicted.
7. The scatter in the galaxy–halo size relation is constrained to lie between 0.33 and 0.91 dex at 68% confidence. For larger values of $\sigma_{\log R}$, faint satellites upscatter to large sizes too frequently, which results in underpredicted luminosity functions. Our 68% confidence lower limit on $\sigma_{\log R}$ of 0.33 dex is consistent with the value estimated in [275]. Lower values of $\sigma_{\log R}$ lead to slightly too many predicted DES and PS1 satellites; however, our results are consistent with $\sigma_{\log R} = 0$ dex at 95% confidence.
8. The power-law index of the galaxy–halo size relation is constrained to lie between 0.5 and 1.45 at 68% confidence. For shallower power-law slopes, satellite sizes do not change appreciably as a function of halo size, which results in a worse joint fit to the observed absolute magnitude and surface brightness distribution. We note that the posterior widens when our Gaussian prior on n is relaxed.

4.8.5 Properties of Halos that Host the Faintest Satellites

We now explore the properties of the lowest-mass halos inferred to host MW satellites. The left panel of Figure 4.6 shows the galaxy occupation fraction derived from our statistical inference, where uncertainties are estimated by drawing from our fiducial posterior. By sampling from our fiducial posterior, we infer that halos with a peak virial mass below $2.5 \times 10^8 M_{\odot}$ and peak circular velocity below 19 km s^{-1} host at least one of the faintest observed satellites. To convert these into maximally conservative upper limits, we account for the uncertainty in MW host halo mass using the procedure described in Appendix C.1.1, which yields limits on the minimum halo mass and peak circular velocity of $\mathcal{M}_{\text{min}} < 3.2 \times 10^8 M_{\odot}$ and $V_{\text{peak,min}} < 21 \text{ km s}^{-1}$ at 95% confidence. Furthermore, we predict that the faintest observed satellite inhabits a halo with $\mathcal{M}_{\text{peak}} = 1.5 \times 10^8 M_{\odot}$, on average.⁸

These results improve upon the minimum halo mass constraint derived from classical and SDSS satellites [336] by a factor of 2, and they are consistent with the constraints reported in [241]. Moreover, these upper limits are not in significant tension with the expected atomic cooling limit of $V_{\text{peak}} \approx 20 \text{ km s}^{-1}$, contrary to recent studies based on the radial MW satellite distribution (e.g., [202]) and consistent with the findings in [78].

We caution that the median galaxy occupation fraction shown in Figure 4.6 is driven by the assumed functional form in Equation 4.3 and is therefore arbitrary. Although the functional form

⁸The faintest observed satellite in our analysis, Cetus II, is detected by DES with $M_V = 0.02$ mag ([150]; Table C.2).

in Equation 4.3 is consistent with results from hydrodynamic simulations for $\mathcal{M}_{\text{peak}} \gtrsim 10^9 M_{\odot}$, this particular functional form is not required to fit the DES and PS1 luminosity functions. Rather, we have evidence that $f_{\text{gal}} > 50\%$ above a peak virial mass of $\sim 10^8 M_{\odot}$. To verify that the assumed form of the galaxy occupation fraction does not impact our constraints, we also test a binned model in which we fit for \mathcal{M}_{50} and a corresponding 90% occupation mass. We find that the resulting 50% and 90% occupation constraints are consistent with those inferred from our fiducial analysis.

A wide range of galaxy occupation fractions have been reported in hydrodynamic simulations, with some placing \mathcal{M}_{50} as high as $\sim 10^9 M_{\odot}$ [410, 175]. However, recent hydrodynamic simulations run at higher resolution result in efficient galaxy formation in significantly lower-mass halos, and some claim that *all* halos down to the resolution limit consistently host star particles [482]. In addition, simulations of galaxy formation at early pre-reionization epochs show that stellar systems form in halos with masses as low as $\sim 10^7 M_{\odot}$ (e.g., see Fig. 13 in [116] for a compilation of recent simulation results). Most recently, high-resolution simulations of high-redshift galaxy formation that include the effects of spatially and temporally inhomogeneous reionization find $\mathcal{M}_{50} \sim 10^8 M_{\odot}$ [257].

Our galaxy occupation fraction constraint implies that models with $\mathcal{M}_{50} > 10^8 M_{\odot}$ are in significant tension with the observed MW satellite population, as long as MW satellite formation is representative of galaxy formation at this halo mass scale, on average. This assumption may not be true if the reionization history of the MW’s Lagrangian volume differs from the average reionization history of an MW-mass halo hosting dwarf galaxies of the masses considered here (however, see [22, 101]). Note that analyses based on H I surveys of Local Group dwarfs indicate a suppression mass scale similar to our \mathcal{M}_{50} constraint [441].

Due to our abundance-matching assumption, the lowest-mass halos in our model host the faintest galaxies, on average. Thus, our constraints on the masses of these halos are conservative, since the most massive halos in our simulations are forced to host more easily observable satellites at fixed distance and size, modulo baryonic disruption effects and abundance-matching scatter. In other words, our abundance-matching model yields a testable prediction: the faintest galaxies should inhabit the halos with the lowest pre-infall virial masses. We expect this correlation to be weakened by post-infall effects, including tidal stripping, but we can nevertheless infer the present-day joint distribution of halo mass and satellite luminosity or stellar mass. We illustrate this stellar mass–halo mass (SMHM) relation in the right panel of Figure 4.6. Our inferred SMHM relation is generally consistent with recent results (e.g., [241]). Like the faint-end slope of the luminosity function, the SMHM relation can be used to discriminate between different subgrid models of star formation and stellar feedback [330]. As in previous studies, we find that the SMHM relation in the ultrafaint regime falls off more steeply than extrapolations of abundance-matching relations derived using higher-mass field galaxies [45]. Interestingly, [14] found that full on-the-fly radiative transfer is necessary to match the steepness and normalization of our inferred SMHM relation for a fixed hydrodynamic feedback prescription.

Ultimately, our predictions must be confronted with the dynamical mass function of observed satellites, measurements of which will improve significantly in the era of upcoming spectroscopic facilities and giant segmented mirror telescopes [423]. A preliminary comparison of our joint predicted distribution of luminosity and V_{\max} with the measured stellar velocity dispersions of DES and PS1 satellites suggests that our model is consistent with the inferred central densities of low-luminosity satellites ($M_V > -6$ mag). Although there is a systematic discrepancy between observed and predicted values of V_{\max} for brighter systems (the “too big to fail” problem), our simple comparison does not account for the conversion from line-of-sight velocity dispersion measured within observed half-light radii to V_{\max} or the tidal effects of the Galactic disk on the density profiles of surviving subhalos. Moreover, the systems for which predicted densities are higher than those inferred observationally are susceptible to baryonic feedback processes that core the inner regions of halos [136], and this effect has been shown to alleviate the too big to fail problem [91, 408, 478, 300, 188].

Finally, we explore the properties of the halos inferred to host the faintest potentially detectable galaxies. In particular, we calculate the minimum peak halo mass necessary for halos to contain a stellar population of at least $100 M_{\odot}$, chosen to represent the approximate threshold for which it would be possible to observationally confirm a stellar system as a dark matter-dominated dwarf galaxy.⁹ By populating a higher-resolution version of one of our fiducial simulations and sampling from the posterior of our abundance-matching relation, we find that systems at the observational threshold occupy halos with $\mathcal{M}_{\text{peak}} > 10^6 M_{\odot}$ at 95% confidence. To detect even lower-mass halos, gravitational probes of dark matter that are independent of baryonic content, e.g. gravitational lensing or stellar streams, must be employed.

4.8.6 Implications for Dark Matter Microphysics

Many deviations from CDM lead to a cutoff in the abundance of low-mass halos. Several authors have used MW satellite abundances to constrain a free-streaming cutoff induced by warm dark matter (e.g., [309, 260, 299, 241]). [335] showed that similar constraints apply to other dark matter models, resulting in limits on the velocity-independent scattering cross section between dark matter and baryons. Our statistical detection of halos with peak virial masses below $3.2 \times 10^8 M_{\odot}$ therefore translates directly into constraints on various microphysical properties of dark matter.

In Chapter 6 [335], we find that the minimum halo mass inferred in this fashion is comparable to the limit on the half-mode mass M_{hm} , which corresponds to the scale at which the matter power spectrum is suppressed by a critical amount relative to CDM due to dark matter free streaming or interactions. Performing a statistical inference in which the half-mode mass is varied will constrain it to lie below our upper limit on \mathcal{M}_{50} , since the abundance of halos at and above the half-mode mass is reduced relative to CDM.

Thus, for a simple and conservative estimate of the dark matter constraints resulting from our

⁹For many MW satellites, this will likely require spectroscopy with giant segmented mirror telescopes [423].

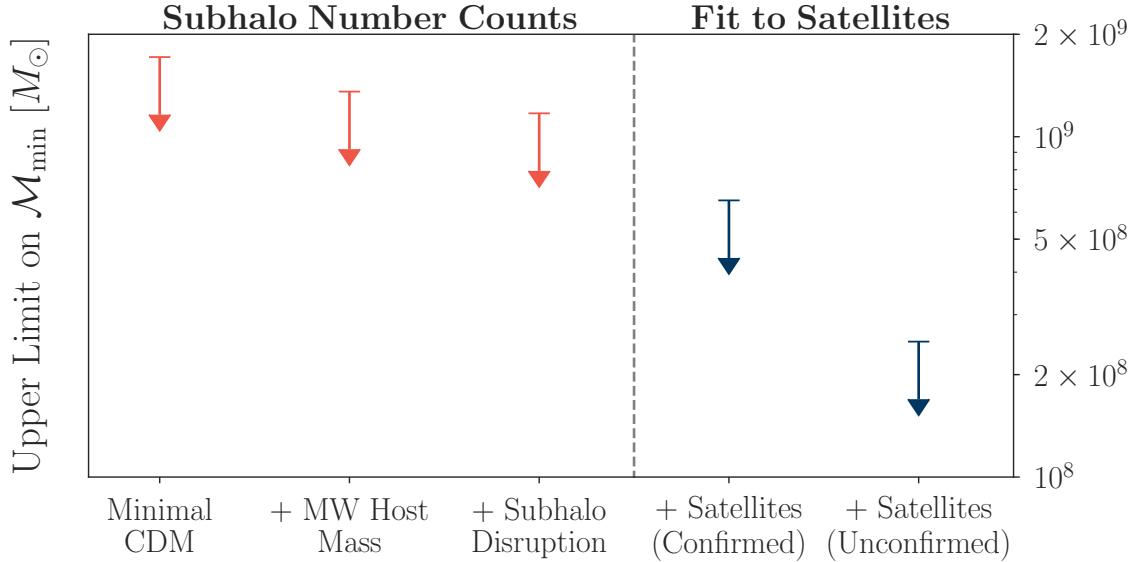


Figure 4.7: Impact of modeling assumptions on the minimum subhalo mass inferred from the observed DES and PS1 satellite populations. The first three models match the number of subhalos to the number of confirmed DES and PS1 satellites, and the last two models populate subhalos with galaxies to fit the position-dependent MW satellite luminosity function and size distribution.

analysis, we set the upper limit on M_{hm} equal to our upper limit on the minimum halo mass, i.e., $M_{\text{hm}} < 3.2 \times 10^8 M_{\odot}$, which corresponds to a lower limit on the half-mode scale of $k_{\text{hm}} > 36 h \text{ Mpc}^{-1}$. Using the relations in [335] with the cosmological parameters $h = 0.7$ and $\Omega_m = 0.286$ corresponding to the simulations used in our analysis (see Section 4.4.1), this yields a lower limit of 3.4 keV on the mass of thermal relic warm dark matter and an upper limit of $6 \times 10^{-30} \text{ cm}^2$ on the velocity-independent dark matter–baryon scattering cross section for a 10 keV dark matter particle mass, both at 95% confidence. We leave a detailed investigation of dark matter constraints to future work.

4.9 Theoretical Uncertainties

We aim to present a thorough galaxy–halo connection model that allows us to marginalize over the most important theoretical uncertainties when modeling the MW satellite population. Nonetheless, our modeling choices necessarily affect the predicted number of detected low-mass halos and thus our upper limit on the minimum halo mass, \mathcal{M}_{min} . In this section, we briefly discuss the main uncertainties in this analysis and their impact on our \mathcal{M}_{min} constraint.

To do so, we consider the upper limit on \mathcal{M}_{min} as a function of modeling assumptions, starting with the most conservative model possible and adding one assumption at a time. We illustrate the results of this exercise in Figure 4.7 for upper limits calculated as follows.

1. Minimal CDM. Assuming a maximally massive MW halo given Gaia constraints (i.e., a virial mass of $1.8 \times 10^{12} M_\odot$; [102, 106, 292, 291]), count the subhalos within the virial radius of the MW in order of decreasing V_{peak} until the number of kinematically confirmed DES and PS1 satellites is matched, and set the lowest corresponding value of $\mathcal{M}_{\text{peak}}$ equal to the upper limit on \mathcal{M}_{min} .
2. MW Host Mass. Repeat the previous step with the MW host mass fixed to its average value in our two fiducial simulations (i.e., an average virial mass of $1.4 \times 10^{12} M_\odot$).
3. Subhalo Disruption. Repeat the previous step many times with the subhalo number weighted by disruption probability, sampling \mathcal{B} from our fiducial posterior, to calculate an upper limit on \mathcal{M}_{min} at 95% confidence.
4. Satellites (confirmed). Repeat the previous step including the observational detection probabilities for mock satellites in the DES and PS1 footprints by drawing satellite properties from our fiducial posterior.
5. Satellites (unconfirmed). Repeat the previous step including the unconfirmed candidate satellites detected by DES and PS1 in the observed tally.

This yields $\mathcal{M}_{\text{min}} < (17, 14, 12, 6.5, 2.5) \times 10^8 M_\odot$ for models (i)–(v), respectively. Note that models (i)–(iii) are extremely conservative, since subhalos are counted in order of decreasing V_{peak} ; however, these models do not reproduce the observed position-dependent MW satellite luminosity function or radial distribution. Model (iv) yields the conservative limit presented in Appendix C.3.1, and model (v) yields our fiducial constraint, uncorrected for MW host halo mass. Although we have not explicitly considered artificial subhalo disruption in this list of theoretical uncertainties (e.g., [453, 454]), our fiducial orphan satellite model effectively assumes that subhalo disruption in dark matter–only simulations is entirely artificial, which is a conservative choice.

Figure 4.7 shows that both fitting the satellite luminosity function and including the population of faint, kinematically unconfirmed satellite galaxies in our fit yield significant increases in constraining power. We emphasize that our galaxy–halo connection model is conservative from the perspective of upper limits on the minimum halo mass, since we assume that high-mass halos host the brightest observed satellite galaxies. Moreover, we marginalize over many uncertainties in the connection between low-mass halos and faint galaxies. Thus, the largest gain in constraining power likely results from our detailed use of observational selection functions, i.e., from the fact that some satellites are not detected in DES or PS1 data. Given our extensive validation of the DES and PS1 selection functions in Paper I, we are therefore confident in our minimum halo mass constraints.

4.10 Conclusions

We have presented the results of a forward-modeling framework for MW satellites applied to recent searches for satellites in photometric surveys over nearly the entire high Galactic latitude sky. Our analysis includes position-dependent observational selection effects that faithfully represent satellite searches in DES and PS1 imaging data, and our galaxy–halo connection model allows us to marginalize over theoretical uncertainties in the relationship between galaxy and halo properties, the effects of baryonic physics on subhalo populations, and the stochastic nature of galaxy formation in low-mass halos. By performing a Bayesian analysis of the observed DES and PS1 satellite populations, we find decisive statistical evidence for the following.

1. The LMC impacts the observed MW satellite population, contributing 4.8 ± 1.7 (1.1 ± 0.9) LMC-associated satellites to the DES (PS1) satellite populations.
2. The LMC fell into the MW within the last 2 Gyr.
3. The faintest satellites currently known occupy halos with peak virial masses less than $3.2 \times 10^8 M_\odot$.
4. The faintest detectable satellites (i.e., dark matter–dominated systems with $M_* > 100 M_\odot$) occupy halos with peak virial masses greater than $10^6 M_\odot$.

These results have broad implications for galaxy formation and dark matter physics. For example, comparing our inferred luminosity function and galaxy occupation fraction to predictions from hydrodynamic simulations will help break degeneracies among subgrid star formation and feedback models. Meanwhile, extending our model to study the evolution of the luminosity function will shed light on high-redshift faint galaxy populations (e.g., [87, 475]) and the MW’s local reionization history (e.g., [101, 304, 257]).

Finally, our statistical detection of low-mass halos translates directly into constraints on a suite of dark matter properties, including warmth in thermal production scenarios, initial velocity distribution in nonthermal production scenarios, self-interaction cross section, interaction strength with the Standard Model, formation redshift, stability, and quantum mechanical behavior on astrophysical scales. Exploring the interplay between galaxy formation physics and alterations to the standard CDM paradigm will be crucial in order to extract these signals from upcoming observations of ultra-faint galaxies, and forward-modeling approaches like the one developed here will drive these studies forward.

Chapter 5

Placing the Milky Way in Context

Abstract

The Satellites Around Galactic Analogs (SAGA) Survey aims to measure the satellite galaxy populations of 100 Milky Way (MW) analogs in the very low-redshift Universe ($z \sim 0.01$) down to an absolute magnitude of $M_{r,o} = -12.3$. Here, we test for consistency between the galaxy–halo connection constraints from previous chapters—which are derived using only the MW satellite population—and SAGA Stage II measurements of 36 MW analog satellite populations. We show that this model provides a reasonable statistical description of SAGA Stage II data, indicating that our galaxy–halo connection constraints are cosmologically representative for satellites brighter than $M_{r,o} = -12.3$. We also forecast the galaxy–halo connection constraining power of the complete SAGA Stage III sample of 100 MW analogs, which will break degeneracies among the baryonic disruption and galaxy–halo size relation components of the model while reducing uncertainties in the inferred stellar mass–halo mass relation in the regime of the brightest MW satellites.

5.1 Paper Status and External Contributions

Section 5.3 of this chapter is published in slightly modified form in the *Astrophysical Journal*, Volume 907, Issue 2, p.85-119 with the title, “The SAGA Survey. II. Building a Statistical Sample of Satellite Systems around Milky Way–like Galaxies,” on which I am a co-author; the rest of this chapter is unpublished and authored by me. This work is the result of a collaboration with members of the Satellites Around Galactic Analogs (SAGA) Survey and particularly Yao-Yuan Mao, Risa Wechsler, and Marla Geha. I am grateful to the entire SAGA team for their input and support.

5.2 Satellites of Milky Way Analogs and the SAGA Survey

The galaxy–halo connection constraints derived in the previous chapters and the dark matter constraints that follow this chapter are based on Milky Way (MW) satellite galaxies alone, which necessitates a modeling approach that accounts for the specific characteristics of the MW system that influence its satellite population. It is therefore crucial to understand the extent to which the MW satellite population is cosmologically representative to draw robust and general inferences about galaxy formation and dark matter physics. For example, several authors have argued that the abundance and properties of MW satellites may be biased relative to satellite galaxies of comparable luminosities around typical MW-mass hosts (e.g., [374, 242]). There are also potential differences between characteristics of MW and M31 satellites including their radial profiles (e.g., [492]) and star formation histories (e.g., [477, 476]), although comparisons between these populations are complicated by their differing host halo and central galaxy properties and formation histories.

Until recently, our census of satellites around MW-mass central galaxies has been dominated by the MW (~ 60 satellites; [149]), M31 (~ 25 satellites; [319]), and a plethora of systems in the Sloan Digital Sky Survey (SDSS) main galaxy sample (~ 1 satellite per host; [401]). Only satellites comparable in luminosity to the Large Magellanic Cloud (LMC; $M_{r,o} \sim -19$) have been detected around these SDSS centrals, while satellites as faint as $M_{r,o} \sim 0$ have been measured nearby in the MW. New observations of satellites around Local Volume (LV) hosts using surface brightness fluctuations, which are sensitive to satellites as faint as $M_{r,o} \sim -9$, have uncovered ~ 5 satellites per MW-mass central down to this luminosity threshold [104, 105]. The luminosity functions of these LV satellite systems are in reasonable agreement with Λ CDM predictions, including for the host-to-host scatter among satellite populations at comparable host halo mass, contrary to some previous claims (e.g., [55]). Meanwhile, their radial distributions are potentially more concentrated than expected, although it is challenging to robustly predict satellite radial distributions due to both numerical challenges in N -body simulations (e.g., [205]) and the nontrivial connection between satellite galaxies and subhalos (e.g., [277] and references therein).

The Satellites Around Galactic Analogs (SAGA) Survey [193, 316] is designed to bridge the gaps in MW analog sample size and satellite luminosity threshold left by current observations. In particular, SAGA aims to measure the satellite populations of 100 MW analog systems between ~ 20 to 40 Mpc down to $M_{r,o} = -12.3$, corresponding to the Milky Way’s fifth-brightest satellite, Leo I. This sample size and satellite luminosity threshold will yield a precise measurement of the cosmological distribution of satellite populations around Milky Way-mass central galaxies outside of the LV, thereby placing MW, M31, and LV measurements in context. Stage I of the SAGA Survey [193] presented complete spectroscopic coverage around 8 MW analogs, discovering 27 new satellite galaxies, and Stage II [316] increased this sample to 36 complete hosts and 127 satellites.

The MW satellite luminosity function is consistent with being drawn from the distribution of

SAGA Stage II luminosity functions, foreshadowing the consistency of the corresponding galaxy–halo connection constraints described below. Several independent hydrodynamic simulations also predict satellite populations that are reasonably consistent with SAGA Stage I (e.g., [404]) and Stage II (e.g., [179, 160]) data. Intriguingly, the fraction of quenched satellites around SAGA centrals is significantly lower than that observed for MW or M31 satellites (and that predicted by hydrodynamic simulations) down to the same luminosity threshold. Placing this finding into a theoretical context while accounting for the disparate observational techniques used to measure the respective quenched fractions is beyond the scope of this work.

Here, we focus on the implications of current and future SAGA satellite abundances for the galaxy–halo connection. Section 5.3 compares the predictions of the galaxy–halo connection model presented in Chapters 3–4, which is constrained only by the MW satellite population, with SAGA Stage II data. Section 5.4 then forecasts the galaxy–halo connection constraining power of the entire SAGA Stage III sample of 100 MW analogs.

5.3 SAGA Stage II: Implications for the Galaxy–Halo Connection

5.3.1 The Model

To place SAGA Stage II results in the context of galaxy–halo connection models, we combine an abundance-matching procedure that associates SAGA host galaxies with dark matter halos with the empirical satellite galaxy model from Chapters 3–4 [336, 333]. This model populates subhalos with satellite galaxies by extrapolating an abundance-matching relation between luminosity and peak maximum circular velocity (calibrated to the GAMA luminosity function from [295] for $M_{r,o} < -13$) and a galaxy size–halo size relation (calibrated to [275] for galaxies with half-light radii $r_{1/2} \gtrsim 1$ kpc) into the regime of faint satellites. In addition, it includes a model for satellite disruption due to a central disk potential, which is calibrated to the Feedback in Realistic Environments (FIRE) baryonic simulations. We update the disruption prescription from [336], which was originally based on the machine-learning algorithm from [337] and the FIRE simulations in [190], using the analytic fit to subhalo disruption in the FIRE simulations presented in [404].

As described in [336, 333], this model accurately describes the luminosity, size, and radial distribution of MW satellite galaxies, including the population of satellites accreted with the LMC, when combined with recent observational MW satellite selection functions [149]. To incorporate theoretical uncertainties in our prediction for SAGA satellite populations, we sample from the posterior distribution over the galaxy–halo connection and baryonic disruption model parameters presented in [333], which is derived by fitting the model to most of the known MW satellite population.

To predict the satellite population for each SAGA host, we first follow the procedure in [193] to

map each SAGA host galaxy to a set of potential dark matter host halos in the `c125-2048` cosmological simulation (a higher-resolution version of the box used in [315]) using abundance matching. We assume that the scatter in central luminosity at fixed halo properties is 0.15 dex, and we select all halos mapped to absolute magnitudes within 0.15 mag of the SAGA host galaxy in question, which roughly corresponds to the quadrature sum of the estimated photometric and distance errors for our hosts. We also impose the host M_K cut and the environment cut described in [316].¹ This procedure yields ~ 2000 halos per SAGA host galaxy, on average, from which we select a random subset of 300 matched halos for each model realization to capture the probabilistic relation between host galaxy and host halo properties. In particular, these mock host halos have the cosmologically representative range of masses and formation histories that is expected for each host galaxy luminosity, which allows us to estimate the uncertainty in satellite populations resulting from the scatter in these quantities at a fixed central luminosity.

For each potential host halo in a given model realization, we then select all other halos above a resolution threshold of $V_{\text{peak}} = 40 \text{ km s}^{-1}$ that satisfy the $\pm 300 \text{ kpc}$ and $\pm 275 \text{ km s}^{-1}$ projected distance and velocity criteria in various projections of our simulation to mimic the SAGA satellite definition; thus, our model self-consistently includes interloping galaxies. We populate (sub)halos with (satellite) galaxies using the model described above, and we measure the “observed” satellite population for each host halo matched to a given SAGA host galaxy for several draws of satellite model parameters from the [333] posterior. We perform three sets of mock observations of galaxies in the $z = 0$ snapshot of the simulation by projecting the simulation along perpendicular axes and imposing the $M_{r,o} < -12.3$ SAGA completeness threshold and a surface brightness completeness limit of $\mu_{\text{eff}} < 26 \text{ mag arcsec}^{-2}$.² Mock-observed satellites are weighted by their survival probability according to our baryonic subhalo disruption model, and mock-observed interlopers are assumed to have survival probabilities of unity.

With these procedures, the model predicts that, at 68% confidence, SAGA hosts inhabit dark matter halos with virial masses in the range $7 \times 10^{11} < M_{\text{vir}}/\text{M}_{\odot} < 2 \times 10^{12}$ and that SAGA satellites are hosted by halos with peak virial masses in the range $2 \times 10^{10} < M_{\text{peak}}/\text{M}_{\odot} < 10^{11}$, and typical present-day virial masses a factor of $\sim 1\text{--}6$ lower.

5.3.2 Comparing Model Predictions with SAGA Results

Figure 5.1 shows the comparison between the observed and predicted distributions of the total number of satellites ($M_{r,o} < -12.3$) among the 36 complete SAGA systems. The error bars on the observed number are generated using the incompleteness correction procedure described in [316].

¹We do not implement the stellar foreground cuts in this mock selection because our simulation does not include stars or a Galactic disk. We do not implement the maximum halo mass cut because the abundance-matching procedure yields very few potential SAGA host halos with present-day virial mass $> 10^{13} \text{ M}_{\odot}$.

²The model predicts a very small number of satellites with $M_{r,o} < -12.3$ and $\mu_{\text{eff}} > 26 \text{ mag arcsec}^{-2}$; however, there are no satellites or satellite candidates at this low surface brightness in the SAGA object catalog or in the matched Hyper Suprime-Cam catalog.

To generate the predicted distribution, we repeatedly draw sets of mock host halos corresponding to the 36 SAGA systems; thus, the “predicted” contour includes both statistical uncertainties (due to the limited host sample size) and systematic uncertainties (due to scatter in the host galaxy–halo connection and our satellite model). Since this plot shows predicted and incompleteness-corrected realizations of the SAGA sample, adjacent data points are anticorrelated. We find good agreement between the total predicted N_{sat} distribution and SAGA observations.

To compare our predictions to the data in more detail, Figure 5.2 shows predicted luminosity functions, and Figure 5.3 shows predicted radial distributions (normalized to the number of satellites with $M_{r,o} < -12.3$ within a projected distance of 300 kpc) compared to the observed radial distribution for each complete SAGA system. Again, contours indicate 68% and 95% confidence intervals due to scatter in the host galaxy–halo connection, draws from our satellite galaxy model parameters, and projections of our simulation. Our predictions are largely consistent with the observed luminosity functions, and they are in excellent agreement with the observed normalized radial distributions. However, for SAGA hosts with the largest numbers of observed satellites, our model underpredicts both the total number of satellites, particularly the number of *bright* ($M_{r,o} < -15$) satellites, for which SAGA observations are highly complete (e.g., see the right panel of Figure 5 in [316]).

To quantify this potential bright-end tension, we estimate that 0.13 (0.09) predicted satellites with $M_{r,o} < -15$ must be added to our fiducial prediction per host to bring it into agreement with the data at 68% (95%) confidence assuming Poisson errors on the observed counts, which corresponds to an $\sim 1\sigma$ – 2σ discrepancy. This is consistent with Figure 5.2, which shows that roughly five SAGA hosts have one additional bright satellite relative to our predicted 95% confidence interval. There is also a hint that the model overpredicts the number of dim satellites, although it is formally consistent with the data given our current incompleteness estimates.

Overall, it is encouraging that the model predictions are in broad agreement with SAGA data, reinforcing our finding that the MW satellite population is not highly atypical.

5.3.3 Implications

Before we discuss the implications of the bright-end tension noted above, we emphasize several caveats associated with our current predictions. The satellite model we employed has specifically been fit to the MW satellite population; thus, our predictions for subhalo and satellite galaxy disruption do not account for the varying masses, morphological properties, and host halo density profiles of SAGA host galaxies relative to the MW. In addition, we have not included “orphan” satellites in the model³, nor have we explored exactly how our predictions depend on galaxy–halo connection parameters within the region of parameter space allowed by the MW satellite population.

³In N -body simulations, a dark matter subhalo may be disrupted by tidal stripping earlier than in reality due to the lack of concentrated baryonic content or spurious numerical effects. Modeling this effect is commonly referred to as including “orphan” satellites.

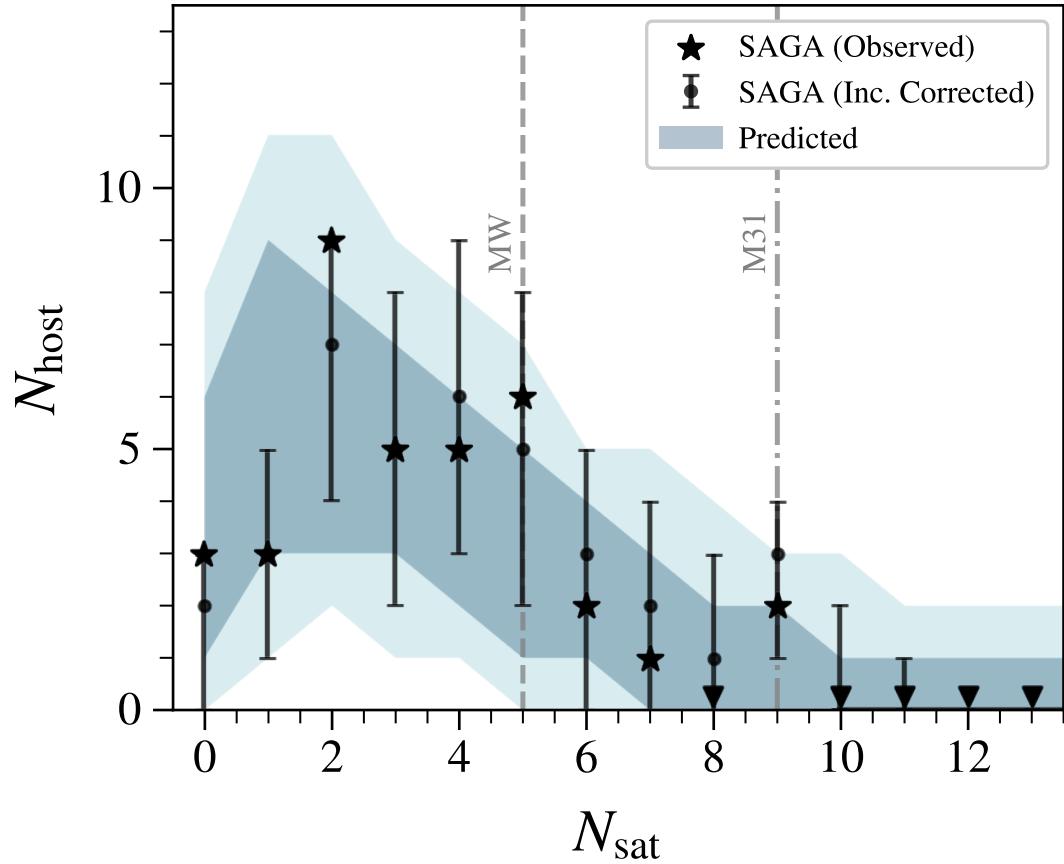


Figure 5.1: Distribution of satellite number in SAGA hosts. The black stars show the data from the SAGA Stage II sample of 36 complete hosts, and the error bars indicate incompleteness corrections as described in [316] (note that the incompleteness-corrected N_{sat} measurements are highly correlated). Dark blue (light blue) contours indicate the predicted 68% (95%) confidence intervals based on our simulation and galaxy–halo connection model. The MW (M31) is shown as a dashed (dashed–dotted) gray line.

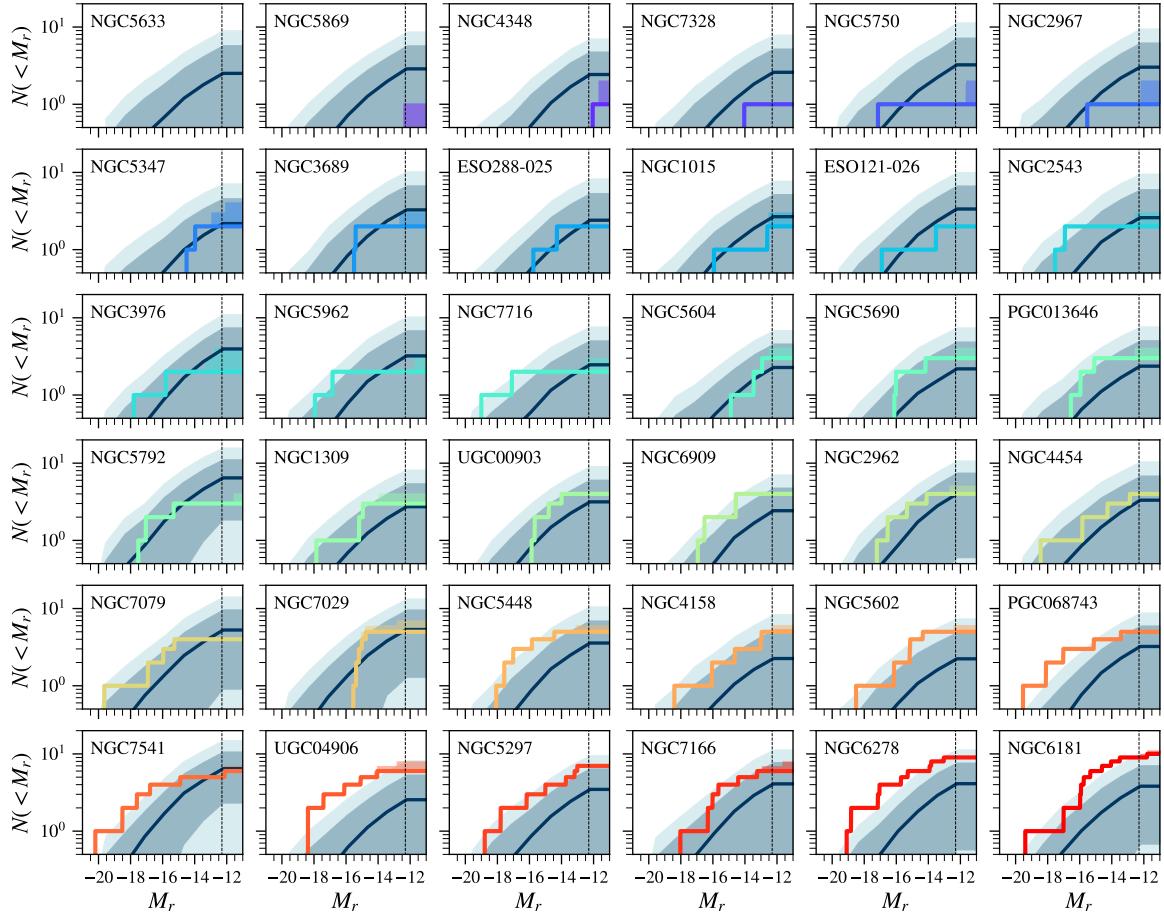


Figure 5.2: SAGA Stage II satellite luminosity functions and incompleteness corrections (colored lines and bands) compared to predictions from a cosmological dark matter–only simulation populated with galaxies using the empirical satellite model in [336, 333], which has been fit to the MW satellite population. Dark blue lines indicate the mean prediction for each satellite population, and dark blue (light blue) contours indicate 68% (95%) confidence intervals, which include the effects of host galaxy–halo abundance-matching scatter, uncertainty in our galaxy–halo connection model (see Section 5.3.1), and projection effects.

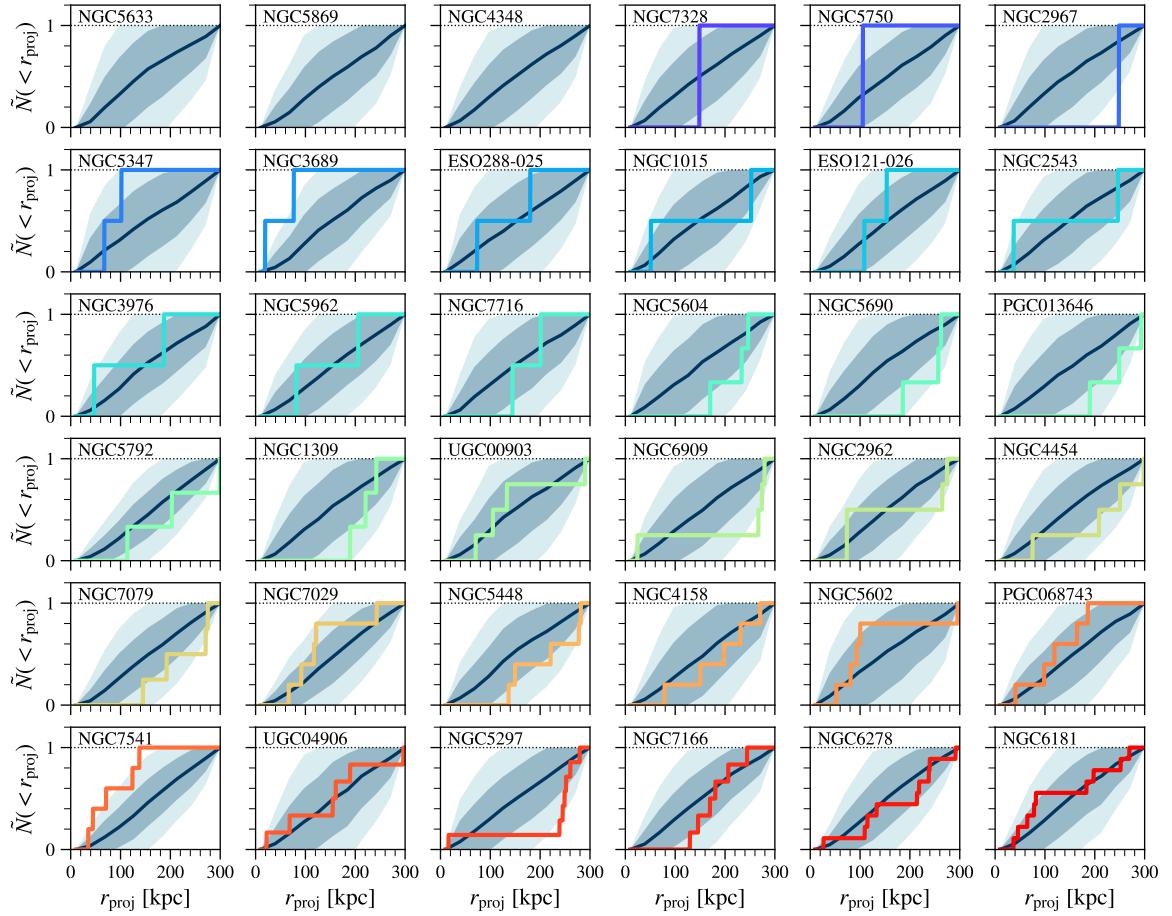


Figure 5.3: Comparison of predicted and observed radial satellite distributions, normalized to the number of satellites within a projected distance of 300 kpc. The predicted mean and confidence interval is identical to Figure 5.2, and the observed radial distributions are computed using all satellites above our $M_{r,o} < -12.3$ absolute magnitude limit. Hence, no observed data are shown for systems that do not have any satellites with $M_{r,o} < -12.3$. We do not correct the SAGA radial distributions because the incompleteness model in [316] predicts that missing satellites should have the same distribution as confirmed ones.

Nonetheless, the potential bright-end discrepancy is reminiscent of similar tensions noted for LV field [339] and satellite [104] galaxies when compared with models that differ from ours in detail, hinting at a more systematic issue that may exist in various theoretical predictions, including those from hydrodynamical simulations. The flexibility of our model allows us to quantify the possible sources of the tension and study potential solutions. We therefore briefly describe possible solutions in the context of our model, leaving a more thorough investigation to future work.

1. *Stellar mass–halo mass relation.* Forcing the halos that host the richest SAGA satellite systems to be a factor of ~ 3 more massive resolves the bright-end discrepancy. However, this will significantly alter the well-validated stellar mass–halo mass relation in the MW-mass regime [474].⁴ Alternatively, this halo mass shift can be achieved if SAGA host magnitudes are systematically biased dim by ~ 0.5 mag, which we also regard as unlikely.
2. *Disruption model.* Removing subhalo disruption from our model reduces the bright-end tension to less than $\sim 1\sigma$. It is unlikely that subhalos that host bright satellites undergo no disruption; however, disruption prescriptions calibrated to hydrodynamical zoom-in simulations are dominated by low-mass subhalos and therefore might overpredict disruption efficiencies for more massive subhalos.
3. *Global luminosity function.* Our abundance-matching prediction is calibrated to the GAMA luminosity function down to $M_{r,o} = -13$; this luminosity function contains both statistical errors (captured by uncertainties in the Schechter function fit to GAMA data) and potential systematic errors (e.g., due to survey incompleteness). While there is no evidence that the GAMA survey is incomplete down to $r = -19.8$, we find that varying the GAMA luminosity function amplitude from [295] within its quoted 2σ error can fully resolve the bright-end tension.

In future work, we plan to address these questions in detail and to explicitly test for consistency between the MW and SAGA satellite populations by refitting the satellite model to the SAGA data. We also plan to compare SAGA results with additional models and hydrodynamical simulations.

5.4 Forecasts for SAGA Stage III

Stage III of the SAGA Survey aims to measure complete luminosity functions for 100 MW analogs down to $M_{r,o} = -12.3$, roughly tripling the Stage II sample size. Using MW, LV, and SAGA satellite populations to jointly constrain our galaxy–halo connection model will require a careful treatment of the modeling discrepancies discussed in Section 5.3.3. Here, we simply forecast the galaxy–halo

⁴We note that varying the 0.15 dex scatter by ± 0.1 dex in our host abundance-matching relation does not significantly affect the bright-end tension.

connection constraining power of SAGA Stage III data by evaluating our model on mock SAGA satellite populations as follows.

First, we populate host halos in `c125-2048` with central galaxies using the abundance-matching procedure described above with a fixed scatter of 0.15 dex. Each abundance-matching realization yields $\sim 10^4$ halos that can correspond to one or more SAGA centrals, assuming the same M_K measurement uncertainty of 0.15 mag described above. For a given set of galaxy–halo connection model parameters, we then predict observed satellite populations (including interlopers) following the same mock procedure used for the SAGA Stage II comparison above. We modify our satellite model to efficiently encompass many mock SAGA satellite populations by writing the likelihood of observing a set of SAGA satellite populations, \mathbf{s}_{SAGA} (specified by their absolute magnitudes and surface brightnesses), given a set of model parameters $\boldsymbol{\theta}$ as

$$P(\mathbf{s}_{\text{SAGA}}|\boldsymbol{\theta}) = \prod_{\text{hosts } i} \prod_{\text{bins } j} P(n_{\text{SAGA},ij}|\hat{\mathbf{n}}_{\text{SAGA},ij}), \quad (5.1)$$

where $n_{\text{SAGA},ij}$ is the observed number of satellites in SAGA host i in bin j , and $\hat{\mathbf{n}}_{\text{SAGA},ij}$ is a vector of the number of mock SAGA satellites corresponding to host i in bin j from a realization of our model at fixed $\boldsymbol{\theta}$. For each SAGA central, we weight mock SAGA hosts by the probability they fall within the host’s measured M_K range assuming either a step function or Gaussian likelihood (our results are not sensitive to the shape of this likelihood). The explicit form of $P(n_{\text{SAGA},ij}|\hat{\mathbf{n}}_{\text{SAGA},ij})$ is given in Equation C.7.

To explore how these mock SAGA predictions respond to our satellite model parameters, we predict observed satellite populations using the planned SAGA Stage III central luminosity distribution, sampling from halos in a sub-volume of $\sim (30 \text{ Mpc } h^{-1})^3$ from `c125-2048` for computational efficiency. Figure 5.4 shows the dependence of the average number of observed satellites per host for a SAGA Stage III-like survey, $\langle N_{\text{sat}}(M_{r,o} < -12.3, \mu < 26) \rangle$, where N_{sat} is fully forward-modeled including projection effects. The baryonic disruption efficiency (\mathcal{B}) and galaxy–halo size relation scatter ($\sigma_{\log R}$) produce the largest variations in $\langle N_{\text{sat}} \rangle$ within the parameter ranges allowed by MW satellites, indicating that SAGA Stage III data will complement or improve current constraints on these parameters.⁵ Nonetheless, we emphasize that comparing the constraints derived from MW, LV, and SAGA data will be fruitful even for parameters that are primarily constrained by nearby satellites to test for consistency.

Finally, we fold these mock satellite populations through our likelihood framework to forecast the constraining power of SAGA Stage III data. As expected from Figure 5.4, we find that the abundance-matching parameters α and σ_M as defined in [333] are not well constrained from SAGA data alone, producing variations in the log-likelihood of $\Delta \ln \mathcal{L} \lesssim -0.5$ relative to its maximum. Nonetheless, SAGA data may inform the stellar mass–halo mass relation slope at higher luminosities

⁵Note that we do not vary galaxy occupation fraction parameters because they are best constrained by galaxies fainter than the SAGA satellite luminosity limit.

than captured by the current definition of α as well as the mass dependence of the luminosity scatter. Meanwhile, our forecast indicates that the input value of the disruption efficiency \mathcal{B} can be measured moderately well ($\Delta \ln \mathcal{L} \sim -1$). This measurement will complement MW constraints because hydrodynamic simulations used to calibrate subhalo disruption models (e.g., Chapter 2) are tuned to match aspects of the MW satellite population. Large values of the satellite size amplitude $\mathcal{A} \gtrsim 200$ pc are ruled out because they produce too few observed satellites, and the remaining size model parameters are not strongly constrained.

Thus, our results indicate that SAGA Stage III results will inform both subhalo disruption efficiency as a function of central galaxy properties and environment at fixed stellar mass and the size–virial radius relation in a regime where it is not well constrained (e.g., see [275, 244]). In addition, addressing the potential tensions discussed in Section 5.3.3 may help calibrate the luminosity function and stellar mass–halo mass relation for $-12 \lesssim M_{r,o} \lesssim -19$; these measurements will be systematically pursued in future work. Finally, we note that these updated galaxy–halo connection measurements are not expected to significantly improve the constraints on microphysical dark matter properties derived in the following chapters, which are driven by the population statistics of ultra-faint dwarf galaxies that have only been detected as satellites of the MW.

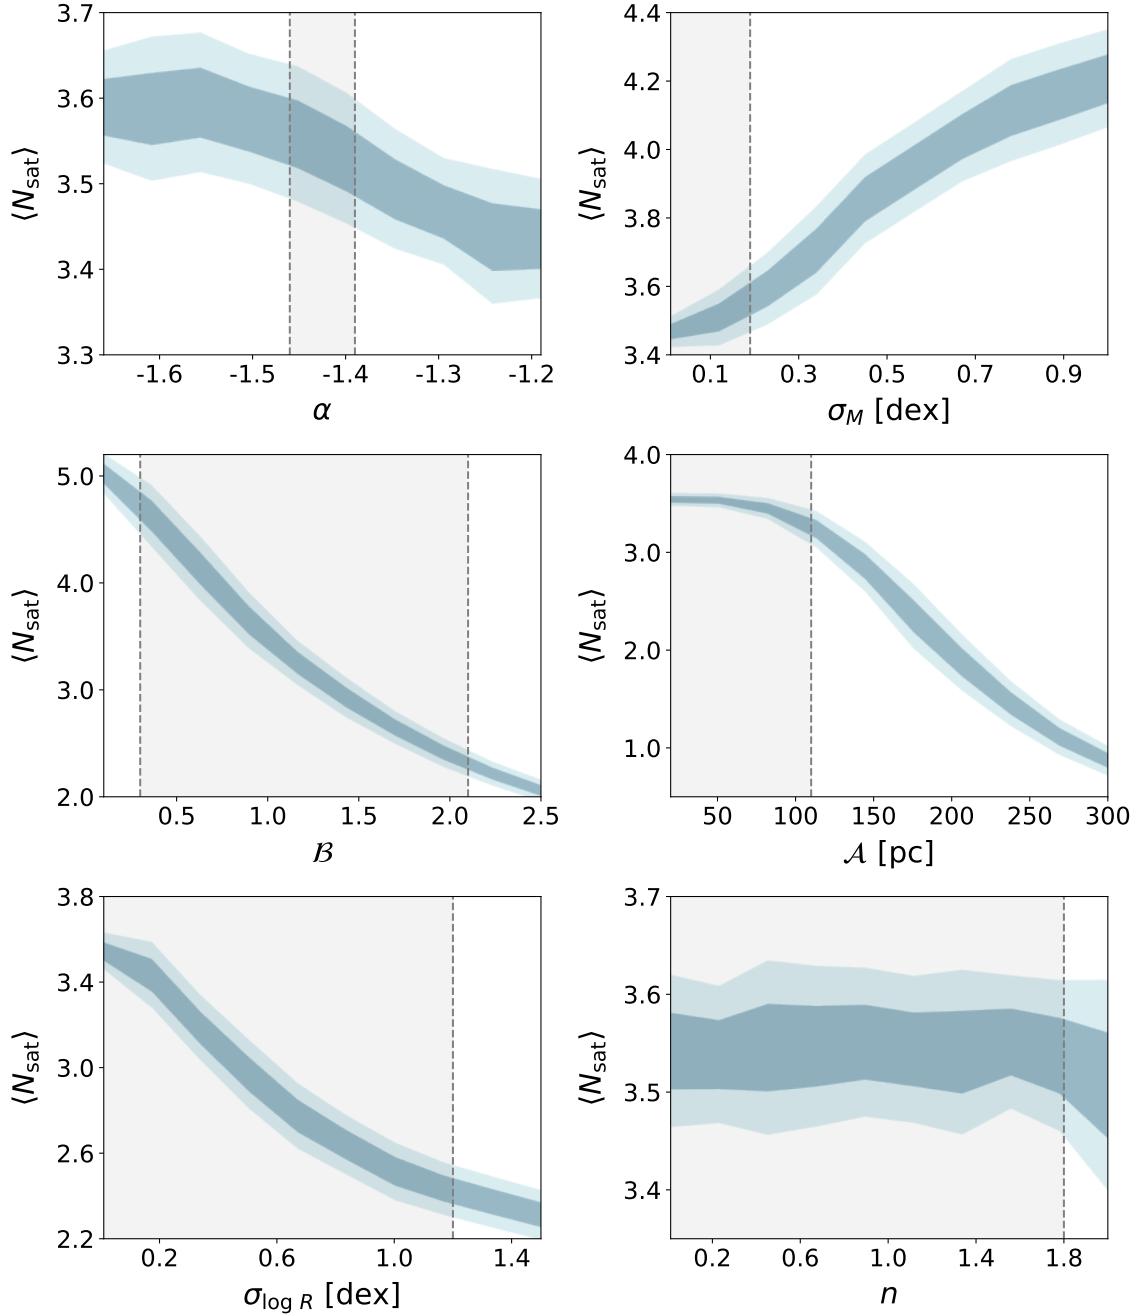


Figure 5.4: Forecasts for the average number of satellites per MW analog detected by a SAGA Stage III-like survey. Dark (light) blue bands show the predicted 68% (95%) confidence intervals for the average number of satellites predicted for 100 mock SAGA observations as a function of one galaxy–halo connection model parameter. In each panel, the remaining parameters are fixed to their best-fit value from our MW satellite analysis in Chapter 4 [333] and vertical dashed lines show the corresponding 95% confidence intervals inferred from MW satellites.

Chapter 6

Dark Matter Constraints. I. DM–SM Interactions

Abstract

Alternatives to the cold, collisionless dark matter (DM) paradigm in which DM behaves as a collisional fluid generically suppress small-scale structure. Herein we use the observed population of Milky Way (MW) satellite galaxies to constrain the collisional nature of DM, focusing on DM–baryon scattering. We first derive analytic upper limits on the velocity-independent DM–baryon scattering cross section by translating the upper bound on the lowest mass of halos inferred to host satellites into a characteristic cutoff scale in the linear matter power spectrum. We then confirm and improve these results through a detailed probabilistic inference of the MW satellite population that marginalizes over relevant astrophysical uncertainties. This yields 95% confidence upper limits on the DM–baryon scattering cross section of $2 \times 10^{-29} \text{ cm}^2$ ($6 \times 10^{-27} \text{ cm}^2$) for DM particle masses m_χ of 10 keV (10 GeV); these limits scale as $m_\chi^{1/4}$ for $m_\chi \ll 1 \text{ GeV}$ and m_χ for $m_\chi \gg 1 \text{ GeV}$. This analysis improves upon cosmological bounds derived from cosmic-microwave-background anisotropy measurements by multiple orders of magnitude over a wide range of DM masses, excluding regions of parameter space previously unexplored by other methods, including direct-detection experiments. Our work reveals a mapping between DM–baryon scattering and other alternative DM models, and we discuss the implications of our results for warm and fuzzy DM scenarios.

6.1 Paper Status and External Contributions

This chapter is published in slightly modified form in the *Astrophysical Journal Letters*, Volume 878, Issue 2, p.L32-37 with the title, “Constraints on Dark Matter Microphysics from the Milky Way Satellite Population,” on which I am the corresponding author. An Erratum that corrects a minor

error in the transfer functions predicted by our Boltzmann solver is published in the *Astrophysical Journal Letters*, Volume 897, Issue 2, p.L46–47, and these corrections are reflected below. This work is the result of a collaboration with Vera Gluscevic, Kimberly Boddy, and Risa Wechsler. Vera and Kim led the development of the DM–proton interaction Boltzmann solver, and all of us worked together to interface these models with Milky Way satellite data; in addition, Vera, Kim, and Risa all made editorial contributions to the text. This work was stimulated by workshops hosted by the Dark Energy Science Collaboration’s Dark Matter Working Group, and I am grateful to all of the members of this group for their support and insights.

6.2 Introduction

In the standard cold, collisionless dark matter (CDM) paradigm, structure forms from initial conditions described by a featureless power spectrum of matter perturbations, giving rise to a present-day dark matter (DM) halo mass function that extends uninterrupted down to sub-solar masses [204, 137]. Non-standard DM scenarios, such as warm DM (WDM; [4, 11]) and fuzzy DM (FDM; [227, 228]), often involve smoothing of matter perturbations and suppression of structure on small scales.

Herein we focus on the framework inspired by the weakly interacting-massive-particle paradigm, in which DM scatters with Standard Model particles (i.e., baryons), through a velocity-independent contact interaction. These interactions transfer heat and momentum between the photon–baryon and DM fluids in the early universe, and damp matter perturbations on scales that enter the cosmological horizon while scattering is efficient (e.g., [67]). The damping scale is set by the interaction strength, and the cumulative effect is largest for the smallest modes, which spend the most time inside the horizon.

As perturbations grow, this distinct fingerprint of DM microphysics is propagated to visible tracers of matter throughout cosmic history as a suppression of small-scale structure relative to CDM. Currently, the best cosmological limits on DM–baryon scattering come from measurements of damping tails in the cosmic-microwave-background (CMB) temperature and polarization power spectra from *Planck* [67, 198, 489] and the Lyman- α forest flux power spectrum from the Sloan Digital Sky Survey (SDSS; [156, 489, 352]). Moving beyond probes of quasi-linear cosmological perturbations, the suppression imprinted on the matter power spectrum at early times leads to an underabundance of collapsed objects—notably, low-mass DM halos in our Galactic neighborhood, and the faint galaxies that reside within them. As low-mass halos arise from matter fluctuations on scales far smaller than those captured by the CMB and other high-redshift probes, population studies of nearby dwarf galaxies could deliver dramatic improvements in sensitivity to DM–baryon interactions.

In this work, we analyze the observed population of Milky Way (MW) satellite galaxies to

place stringent limits on velocity-independent DM–baryon elastic scattering.¹ We present both an analytic derivation of these limits and a rigorous likelihood analysis in which we marginalize over uncertainties related to the connection between galaxies and halos and the impact of baryonic physics on subhalo abundances. Our population analysis yields 95% confidence upper limits on the DM–baryon scattering cross section of $(2 \times 10^{-29}, 8 \times 10^{-29}, 3 \times 10^{-28}, 6 \times 10^{-27}) \text{ cm}^2$ for DM particle masses of $(10^{-5}, 10^{-3}, 10^{-1}, 10) \text{ GeV}$, improving upon CMB limits by multiple orders of magnitude (see Figure 6.1). This analysis probes unexplored regions of DM parameter space, and it is complementary to direct-detection constraints that rely on the local DM distribution.

This Letter is organized as follows: in Section 6.3, we derive analytic limits on DM–baryon scattering from the existence of low-mass halos; in Section 6.4, we present our likelihood analysis of the MW satellite population and the corresponding DM–baryon scattering limits; in Section 6.5, we translate our results into constraints on other non-CDM models; we discuss caveats and conclude in Section 6.6. Throughout, we adopt the best-fit *Planck* cosmology with Hubble parameter $h = 0.6727$, DM density $\Omega_m h^2 = 0.1199$, and baryon density $\Omega_b h^2 = 0.0222$ [367], and we set $c = k_B = 1$.

6.3 Analytic Estimate from Individual Halos

To develop physical intuition for the effects of DM–baryon interactions on the late-time population of DM halos, we first derive an analytic estimate for the mass of the smallest halo allowed to form in an interacting cosmology.² We then translate upper limits on the minimum halo mass into upper limits on the DM–baryon scattering cross section. The limits we obtain in this Section do not depend strongly on observational completeness corrections or on galaxy–halo connection modeling, as demonstrated in Section 6.3.2.

6.3.1 Minimum Halo Mass in an Interacting Cosmology

In a non-standard cosmology with DM–baryon interactions, linear matter perturbations smaller than a critical length scale $\lambda_{\text{crit}} = 2\pi/k_{\text{crit}}$ are substantially suppressed relative to CDM. As the universe expands, collision damping affects progressively larger scales; when scattering becomes inefficient, the DM and baryon fluids kinetically decouple. The size of the largest perturbation entirely erased by DM collisions corresponds to the size of the cosmological horizon when the rate of momentum transfer R_χ between the DM and baryon fluids drops below the Hubble rate aH

$$aH = R_\chi \mid_{z=z_{\text{crit}}}, \quad (6.1)$$

¹We consider DM–proton interactions, and therefore constrain both the spin-independent and spin-dependent DM–nucleon scattering cross section. Because we neglect helium, our constraints for spin-independent scattering are conservative for DM masses above $\sim 1 \text{ GeV}$ [68].

²Note that [70, 71] perform similar derivations for DM–photon and DM–neutrino scattering models.

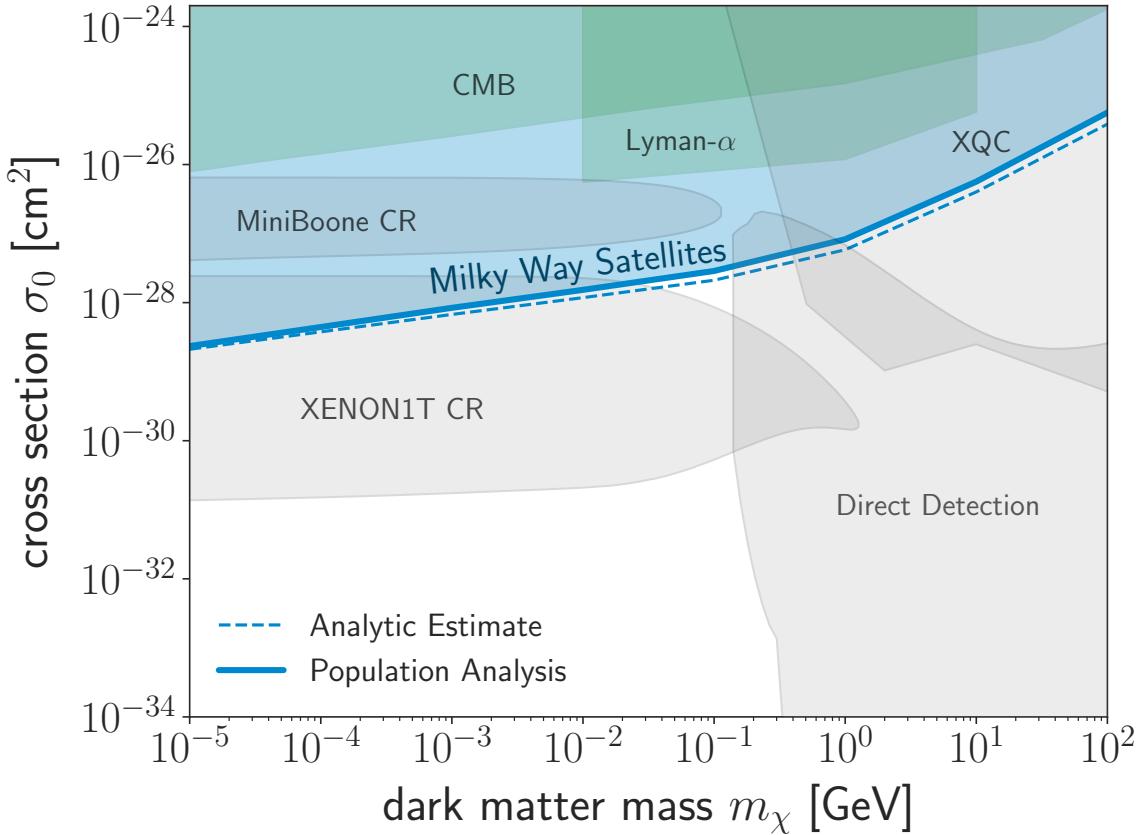


Figure 6.1: Upper limits on the velocity-independent DM–proton scattering cross section as a function of DM particle mass. The blue shaded region is excluded by the population of classical and SDSS-discovered MW satellites with 95% confidence by our likelihood analysis, which marginalizes over relevant astrophysical uncertainties (Section 6.4). The dashed line shows analytic upper limits derived from the existence of the lowest-mass halos hosting satellites (Section 6.3). Green contours show cosmological constraints from the CMB [67, 198] and the Lyman- α forest [489]. Gray contours show experimental constraints from cosmic-ray scattering [88], the X-ray Quantum Calorimeter (XQC; [162]), and direct-detection experiments including CRESST-III [118], the CRESST 2017 surface run [27], and XENON1T [31], as interpreted by [159]. Limits from Galactic center gas clouds [61, 463] overlap with parts of the XQC, CMB, and Lyman- α contours for $10^{-3} \text{ GeV} \lesssim m_\chi \lesssim 100 \text{ GeV}$, and are omitted for clarity.

where z_{crit} is the corresponding redshift. To calculate z_{crit} , we use the expression for the velocity-averaged momentum transfer rate for velocity-independent DM–proton scattering [67]

$$R_\chi = \mathcal{N}_0 a \rho_b Y_p \frac{\sigma_0}{m_\chi + m_p} \left(\frac{T_b}{m_p} + \frac{T_\chi}{m_\chi} \right)^{1/2}, \quad (6.2)$$

where $\mathcal{N}_0 \equiv 2^{7/2}/(3\sqrt{\pi})$; a is the scale factor; ρ_b is the baryon energy density; Y_p is the proton mass fraction; σ_0 is the velocity-independent DM–proton scattering cross section; m_χ is the DM particle mass; m_p is the proton mass; and T_b and T_χ are the temperatures of the baryon and DM fluids, respectively. The term $(T_b/m_p + T_\chi/m_\chi)^{1/2}$ is the thermal dispersion of the relative velocity between the DM and baryon fluids. At early times, the interactions keep the fluids in thermal equilibrium. In particular, the heat transfer rate $R'_\chi \equiv (m_\chi/(m_\chi + m_p))R_\chi$ exceeds the Hubble rate, driving the DM temperature to that of the photon–baryon fluid, $T_\chi = T_b = T_0(1 + z)$, where T_0 is the CMB temperature today. Thermal decoupling occurs at z_{th} , when the heat transfer rate decreases sufficiently, such that

$$aH = R'_\chi \Big|_{z=z_{\text{th}}}. \quad (6.3)$$

Note that z_{th} occurs deep within the radiation-dominated era, when $H \approx H_0 \sqrt{\Omega_{\text{rad}}} z^2$, where $\Omega_{\text{rad}} \approx 10^{-4}$ is the radiation energy density divided by the critical density today, $\bar{\rho}$. After thermal decoupling ($z < z_{\text{th}}$), DM cools adiabatically and $T_\chi = T_0(1 + z)^2/(1 + z_{\text{th}})$.

Solving Equation (6.3) for z_{th} , accounting for the thermal evolution of the DM and baryon fluids in Equation (6.2), and making appropriate substitutions in Equation (6.1), we can find z_{crit} given the parameters σ_0 and m_χ that describe our DM–baryon scattering model. Next, requiring the critical mode to undergo a full oscillation within the cosmological horizon, we compute the corresponding wavenumber k_{crit} via

$$k_{\text{crit}} = 2 \left(\frac{1}{aH} \right)^{-1} \Big|_{z=z_{\text{crit}}}. \quad (6.4)$$

Modes with $k > k_{\text{crit}}$ in the linear matter power spectrum are extremely suppressed (see Figure 6.2).

The next step is to determine the mass of a collapsed halo arising from perturbation modes that correspond to k_{crit} . For this purpose, we use the relation

$$M_{\text{crit}} = \frac{4\pi}{3} \rho_m \left(\frac{\lambda_{\text{crit}}}{2} \right)^3 = \frac{4\pi}{3} \Omega_m \bar{\rho} \left(\frac{\pi}{k_{\text{crit}}} \right)^3. \quad (6.5)$$

Combining with Equation (6.4), this yields

$$M_{\text{crit}}(\sigma_0, m_\chi) \approx \frac{\pi^4}{6} \frac{(\mathcal{N}_0 Y_p \Omega_b)^2 \Omega_m \bar{\rho}^3}{(H_0 \sqrt{\Omega_{\text{rad}}})^5} \frac{T_0}{m_p} \left(\frac{\sigma_0}{\tilde{m}_\chi} \right)^2, \quad (6.6)$$

where \tilde{m}_χ asymptotes to m_χ for $m_\chi \gg m_p$ and to $(m_\chi m_p^3)^{1/4}$ for $m_\chi \ll m_p$; the expression for intermediate DM mass does not have a closed form. Halos that are less massive than $M_{\text{crit}}(\sigma_0, m_\chi)$ do

not form in a cosmology with DM–baryon interactions because power vanishes on the corresponding scales due to early-time scattering.

6.3.2 Limits on the Interaction Cross Section

If halos are detected down to a minimum mass M_{\min} , interpreting an upper limit on M_{\min} as an upper limit on M_{crit} yields an upper bound on σ_0 , at fixed m_χ . Thus, detecting halos with $M < M_{\text{crit}}$ excludes cross sections that correspond to k_{crit} .

Many independent astrophysical probes testify to the existence of low-mass halos and set upper limits on M_{\min} . For example, substructure detections in strongly lensed systems [456, 217] and the dynamical masses of dwarf galaxies obtained from spectroscopy (e.g., [426]) both imply that halos exist down to a mass of $\sim 10^9 M_\odot$. Recent studies of the MW satellite population that model the galaxy–halo connection, completeness corrections of observed satellites, and the impact of baryonic physics on galaxy formation and subhalo abundances have pushed the upper bound on M_{\min} even lower. For example, Chapter 3 [336] reports $M_{\min} < 5.4 \times 10^8 M_\odot$ with 95% confidence using classical and SDSS-discovered MW satellites, and [241] derived consistent results. These studies do not include ultra-faint satellites discovered in recent years [42, 150, 148, 272, 274, 282, 283, 222]; accounting for these systems will further improve limits on M_{\min} . In addition, upcoming surveys including The Rubin Observatory Legacy Survey of Space and Time (LSST; [301]) are expected to discover even more faint MW satellites, which could lower M_{\min} by a factor of ~ 5 compared to current constraints [151].

The observed MW satellite population likely contains several halos that are near the current limit on the minimum halo mass. Thus, we set $M_{\text{crit}} < 5.4 \times 10^8 M_\odot$ to derive a bound of $k_{\text{crit}} > 30 h \text{ Mpc}^{-1}$. Using this scale in our analytic prescription yields $\sigma_0 < 2 \times 10^{-29} \text{ cm}^2$ for a DM particle mass of 10 keV. Constraints for other DM particle masses are shown in Figure 6.1.

It is important to note that these limits scale weakly with M_{\min} . In particular, the upper limit on σ_0 scales as $M_{\text{crit}}^{1/2}$, so increasing M_{\min} by an order of magnitude only weakens the limit on σ_0 by a factor of ~ 3 . Thus, these constraints do not depend sensitively on the galaxy–halo connection and completeness correction modeling used to derive M_{\min} , and they shift by small amounts if halo masses obtained from strong lensing or spectroscopic measurements are instead chosen as the reference.

6.4 Satellite Population Likelihood Analysis

Halo formation is affected at masses well above M_{crit} due to the gradual power suppression caused by DM–baryon interactions (see Figure 6.2). Thus, using a *population* of low-mass halos should yield more stringent limits than the existence of a single low-mass halo. In this section, we therefore perform a probabilistic analysis of the MW satellite population to place more realistic constraints

on DM–baryon scattering; our results are shown in Figure 6.1.

In principle, accounting for the detailed effects of DM–baryon scattering on the late-time halo population requires simulations that self-consistently include both the initial linear power suppression described above as well as late-time DM–baryon interactions. However, we find that the power suppression in the DM–baryon scattering case is remarkably similar to that in WDM, as shown in Figure 6.2. Moreover, we expect late-time interactions to be a small effect for the interaction model that we consider. Taken together, these facts allow us to use the results of WDM simulations run with nearly identical initial conditions as the DM–baryon scattering model under consideration.

To derive the correspondence between DM–baryon scattering and WDM, we use a modified version of the Boltzmann solver **CLASS** (described in [67, 68, 198]), which evolves linear cosmological perturbations in the presence of DM–baryon interactions. We generate linear matter power spectra as a function of σ_0 and m_χ , and we compare these to WDM power spectra using the transfer function in [415] by matching the *half-mode scale* k_{hm} , i.e., the wavenumber at which the transfer function $T(k) = (P_{\text{collisional}}(k)/P_{\text{CDM}}(k))^{1/2}$ is equal to 50%. We compute $k_{\text{hm}}(\sigma_0, m_\chi)$ numerically using our Boltzmann solver by varying σ_0 at several values of m_χ .

The correspondence shown in Figure 6.2 allows us to map the suppression in MW subhalo abundances found in WDM simulations to our interacting cosmology. In particular, we use the subhalo mass function from [299], which is fit to cosmological zoom-in simulations of thermal relic sterile neutrino WDM

$$\frac{dN}{dM} \Big|_{\text{collisional}} = \frac{dN}{dM} \Big|_{\text{CDM}} \left(1 + \gamma \frac{M_{\text{hm}}}{M}\right)^{-\beta}, \quad (6.7)$$

where $\gamma = 2.7$, $\beta = 0.99$, M is the peak subhalo virial mass output by the halo finder, and M_{hm} is the mass corresponding to k_{hm} via a relation equivalent to Equation (6.5).

Next, to forward-model the MW satellite population, we modify the framework presented in Chapter 3 [336]. In particular, we supplement high-resolution DM-only simulations of MW-mass host halos [315] with a flexible model for the galaxy–halo connection and the impact of both baryonic physics and DM–baryon scattering on subhalo populations. We then fit the luminosity function of classical and SDSS-discovered MW satellites using a Poisson likelihood in bins of satellite luminosity. As in [336], free parameters in our fit include the slope and scatter in the galaxy–halo connection and the strength of subhalo disruption due to baryonic effects. These are nuisance parameters with large uncertainties that should be marginalized over for the purpose of placing robust limits on DM microphysics. We use the ratio of the collisional-to-CDM subhalo mass functions in Equation (6.7) to assign a “survival probability” to each subhalo in our CDM simulations, following [241]. Thus, the final free parameter in our fit is M_{hm} , and we obtain a marginalized posterior distribution $P(M_{\text{hm}})$ using a flat prior on $\log(M_{\text{hm}})$. To be conservative, we assume that *all* subhalos host galaxies, even though the galaxy occupation fraction is likely low and mass-dependent in this regime (e.g., [410, 175]). A non-trivial occupation fraction could force lighter subhalos to host observed satellites, further strengthening our constraints. However, we find that marginalizing over a step-function

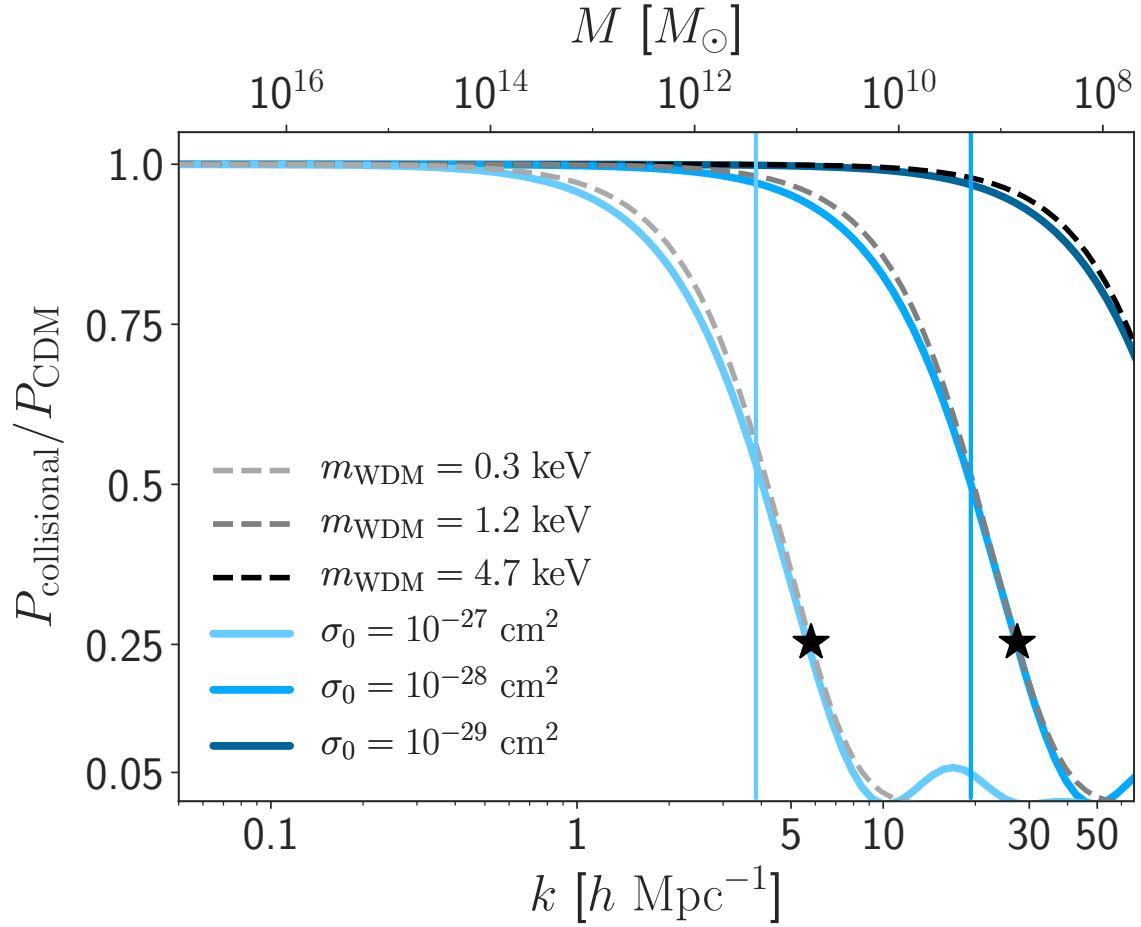


Figure 6.2: Ratio of the linear matter power spectrum in a DM–baryon scattering cosmology to that in CDM (solid lines), for a range of interaction cross sections (for 1 MeV DM particles). Dashed lines show the same quantity for WDM models with matching half-mode scales (denoted as black stars). Vertical lines indicate the critical scale discussed in Section 6.3.

galaxy formation threshold does not significantly affect our results.

Finally, we map $P(M_{\text{hm}})$ to $P(\sigma_0|m_\chi)$ using our half-mode scale calculation. We find $M_{\text{hm}} < 3.1 \times 10^8 M_\odot$ (corresponding to $k_{\text{hm}} > 36 h \text{ Mpc}^{-1}$) with 95% confidence, yielding upper limits on σ_0 of $(2 \times 10^{-29}, 8 \times 10^{-29}, 3 \times 10^{-28}, 6 \times 10^{-27}) \text{ cm}^2$ for DM particle masses of $(10^{-5}, 10^{-3}, 10^{-1}, 10) \text{ GeV}$, as shown in Figure 6.1. These limits improve upon CMB constraints by multiple of magnitude for $m_\chi \lesssim 1 \text{ GeV}$, and can be extrapolated to higher masses. However, for $m_\chi \lesssim 10 \text{ keV}$, relativistic effects become important. Moreover, in the following section we show that thermal DM lighter than $\sim 3 \text{ keV}$ is ruled out at all cross sections.

6.5 Implications for Other Non-CDM Models

The methods developed in this work apply to *any* model in which DM behaves similarly to a collisional fluid on small scales. For example, we can translate our M_{hm} constraint to a lower limit on WDM mass by combining the transfer function derived in [415] for a thermal relic sterile neutrino of mass m_{WDM} with the half-mode mass equivalent of Equation (6.5), which yields

$$m_{\text{WDM}} = 2.32 \left(\frac{\Omega_m}{0.25} \right)^{0.4} \left(\frac{h}{0.7} \right)^{0.8} \left(\frac{M_{\text{hm}}}{10^9 M_\odot} \right)^{-0.3} \text{ keV.} \quad (6.8)$$

We find $m_{\text{WDM}} > 3.26 \text{ keV}$ with 95% confidence, which is consistent with previous MW satellite results [241] and competitive with Lyman- α forest constraints [459, 232].

Our results also constrain FDM models, in which ultra-light axions comprise DM and small-scale structure is suppressed due to quantum interference effects [227, 228]. We translate our M_{hm} constraint into a lower limit on FDM mass m_ϕ using the $m_{\text{WDM}} - m_\phi$ relation from [33] in Equation (6.8), which gives

$$m_\phi = 1.3 \times 10^{-21} \left(\frac{\Omega_m}{0.25} \right)^{0.95} \left(\frac{h}{0.7} \right)^{1.9} \left(\frac{M_{\text{hm}}}{10^9 M_\odot} \right)^{-0.71} \text{ eV.} \quad (6.9)$$

We find $m_\phi > 2.9 \times 10^{-21} \text{ eV}$, which is again competitive with Lyman- α forest constraints [233]. We note that the high-redshift galaxy luminosity function and the epoch of reionization provide WDM and FDM constraints that are complementary to—but currently weaker than—our limits (e.g., [115]).

We expect that a re-analysis of the MW satellite population using our method will yield even more stringent limits for velocity-dependent DM–baryon scattering, i.e., for any scattering model in which the cross section scales as the relative particle velocity to a positive power. Such interactions arise in DM effective theory (e.g., [67]), and we leave an investigation of this scenario to future work.

6.6 Conclusions and Discussion

Small-scale tracers of the matter distribution in the universe provide insights into DM microphysics that are complementary to other probes. In this work, we report stringent upper bounds on DM–baryon scattering for a velocity-independent contact interaction by analyzing the population of classical and SDSS-discovered MW satellite galaxies. Based on the consistency of the satellite population with CDM expectations down to a halo mass scale of $\sim 10^8 M_\odot$, we place robust limits on a range of DM interaction cross sections and particle masses. Our analysis closes unexplored gaps in DM–baryon scattering parameter space, improving upon previous cosmological bounds by several orders of magnitude, and it is not subject to uncertainties in the local DM distribution or cosmic-ray propagation.

In our likelihood analysis, we have accounted for several astrophysical uncertainties, including the slope and scatter of the low-mass galaxy–halo relation and the impact of baryonic physics on subhalo abundances. We have made the conservative assumption that *all* subhalos host galaxies, though marginalizing over a step-function galaxy formation threshold does not affect our results. Moreover, [336] verified that potential spurious numerical effects in our simulations (e.g., artificial subhalo disruption) do not alter the inferred minimum halo mass.

There are several possible caveats to our analysis that we leave for future work. First, we have not investigated how our results depend on cosmological parameters. However, we expect uncertainties on cosmological parameters to play a minimal role; for example, Equation (6.6) implies that limits on σ_0 scale as $\Omega_m^{-1/2}$. Future analyses that marginalize over cosmological parameters are needed to confirm this weak dependence. Next, we relied on simulations with a narrow range of host halo mass, even though the uncertainty on the mass of MW is relatively large (e.g., [128]). We do not expect this uncertainty to affect our results significantly, as M_{\min} scales linearly with host mass and our limits depend weakly on M_{\min} . To further improve our constraints, it will be necessary to understand how a mass-dependent galaxy occupation fraction and baryonic effects beyond enhanced subhalo disruption affect satellite populations in detail.

This work demonstrates that the MW satellite population places competitive constraints on a class of non-CDM models. Moreover, it paves the way for joint probabilistic analyses of forthcoming small-scale structure datasets. For example, upcoming imaging surveys such as LSST will facilitate systematic searches for density gaps in Galactic stellar streams, which potentially trace even lower-mass subhalos than those inferred from satellites (e.g., [76]). In addition, gravitational lensing measurements with the Atacama Large Millimeter/submillimeter Array will map out low-mass DM substructure in a complementary manner to faint galaxies (e.g., [216]). Finally, future spectroscopic surveys like DESI [131] will provide improved Lyman- α forest measurements, allowing for important high-redshift consistency tests of local small-scale structure constraints. Distinct systematic and modeling uncertainties accompany each of these DM probes, and joint likelihood analyses of all available datasets—enabled by approaches similar to ours—will be crucial in order to study the

nature of DM in the era of next-generation surveys.

Chapter 7

Dark Matter Constraints. II. Warm, Interacting, & Fuzzy DM

Abstract

We perform a comprehensive study of Milky Way (MW) satellite galaxies to constrain the fundamental properties of dark matter (DM). This analysis fully incorporates inhomogeneities in the spatial distribution and detectability of MW satellites and marginalizes over uncertainties in the mapping between galaxies and DM halos, the properties of the MW system, and the disruption of subhalos by the MW disk. Our results are consistent with the cold, collisionless DM paradigm and yield the strongest cosmological constraints to date on particle models of warm, interacting, and fuzzy dark matter. At 95% confidence, we report limits on (i) the mass of thermal relic warm DM, $m_{\text{WDM}} > 6.5 \text{ keV}$ (free-streaming length, $\lambda_{\text{fs}} \lesssim 10 h^{-1} \text{ kpc}$), (ii) the velocity-independent DM-proton scattering cross section, $\sigma_0 < 8.8 \times 10^{-29} \text{ cm}^2$ for a 100 MeV DM particle mass (DM-proton coupling, $c_p \lesssim (0.3 \text{ GeV})^{-2}$), and (iii) the mass of fuzzy DM, $m_\phi > 2.9 \times 10^{-21} \text{ eV}$ (de Broglie wavelength, $\lambda_{\text{dB}} \lesssim 0.5 \text{ kpc}$). These constraints are complementary to other observational and laboratory constraints on DM properties.

7.1 Paper Status and External Contributions

This chapter is published in slightly modified form in Physical Review Letters, Volume 126, Issue 9, p.091101-091111 with the title, “Constraints on Dark Matter Properties from Observations of Milky Way Satellite Galaxies,” on which I am the corresponding author. It is the result of a collaboration with the Dark Energy Survey (DES) Milky Way Working Group (MWWG) and particularly Alex Drlica-Wagner, Keith Bechtol, Sidney Mau, Risa Wechsler, Vera Gluscevic, and Kimberly Boddy

and culminates the Milky Way satellite census led by Alex and Keith and the galaxy–halo connection interpretation presented in Chapter 4. Alex, Keith, Sid, and Risa all provided helpful input throughout the project and made editorial contributions to the text. As with other DES papers, this work would not have been possible without significant collaboration, and I am particularly grateful to the conveners and members of the MWWG (both past and present) for facilitating these efforts.

7.2 Introduction

In the concordance model of cosmology, collisionless cold dark matter (CDM) makes up $\sim 25\%$ of the matter-energy density of the Universe [15]. While dark matter (DM) has the potential to solve a number of outstanding challenges in the standard model (SM) of particle physics [361, 246, 59], the only positive empirical evidence for DM comes from cosmological and astrophysical observations. Furthermore, by studying the astrophysical distribution of DM, it is possible to probe its particle nature [97, 95]. Specifically, the formation, abundance, and structure of gravitationally bound DM structures, known as “halos,” provide valuable information about viable ranges of the DM particle mass, production mechanism, and couplings to the SM. In particular, the abundance and properties of the smallest DM halos have the potential to indicate a departure from the CDM paradigm [97, 95].

The smallest known DM halos host the ultrafaint dwarf satellite galaxies of the Milky Way (MW) [425]. In these systems, star formation is highly suppressed by reionization and stellar feedback, leading to mass-to-light ratios that are hundreds of times larger than the universal average [98, 425]. Ultrafaint satellite galaxies are, thus, pristine laboratories for studying DM; in particular, the abundance of these systems is a sensitive probe of *any* DM physics that suppresses the formation or present-day abundance of small halos [309, 364, 368, 26, 260, 3, 72, 398].

Here, we study the following theoretical paradigms for DM that affect the properties of the MW satellite population:

- (i) *Warm dark matter (WDM)* is produced in the early Universe with a temperature of $\mathcal{O}(1\text{ keV})$, although its momentum distribution can be nonthermal. Any viable WDM candidate must be cold enough to reproduce the observed large-scale structure, but its non-negligible free-streaming length suppresses the formation of the low-mass halos that host MW satellite galaxies [309, 368, 26, 3, 299, 260]. One of the most popular WDM candidates is a sterile neutrino [4, 11].
- (ii) *Interacting dark matter (IDM)* couples strongly enough to the SM to be heated by interactions with the photon-baryon fluid before recombination. This collisional damping washes out small-scale structure, even if the DM is produced nonthermally [70, 71, 335]. DM-nucleon interactions arise in generalizations of the weakly-interacting-massive-particle (WIMP) scenario [67, 68, 198], and the impact of DM-radiation interactions on low-mass halos has also been studied [72, 412, 167]. Here, we consider a velocity- and spin-independent DM-proton coupling, c_p .

- (iii) *Fuzzy dark matter (FDM)* consists of an ultralight boson with a sufficiently small mass, $\mathcal{O}(10^{-22} \text{ eV})$, such that its de Broglie wavelength is comparable to the sizes of dwarf galaxies, $\mathcal{O}(1 \text{ kpc})$; this inhibits the formation of low-mass halos due to the uncertainty principle [227, 228, 154, 153]. Ultralight axions constitute one popular class of FDM [317].

In this Chapter, we use novel measurements and modeling of the MW satellite galaxy population to constrain each DM paradigm described above. Specifically, we combine a census of MW satellites [149] from the Dark Energy Survey (DES) [5] and Pan-STARRS1 (PS1) [108] with a rigorous forward-modeling framework [333] to fit the position-dependent MW satellite luminosity function in each of these DM paradigms. This procedure fully incorporates inhomogeneities in the observed MW satellite population and marginalizes over uncertainties in the mapping between MW satellite galaxies and DM halos, the efficiency of subhalo disruption due to the MW disk, and the properties of the MW system.

Our analysis yields stringent constraints on each DM paradigm based on the abundance of observed MW satellites. These limits are complementary to constraints from the Lyman- α forest [459, 232, 233, 395], strongly lensed systems [226, 195], and MW stellar streams [40]. Our results imply that CDM is consistent with astrophysical observations down to the smallest currently accessible scales ($k \sim 40h \text{ Mpc}^{-1}$) and strongly reinforce previous work demonstrating that there is no discrepancy between the number of MW satellites predicted by CDM and current observations [264]. Throughout this work, we fix cosmological parameters at $h = 0.7$, $\Omega_m = 0.286$, $\Omega_\Lambda = 0.714$, $\sigma_8 = 0.82$, and $n_s = 0.96$ [218].

7.3 Analysis Overview

Before discussing our treatment of each DM paradigm in detail, we describe the main components of our analysis used to connect non-CDM scenarios to the observed MW satellite population. For each paradigm, we assume that the non-CDM component constitutes the entirety of the DM. Figure 7.1 illustrates how our analysis proceeds: Non-CDM physics suppresses the linear matter power spectrum on small scales (left panel), which manifests as an underabundance of subhalos (middle panel) and faint MW satellite galaxies (right panel) relative to CDM predictions.

7.3.1 Transfer Function

The linear matter power spectrum, normalized to that of CDM, is used to generate initial conditions for simulations of structure formation. In particular, the transfer function is defined as

$$T^2(k) \equiv \frac{P_{\text{DM}}(k)}{P_{\text{CDM}}(k)}, \quad (7.1)$$

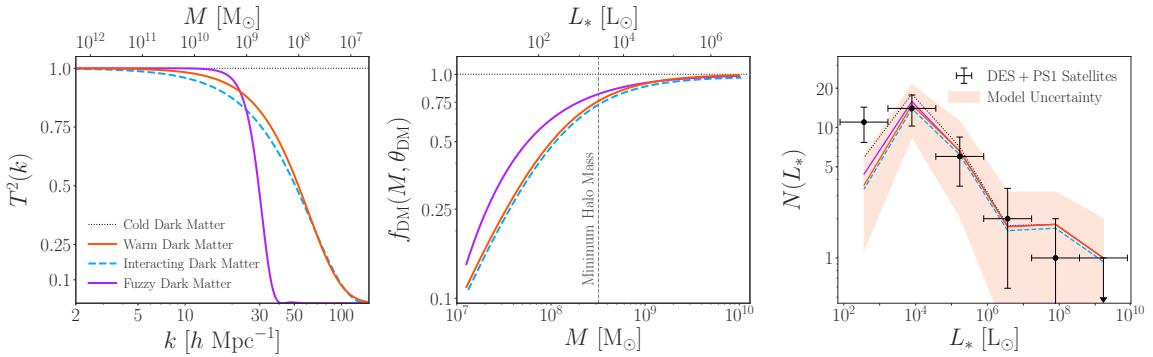


Figure 7.1: Left panel: Transfer functions for the WDM (orange), IDM (blue), and FDM (magenta) models that are ruled out by our analysis at 95% confidence, corresponding to $m_{\text{WDM}} = 6.5 \text{ keV}$, $\sigma_0 = 8.8 \times 10^{-29} \text{ cm}^2$ (for DM particle mass $m_\chi = 100 \text{ MeV}$), and $m_\phi = 2.9 \times 10^{-21} \text{ eV}$, respectively. These constraints are marginalized over our MW satellite model and the properties of the MW system. Middle panel: SHMF suppression relative to CDM for each ruled-out non-CDM model. The vertical dashed line indicates the 95% confidence upper limit on the lowest-mass halo inferred to host MW satellite galaxies [333]. Note that the IDM SHMF is assumed to be identical to the WDM SHMF in our analysis, and is offset slightly for visual clarity. Right panel: Predicted MW satellite galaxy luminosity functions for each ruled-out non-CDM model compared to DES and PS1 observations, evaluated at the best-fit MW satellite model parameters from Ref. [333]. The shaded band illustrates the uncertainty of our WDM prediction due to the stochasticity of our galaxy-halo connection model and the limited number of simulations used in our analysis; the size of this uncertainty is very similar to that in CDM and the other alternative DM models shown. This panel is a simple one-dimensional representation of our MW satellite and DM model fit to the luminosity, size, and spatial distribution of satellites in the DES and PS1 survey footprints. The comparison of our CDM model to data is described in Ref. [333], and full posterior distributions for our non-CDM analyses are provided in Supplemental Material [188, 30, 219].

where k is the cosmological wave number, $P_{\text{CDM}}(k)$ is the CDM linear matter power spectrum, and $P_{\text{DM}}(k)$ is the linear matter power spectrum of a non-CDM model [69]. $P_{\text{DM}}(k)$ is obtained by integrating the relevant Boltzmann equation (which may include DM-SM interactions) given the initial DM phase-space distribution. The left panel in Figure 7.1 illustrates the transfer function for the three DM paradigms we consider.

It is convenient to define the *half-mode scale* k_{hm} as the wave number satisfying $T^2(k_{\text{hm}}) = 0.25$ [415]. The corresponding *half-mode mass*,

$$M_{\text{hm}} = \frac{4\pi}{3} \Omega_m \bar{\rho} \left(\frac{\pi}{k_{\text{hm}}} \right)^3 \quad (7.2)$$

is a characteristic mass scale below which the abundance of DM halos is significantly suppressed relative to CDM. Here, $\bar{\rho}$ is the critical density of the Universe today.

7.3.2 Subhalo Mass Function

The abundance of subhalos within the virial radius of the MW is expressed as the cumulative number of subhalos as a function of subhalo mass M . We follow Chapter 6 [335] by using peak virial mass, defined according to the Bryan-Norman overdensity [93] with $\Delta_{\text{vir}} \simeq 99.2$ (consistent with our cosmological parameters). We define

$$\left(\frac{dN_{\text{sub}}}{dM} \right)_{\text{DM}} \equiv f_{\text{DM}}(M, \boldsymbol{\theta}_{\text{DM}}) \left(\frac{dN_{\text{sub}}}{dM} \right)_{\text{CDM}}, \quad (7.3)$$

where $f_{\text{DM}}(M, \boldsymbol{\theta}_{\text{DM}})$ is the suppression of the SHMF relative to CDM and $\boldsymbol{\theta}_{\text{DM}}$ are DM model parameters; both f_{DM} and $\boldsymbol{\theta}_{\text{DM}}$ depend on the DM model in question. The middle panel in Figure 7.1 shows SHMF suppression for the three DM paradigms we consider.

7.3.3 MW Satellite Model

Here, we describe the additions to our MW satellite model pertaining to the non-CDM paradigms described above. We comprehensively discuss the underlying galaxy-halo connection model in Supplemental Material. We combine the SHMF suppression in Equation 7.3 with a state-of-the-art satellite modeling framework [333] to predict the abundance of observed MW satellites in each DM paradigm. Our modeling framework combines cosmological zoom-in simulations of two halos from Ref. [315]—which are chosen to have masses, concentrations, and assembly histories similar to those inferred for the MW halo and include realistic analogs of the Large Magellanic Cloud system—with a statistical model of the galaxy-halo connection in order to populate subhalos with satellite galaxies.

We implement SHMF suppression by multiplying the detection probability of each mock satellite,

which includes terms that model tidal disruption due to the MW disk, the efficiency of galaxy formation, and observational detectability, by a factor of $f_{\text{DM}}(M, \theta_{\text{DM}})$, following Refs. [241, 335]. This procedure assumes that the shape of the observed radial satellite distribution (which our model predicts reasonably well [333]) is unchanged in alternative DM scenarios, which is consistent with results from cosmological WDM simulations of MW-mass halos [299, 80]. The validity of this assumption is less certain for FDM, because dynamical friction operates differently for wavelike versus particle DM [284], although this results in negligible differences in disruption timescales for the $\sim 10^8 M_{\odot}$ subhalos that drive our constraints [154]. The right panel in Figure 7.1 shows the predicted satellite luminosity function for each non-CDM model under consideration evaluated with model parameters that are ruled out at 95% confidence.

7.3.4 Fitting Procedure

We fit predicted satellite populations to the observed satellite population from DES and PS1 using the observational selection functions derived in Ref. [149], assuming that satellite surface brightness is distributed according to a Poisson point process in each survey footprint following Chapters 3–4 [336, 333]. We use the Markov chain Monte Carlo (MCMC) code `EMCEE` [180] to simultaneously fit for seven parameters governing the galaxy-halo connection, one parameter governing the impact of the MW disk on subhalo disruption, and one parameter governing the impact of the DM model in question, which we express as a subhalo mass scale. In particular, our thermal relic WDM constraint is derived by fitting for M_{hm} , and our FDM limit is derived by fitting for a characteristic mass scale M_0 . Further details on our fitting procedure are provided in Supplemental Material.

Subhalo abundance is known to scale linearly with host halo mass [315], and we assume that satellite luminosity is a monotonic function of subhalo mass, modulo scatter [333]. We therefore expect a higher-mass MW host halo to yield weaker constraints on non-CDM models, because observed satellites would inhabit correspondingly higher-mass subhalos. The average virial mass of the host halos in our two realistic MW-like simulations is $1.4 \times 10^{12} M_{\odot}$, which is consistent with the 95% confidence range for the virial mass of the MW halo inferred from *Gaia* measurements of satellite kinematics [102, 106]. To be conservative, we account for the uncertainty in MW halo mass on our DM constraints by assuming that the mass scale describing the suppression of the SHMF in each DM paradigm is linearly related to the virial mass of the MW halo, following the scaling for minimum halo mass derived in Ref. [333]. In particular, we multiply the upper limit on the characteristic mass scale in each of our non-CDM fits by the ratio of the largest allowed MW halo mass to the average host halo mass in our simulations. We validate this procedure by fitting the observed satellite population using each of our two MW-like simulations separately, which yields reasonable agreement with the linear scaling expectation. This conservative scaling mitigates the largest uncertainty associated with the limited statistics of our two realistic simulations.

In summary, our fit to the MW satellite population incorporates both intrinsic inhomogeneities

DM Model	Parameter	Constraint	Derived property	Constraint
WDM	m_{WDM}	$m_{\text{WDM}} > 6.5 \text{ keV}$	Free-streaming length	$\lambda_{\text{fs}} \lesssim 10 h^{-1} \text{ kpc}$
IDM	σ_0	$\sigma_0 < 8.8 \times 10^{-29} \text{ cm}^2$	DM-Proton coupling	$c_p \lesssim (0.3 \text{ GeV})^{-2}$
FDM	m_ϕ	$m_\phi > 2.9 \times 10^{-21} \text{ eV}$	de Broglie wavelength	$\lambda_{\text{dB}} \lesssim 0.5 \text{ kpc}$

Table 7.1: Constraints on the WDM, IDM, and FDM paradigms from observations of MW satellite galaxies. Limits for each non-CDM model are derived by assuming that it constitutes the entirety of the DM. The first column lists the DM paradigm, the second column describes the particle physics parameters constrained by this analysis, the third column lists the corresponding constraints at 95% confidence, the fourth column describes the derived property constrained for each DM model, and the fifth column lists constraints on the derived parameters. Limits on the DM-proton scattering cross sections depend on the DM particle mass, m_χ (see Figure 7.2); for simplicity, we present our constraint for $m_\chi = 100 \text{ MeV}$.

in the spatial distribution of MW satellites and those introduced by the varying coverage and depth of current surveys. We assume that alternative DM physics modifies only the SHMF, via Equation 7.3, and we report 95% confidence limits on DM model parameters that are marginalized over uncertainties in our MW satellite model and the properties of the MW system.

7.4 WDM Analysis

Thermal relic WDM with particle mass, m_{WDM} , has been studied extensively in the literature (e.g., Refs. [460, 299]) and serves as a benchmark model for our analysis.

Transfer function. The transfer function for thermal relic WDM is given as a function of m_{WDM} by Ref. [460]. This transfer function is commonly assumed in cosmological studies of WDM and facilitates a well-defined comparison to other small-scale structure results [459, 232, 226, 195, 40]. However, the simple thermal relic transfer function is inadequate to describe specific particle models of WDM, such as resonantly produced sterile neutrinos [296]. Thus, constraints on specific DM candidates must be inferred using transfer functions appropriate for the particle model in question, as we discuss below.

SHMF. Several authors have implemented the thermal relic WDM transfer function from Ref. [460] in cosmological zoom-in simulations to estimate the suppression of the SHMF in MW-mass host halos [415, 28, 299, 80]. These results depend on the algorithm used to remove spurious halos [467, 28] and, therefore, vary among studies. Following Ref. [297], SHMF suppression for thermal relic WDM can be expressed as

$$f_{\text{WDM}}(M, m_{\text{WDM}}) = \left[1 + \left(\frac{\alpha M_{\text{hm}}(m_{\text{WDM}})}{M} \right)^\beta \right]^\gamma, \quad (7.4)$$

where α , β , and γ are constants and M_{hm} is related to m_{WDM} in our fiducial cosmology via

$$M_{\text{hm}}(m_{\text{WDM}}) = 5 \times 10^8 \left(\frac{m_{\text{WDM}}}{3 \text{ keV}} \right)^{-10/3} M_{\odot}. \quad (7.5)$$

To facilitate comparison with recent WDM constraints from analyses of the MW satellite population [335], strong gravitational lenses [226, 195], and stellar streams [40], we adopt the SHMF from Ref. [299], which corresponds to Equation 7.4 with $\alpha = 2.7$, $\beta = 1.0$, and $\gamma = -0.99$. We note that the recent estimate of the SHMF from Ref. [297]—which specifically models resonantly produced sterile neutrino WDM—is significantly *less* suppressed than the thermal relic SHMF from Ref. [299]. Thus, our fiducial WDM constraint applies directly only to thermal relic DM.

Fitting procedure. We implement Equation 7.4 in our fit to the MW satellite population to obtain a marginalized posterior distribution over M_{hm} . In particular, we fit for $\log_{10}(M_{\text{hm}})$ using a uniform prior on this logarithmic quantity, and we translate the resulting limit to m_{WDM} using Equation 7.5. We translate our thermal relic WDM limit into constraints on resonantly produced sterile neutrinos by following Refs. [414, 308]. Specifically, we analyze sterile neutrino transfer functions over a grid of mass and mixing angle values [111], and we constrain sterile neutrino models that produce transfer functions which are strictly more suppressed than our 95% confidence ruled-out thermal relic WDM model. This procedure is described in detail in Supplemental Material.

7.5 IDM Analysis

Our treatment of IDM follows the prescription of Chapter 6 [335]. For concreteness, we focus on the case of velocity-independent DM-proton scattering.

Transfer function. Following Chapter 6 [335], the transfer function in our fiducial IDM model is obtained using the modified version of the Boltzmann solver CLASS described in Refs. [67, 68, 198], which we use to evolve linear cosmological perturbations in the presence of velocity-independent DM-proton interactions. These interactions are described by the velocity-independent scattering cross section σ_0 and the DM particle mass m_χ . As noted in Chapter 6 [335], transfer functions for this model are very similar to those of thermal relic WDM, modulo dark acoustic oscillations that occur at very small scales and are significantly suppressed for our parameter space of interest.

SHMF. Because cosmological zoom-in simulations including DM-proton scattering have not been performed, we follow Chapter 6 [335] by mapping the SHMF suppression of IDM to that of WDM based on the correspondence of the transfer functions. In particular, we match the half-mode scales in the transfer functions to construct a relation between m_{WDM} and (σ_0, m_χ) , and we assume that the IDM SHMF is identical to the corresponding thermal relic WDM SHMF from Ref. [299]. This procedure neglects late-time DM-proton scattering, which has a negligible impact on subhalo abundances in our IDM model, even in regions with high baryon densities.

Fitting procedure. Following Chapter 6 [335], we use the mapping procedure described above to

translate our 95% confidence limit on thermal relic WDM into limits on σ_0 for several values of m_χ in our fiducial IDM model.

7.6 FDM Analysis

Finally, we provide details on each step for the FDM paradigm. We focus on the case of ultralight scalar field DM with negligible self-interactions and SM couplings.

Transfer function. The FDM transfer function is given as a function of the FDM mass m_ϕ by Ref. [227]. We note that this transfer function features steeper power suppression than thermal relic WDM for a fixed half-mode scale.

SHMF. We assume that the FDM SHMF suppression takes the form of Equation 7.3, and we fit the results of the semianalytic model in Refs. [154, 153] with a function of the form

$$f_{\text{FDM}}(M, m_\phi) = \left[1 + \left(\frac{M_0(m_\phi)}{M} \right)^{\tilde{\beta}(m_\phi)} \right]^{\tilde{\gamma}(m_\phi)}, \quad (7.6)$$

where $\tilde{\beta}(m_\phi)$ and $\tilde{\gamma}(m_\phi)$ are provided in Supplemental Material. The characteristic subhalo mass scale M_0 is related to the FDM mass via [413]

$$M_0(m_\phi) = 1.6 \times 10^{10} \left(\frac{m_\phi}{10^{-22} \text{ eV}} \right)^{-4/3} \text{ M}_\odot. \quad (7.7)$$

The SHMF suppression in Equation 7.6 encapsulates the effects of tidal stripping on subhalos with solitonic cores, which was explicitly included by Refs. [154, 153]. This SHMF suppression is significantly less severe than that estimated from the FDM simulations in Ref. [413]. As described in Supplemental Material, using the SHMF from Ref. [413] in our fit yields a limit on the FDM mass that is roughly 3 times more stringent than our fiducial result. This confirms that the FDM SHMF is a key theoretical uncertainty that must be addressed [228].

Fitting procedure. We implement the SHMF in Equation 7.6 in our fit to the MW satellite population to obtain a marginalized posterior distribution over M_0 . In particular, we fit for $\log_{10}(M_0)$ using a uniform prior on this logarithmic quantity, and we translate the resulting limit to m_ϕ using Equation 7.7. We note that our procedure for constraining FDM uses the detailed *shape* of the SHMF suppression in this model rather than mapping the half-mode scale of the FDM transfer function to that of thermal relic WDM as in Chapter 6 [335] or bounding the FDM SHMF by ruled-out thermal relic WDM SHMFs as in Ref. [418]. This is necessary because both the shape of the FDM transfer function and the resulting suppression of the SHMF differ in detail from thermal relic WDM (see Figure 7.1).

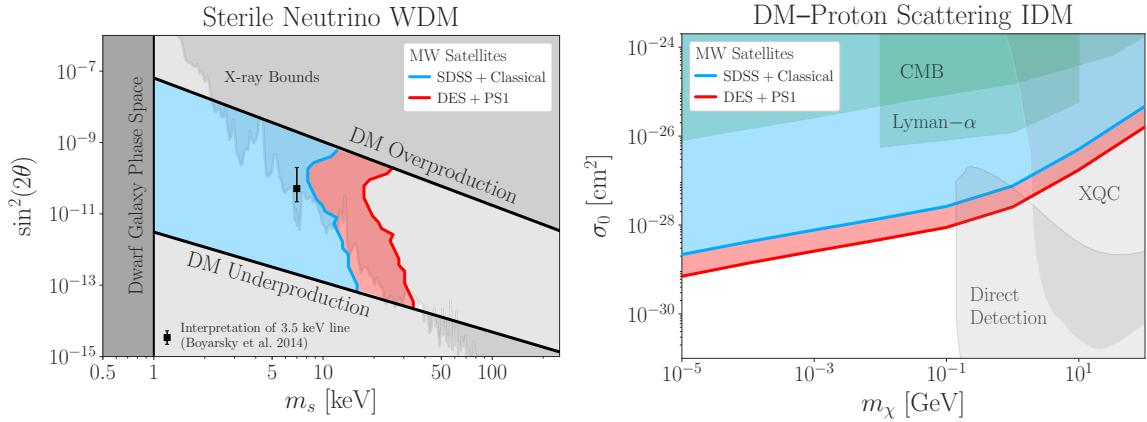


Figure 7.2: Exclusion regions for WDM and IDM models from our analysis of MW satellites observed with DES and PS1 (red) compared to previous constraints from classical and SDSS satellites [335] (blue) and other experimental results. Left panel: Constraints on the mass and mixing angle of resonantly produced sterile neutrino DM. These constraints are derived by finding mass and mixing angle combinations that suppress the linear matter power spectrum more strongly than the $m_{\text{WDM}} = 6.5$ keV thermal relic ruled out at 95% confidence by our analysis. The black point with error bars shows the sterile neutrino interpretation of the 3.5 keV x-ray line [82]. The dark gray region is ruled out by dwarf galaxy internal dynamics [83], and the gray contour shows x-ray constraints [225, 363, 134]. Solid black lines indicate regions of parameter space in which resonantly produced sterile neutrinos cannot constitute all of the DM in the neutrino minimal standard model [34, 414]. Right panel: Constraints on the interaction cross section and DM mass for velocity-independent DM-proton scattering. Green contours show cosmological limits from the CMB [67, 198] and the Lyman- α forest [489]. Light gray contours show experimental limits from the x-ray quantum calorimeter [314] and direct detection results as interpreted by Ref. [159].

7.7 Results

Table 7.1 presents our constraints on the WDM, IDM, and FDM paradigms. We describe these results below and translate the limits into constraints on specific models corresponding to each DM paradigm.

(i) *WDM*.—Our fit using the thermal relic WDM SHMF suppression from Ref. [299] yields $M_{\text{hm}} < 3.0 \times 10^7 M_{\odot}$, or $m_{\text{WDM}} > 7.0 \text{ keV}$, at 95% confidence. Linear scaling with MW halo mass yields our fiducial constraint of $M_{\text{hm}} < 3.8 \times 10^7 M_{\odot}$, corresponding to $m_{\text{WDM}} > 6.5 \text{ keV}$. This translates to an upper limit on the free-streaming length of $\lambda_{\text{fs}} \lesssim 10 h^{-1} \text{ kpc}$, corresponding to the virial radii of the smallest halos that host MW satellite galaxies, and improves on previous m_{WDM} constraints from the MW satellite population by a factor of ~ 2 [335].

Our constraint on thermal relic WDM translates to a lower limit of 50 keV on the mass of a nonresonant Dodelson-Widrow sterile neutrino [141, 460]. We also translate our thermal relic WDM limit into constraints on the mass and mixing angle of resonantly produced sterile neutrinos assuming a Shi-Fuller production mechanism [422], following the conservative procedure described above. As shown by the red exclusion region in the left panel in Figure 7.2, our analysis rules out nearly the entire remaining parameter space for resonantly produced sterile neutrinos in the neutrino minimal standard model [34] at greater than 95% confidence. (A small region of parameter space is not excluded at the lowest viable mixing angles and $m_s \gtrsim 30 \text{ keV}$.) In addition, we robustly rule out the resonantly produced sterile neutrino interpretation of the 3.5 keV x-ray line [82].

(ii) *IDM*.—Mapping our $m_{\text{WDM}} > 6.5 \text{ keV}$ constraint to the DM-proton scattering model following the procedure in Chapter 6 [335] yields constraints on the velocity-independent interaction cross section of $(7.0 \times 10^{-30}, 2.6 \times 10^{-29}, 8.8 \times 10^{-29}, 1.7 \times 10^{-27}) \text{ cm}^2$ for DM particle masses of $(10^{-5}, 10^{-3}, 10^{-1}, 10) \text{ GeV}$, respectively, at 95% confidence. As shown by the red exclusion region in the right panel in Figure 7.2, these constraints are highly complementary to direct detection limits, particularly at low DM masses [335]. We note that these constraints scale as $m_{\chi}^{1/4}$ (m_{χ}) for $m_{\chi} \ll 1 \text{ GeV}$ ($m_{\chi} \gg 1 \text{ GeV}$). At a DM mass of 100 MeV, our limit translates into an upper bound on the DM-proton coupling of $c_p \lesssim (0.3 \text{ GeV})^{-2}$ [67].

Despite our conservative marginalization over MW halo mass, these results improve upon those in Chapter 6 [335] by a factor of ~ 3 at all DM masses. This is stronger than the improvement expected from the analytic prediction for cross section constraints derived in Chapter 6 [335] due to a more precise determination of the SHMF, resulting from the sky coverage and sensitivity of DES and PS1.

Several complementary astrophysical and cosmological measurements probe the DM-proton scattering cross section. Stringent limits have been derived by reinterpreting direct detection

constraints in the context of cosmic ray upscattering [88]. We do not show these results in Figure 7.2, because they constrain the DM-proton scattering at relativistic energies, which precludes a straightforward mapping to the velocity-independent cross section constrained here. The IDM model we consider contributes to the energy density of relativistic species at big bang nucleosynthesis, which sets a lower on its mass that depends on the spin statistics of the DM particle [73, 346, 279]. Understanding the interplay of these results with our limits is an important area for future work.

- (iii) *FDM*.—We obtain $M_0 < 1.4 \times 10^8 M_\odot$ at 95% confidence from our fiducial FDM fit. Applying linear MW-host mass scaling yields $M_0 < 1.8 \times 10^8 M_\odot$ at 95% confidence, or $m_\phi > 2.9 \times 10^{-21}$ eV. This translates to an upper limit on the de Broglie wavelength of $\lambda_{\text{dB}} \lesssim 0.5 h^{-1}$ kpc, roughly corresponding to the sizes of the smallest MW satellite galaxies. Thus, the 10^{-22} eV FDM model invoked to reconcile the apparent mismatch between the predicted and observed inner dark matter density profiles of dwarf galaxies [228], and to fit the internal dynamics of low-surface-brightness [58, 89] and ultradiffuse [470] galaxies, is strongly disfavored by MW satellite abundances.

To connect to particle models of FDM, we plot this limit in the well-motivated parameter space of ultralight axion mass versus axion-photon coupling in Figure 7.3. For the range of axion-photon couplings that we consider, this mixing has a negligible effect on structure formation. We reiterate that our constraint was derived assuming a light scalar field without self-interactions; this assumption may be violated in specific ultralight axion models. Although our analysis and Lyman- α forest studies exclude a similar region of parameter space [233, 395], our work probes structure on complementary physical scales with distinct theoretical and observational systematics.

7.8 Discussion

In this Letter, we used a state-of-the-art model of the MW satellite galaxy population to place stringent and robust limits on three fundamental DM paradigms: WDM, IDM, and FDM. Although some of these alternative DM models gained popularity by solving apparent small-scale structure “challenges” facing CDM, recent observational and theoretical advances have reversed this scenario. In particular, astrophysical and cosmological observations of the smallest DM structures are now among the strongest constraints on the microphysical properties of DM.

This analysis improves upon previous work by using MW satellite observations over nearly the entire sky and rigorously accounting for both satellite detectability and uncertainties in the galaxy-halo connection. Our constraints are comparable in sensitivity to Lyman- α forest, strong lensing, and stellar stream perturbation analyses. Future cosmic surveys promise to further improve these

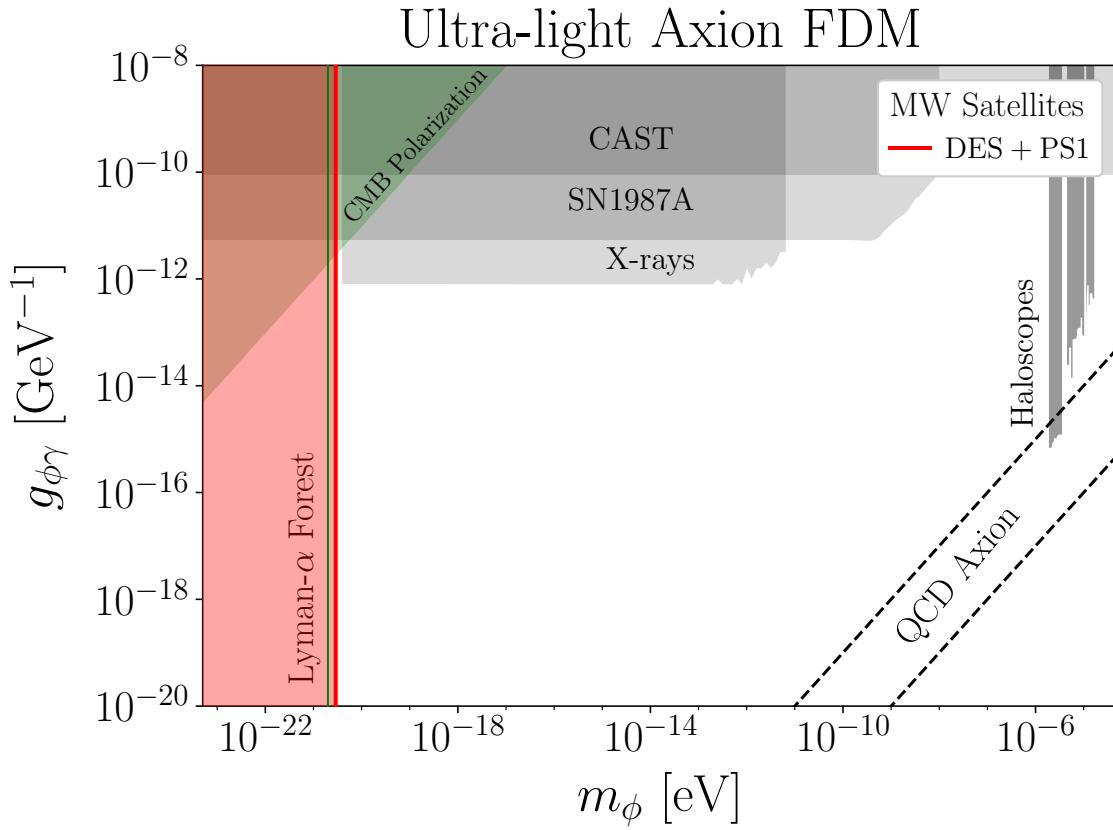


Figure 7.3: Constraints on ultralight axion particle mass versus axion-photon coupling from our analysis of the MW satellite population (red). Limits from CMB polarization washout [170] and the Lyman-\$\alpha\$ forest [233] are shown in green, and haloscope limits are shown as gray vertical bands. Experimental constraints from the CAST experiment [32], the lack of a \$\gamma\$-ray signal from SN1987A [357], and the x-ray transparency of the intracluster medium [382] are shown in gray and do not require that the ultralight axion makes up all of the DM. The dashed lines [386] span canonical QCD axion models [263, 140].

measurements and to enable a detailed comparisons to the internal dynamics of these galaxies [151, 438].

As the observational data improve, there are several uncertainties in the current modeling framework that are important to address. In particular, our use of only two realistic MW simulations limits the range of host halos and subhalo populations that enter our analysis; sampling a wider range of host halo masses, formation histories, and environments will improve the accuracy and precision of DM constraints derived from MW satellite galaxies. We describe other model uncertainties in Supplemental Material.

The breadth of DM models constrained by observations of MW satellites is particularly important given the growing interest in a wide range of theoretical possibilities following nondetections in collider, direct, and indirect searches for canonical WIMPs. In addition to the three DM paradigms considered in this work, small-scale structure measurements are also sensitive to the initial DM velocity distribution in nonthermal production scenarios [322], the DM formation epoch [407, 123], the DM self-interaction cross section [461, 493, 448, 334], and the DM particle lifetime [364, 469].

Future work could generalize our approach by measuring deviations in the small-scale linear matter power spectrum relative to a baseline CDM scenario rather than setting constraints in the context of particular DM models. Features in the power spectrum on extremely small scales are a hallmark of many inflationary models [252, 483], and it is conceivable that DM substructure measurements can be used to infer the nature of the corresponding primordial density fluctuations.

Chapter 8

Dark Matter Constraints. III. DM Formation Epoch

Abstract

A small fraction of thermalized dark radiation that transitions into cold dark matter (CDM) between big bang nucleosynthesis and matter-radiation equality can account for the entire dark matter relic density. Because of its transition from dark radiation, “late-forming dark matter” (LFDM) suppresses the growth of linear matter perturbations and imprints the oscillatory signatures of dark radiation perturbations on small scales. The cutoff scale in the linear matter power spectrum is set by the redshift z_T of the phase transition; tracers of small-scale structure can therefore be used to infer the LFDM formation epoch. Here, we use a forward model of the Milky Way (MW) satellite galaxy population to address the question: How late can dark matter form? For dark radiation with strong self-interactions, which arises in theories of neutrino-like LFDM, we report $z_T > 5.5 \times 10^6$ at 95% confidence based on the abundance of known MW satellite galaxies. This limit rigorously accounts for observational incompleteness corrections, marginalizes over uncertainties in the connection between dwarf galaxies and dark matter halos, and improves upon galaxy clustering and Lyman- α forest constraints by nearly an order of magnitude. We show that this limit can also be interpreted as a lower bound on z_T for LFDM that free-streams prior to its phase transition, although dedicated simulations will be needed to analyze this case in detail. Thus, dark matter created by a transition from dark radiation must form no later than one week after the big bang.

8.1 Paper Status and External Contributions

This chapter is published in slightly modified form in Physical Review D, Volume 103, Issue 4, p.043517-043526 with the title, “Constraints on the epoch of dark matter formation from Milky

Way satellites,” on which I am the corresponding author. It is the result of a collaboration with Subinoy, who developed the theory and linear Boltzmann solver for late-forming dark matter and made editorial contributions to the text.

8.2 Introduction

Despite intensive experimental searches in recent decades, the nature of dark matter (DM) remains a mystery. Combined with a cosmological constant (Λ), the simple hypothesis of a cold, collisionless dark matter (CDM) particle that interacts extremely weakly with Standard Model (SM) particles is consistent with all cosmological observations to date, on scales ranging from individual galaxies [397], to galaxy clusters [504], to the cosmological horizon as probed by large-scale structure [35] and cosmic microwave background (CMB) measurements [16, 59]. However, particle physics experiments have not detected canonical weakly interacting-massive-particle (WIMP) CDM, and several astrophysical anomalies have been claimed to provide evidence for physics beyond the collisionless CDM paradigm [97].

In this work, we explore and strongly constrain one such alternative scenario, known as “late-forming dark matter” (LFDM), where DM appears much later in cosmic history than WIMPs and other popular DM candidates [124, 199]. Instead of focusing on a specific particle physics construction of LFDM, we consider a general class of models in which DM is produced from an excess (dark) radiation component that undergoes a phase transition due to nontrivial interactions in the dark sector. Measurements from the *Planck* mission rule out the existence of a fully thermalized extra radiation component during the epoch of the CMB [16]. However, as we will demonstrate, LFDM can account for the entire DM content of the Universe while remaining compatible with *Planck* limits on the number of excess light degrees of freedom if even a tiny fraction of dark radiation transitions into CDM between the epoch of big bang nucleosynthesis (BBN) and the CMB.

LFDM is intriguing because it can be realized as a light, neutrino-like particle [124, 199], reviving the possibility of \sim eV-mass neutrino-like DM, which is incompatible with structure formation constraints if produced thermally [49, 6, 15, 332]. Intriguingly, there are tentative hints of a fourth sterile neutrino generation from short-baseline neutrino oscillation experiments [17, 18, 19]. However, this signal does not appear ubiquitously (e.g., [8]) and its interpretation as a sterile neutrino is difficult to reconcile with cosmological observables (e.g., [212]). Moreover, within the “3+1” neutrino oscillation framework, these results are difficult to reconcile with the absence of anomalies in ν_μ disappearance as probed by recent atmospheric [2, 1] and short-baseline [9, 2, 25] experiments. Thus, if the existence of a fourth sterile neutrino generation is confirmed by future analyses, it is likely that new physics beyond sterile-plus-active oscillation models is necessary to resolve the tension between neutrino appearance and disappearance data. Whether LFDM models can be connected to these anomalies is a compelling question for sterile neutrino model building, and is not the aim of this

paper. Instead, we focus on cosmological signatures of the LFDM phase transition.

The LFDM phase transition affects linear matter perturbations and imprints its effects on various tracers of the DM density field throughout cosmic history. In particular, the linear matter power spectrum $P(k)$ is suppressed on scales smaller than the size of the cosmological horizon at the LFDM transition redshift, z_T , because the corresponding modes entered the horizon while LFDM behaved like radiation. Thus, later phase transitions suppress power on larger scales. This phenomenology pertains to *any* cosmic fluid that transitions into CDM from a (dark) radiation component. Moreover, because the absence of cold, heavy DM particles always dilutes gravitational potentials, it also pertains to any scenario in which DM is absent until late times.

In this work, we leverage this power suppression signal to address the question: “What is the latest epoch after which dark matter must behave exactly like CDM?” We show that the answer depends on whether the LFDM fluid has strong self-interactions prior to its transition into CDM (we refer to this case as *self-interacting*, or SI), or whether it free-streams prior to the phase transition (we refer to this case as *free-streaming*, or FS). In the SI LFDM case, the linear matter power spectrum contains the oscillatory signatures of dark radiation perturbations, the amplitude of which depends on the strength of the LFDM self-interactions prior to the phase transition [124, 121]. These self-interactions are expected in neutrino-like LFDM models, including in theories of neutrino dark energy [124] and in a model of sterile fermion DM that has been proposed to have some observable effects on CMB [199]. Meanwhile, the limit in which LFDM transitions to CDM from a free-streaming dark radiation component without self-interactions yields a sharper cutoff in the matter power spectrum.

Analyses of the Lyman- α forest, galaxy clustering, and the high-redshift galaxy luminosity function have set a lower limit on the SI LFDM transition redshift of $z_{T,SI} \gtrsim 9 \times 10^5$ based on the lack of observed power spectrum suppression relative to CDM on quasilinear scales corresponding to wave numbers $k \sim 1h \text{ Mpc}^{-1}$ [407, 115]. Following the reasoning above, tracers of matter fluctuations on even smaller scales contain information about earlier LFDM transition redshifts. Indeed, LFDM initially gained popularity because of its ability to address several “small-scale crises” historically attributed to CDM, including the “missing satellites” [266, 325] and “too big to fail” [84, 186] problems for Milky Way (MW) satellite galaxies, which occupy DM halos that arise from fluctuations on nonlinear scales of $k \gtrsim 10h \text{ Mpc}^{-1}$.

State-of-the-art empirical models [241, 340, 264, 336, 333] and hydrodynamic simulations [408, 478, 188] combined with rigorous estimates for the incompleteness of current MW satellite searches provide strong evidence that the observed MW satellite population is consistent with CDM predictions. In Chapter 7 [332], we used the MW satellite model in Chapters 3–4 [336, 333]—which accurately describes the observed MW satellite population over nearly three-fourths of the sky, including satellites associated with the Large Magellanic Cloud—to derive constraints on a variety of non-CDM models that suppress the linear matter power spectrum on small scales. In particular,

[332] reported that the observed MW satellite population is consistent with CDM predictions down to a halo mass scale of $\sim 3 \times 10^8 M_\odot$, corresponding to characteristic wave numbers $k \sim 40h \text{ Mpc}^{-1}$, and ruled out thermal relic warm dark matter (WDM) lighter than 6.5 keV at 95% confidence. Importantly, this constraint is marginalized over uncertainties in the connection between faint galaxies and low-mass halos and the properties of the MW system. Independent studies of other small-scale structure probes, including the Lyman- α forest, strong gravitational lenses, and stellar streams, have derived consistent WDM constraints [459, 232, 226, 195, 41].

Here, we extend the analysis of Chapter 7 [332] to place limits on the LFDM formation epoch. We show that SI LFDM imprints a cutoff in the linear matter power spectrum that is very similar to thermal relic WDM, and we exploit this correspondence to constrain the model. Based on the abundance of MW satellite galaxies, our analysis yields a lower bound of $z_{T,SI} > 5.5 \times 10^6$ on the SI LFDM transition redshift at 95% confidence, which improves upon previous results [407, 115] by a factor of ~ 6 . This implies that SI LFDM must form no later than one week after the big bang. In addition, we show that our constraint on $z_{T,SI}$ can be interpreted as a lower limit on the FS LFDM transition redshift, and we estimate the improvement that future simulation-based analyses can provide for this model.

Throughout, we assume that LFDM constitutes the entire DM relic density, and we hold cosmological parameters fixed at the Λ CDM best-fit values from [10].

8.3 Late-forming Dark Matter Models

We begin with a brief overview of LFDM physics. We consider LFDM models in which an excess radiation component ΔN_{eff} undergoes a phase transition to a CDM state at redshift z_T . In this scenario, the initial number of relativistic degrees of freedom N_{eff} is generically larger than in a standard Λ CDM cosmology. However, we will see that even a tiny fractional increase in N_{eff} suffices to produce the observed CDM relic density, provided that the LFDM phase transition occurs a few e -foldings before matter-radiation equality (MRE).

Since the epoch of its phase transition to the present, LFDM redshifts identically to CDM, implying that

$$\rho_{\text{LFDM}}(z) = \rho_{\text{LFDM}}(z_T) \frac{(1+z)^3}{(1+z_T)^3}, \quad (8.1)$$

where $\rho_{\text{LFDM}}(z)$ is the LFDM density evaluated at redshift z . Assuming that a fraction of excess radiation is converted into the entire CDM density at redshift z_T , this yields the following decrement in the effective number of neutrino degrees of freedom:

$$\Delta N_{\text{eff}} \rho_\nu(z_T) = \rho_{\text{LFDM}}(0) (1+z_T)^3, \quad (8.2)$$

where $\rho_\nu(z_T)$ is the energy density of one neutrino-like radiation species at the formation epoch.

Thus, we have

$$\Delta N_{\text{eff}} = \frac{\rho_{\text{CDM}}(0)}{\rho_{\nu}(0)} \approx 0.2 \left(\frac{\Omega_{\text{CDM}} h^2}{0.1199} \right) \left(\frac{10^5}{1 + z_T} \right). \quad (8.3)$$

Note that ΔN_{eff} is inversely proportional to the redshift of the LFDM phase transition. Because the effective number of neutrino degrees of freedom changes dynamically in this model, observational constraints on N_{eff} must be interpreted with caution.

For most LFDM phase transition epochs between BBN and the CMB, the resulting value of ΔN_{eff} is smaller than the precision of current observational constraints on this quantity; for instance, Equation 8.3 implies that $z_T = 10^5$ corresponds to $\Delta N_{\text{eff}} = 0.2$, assuming the best-fit *Planck* value of $\Omega_{\text{CDM}} h^2 = 0.1199$ [16]. Recent constraints on N_{eff} from *Planck* and WMAP prefer the existence of a fractional dark radiation component, with $\Delta N_{\text{eff}} = 0.15$ at 95% confidence [15]. This bound is relaxed in the presence of nontrivial dark radiation self-interactions, which modify standard cosmological behavior during the radiation-dominated epoch [278]. Thus, LFDM is in complete agreement with ΔN_{eff} constraints if the phase transition occurs before $z \sim 10^5$, in which case $\Delta N_{\text{eff}} \ll 0.2$ is sufficient to account for the entire DM relic density. Such a small fractional change in ΔN_{eff} from an \sim eV neutrino-like particle also affects CMB density perturbations; in particular, modes with $\ell > 200$ that enter the horizon between BBN and the CMB respond to the presence of this tiny dark radiation excess. Constraints from this effect are compatible with the typical values of ΔN_{eff} required for LFDM to constitute the entirety of DM [417].

Importantly, unlike WIMPs (which couple to the SM through the weak interaction) or QCD axions (which primarily couple to the SM through electromagnetic interactions), LFDM need not have any interactions with the visible sector. Direct detection signatures are therefore not guaranteed for LFDM, although they are possible for specific constructions of the model. On the other hand, the suppression of the linear matter power spectrum, which manifests as a suppression of the power inferred from various tracers throughout cosmic history (e.g., [407, 115]), is *inevitable* in LFDM. In addition, dark acoustic oscillations (DAOs) imprinted prior to the phase transition can leave distinct signatures; for example, the 21-cm brightness power spectrum may be enhanced in LFDM models relative to CDM [122].

8.3.1 Self-interacting LFDM

SI LFDM is a natural model in which the phase transition from a dark radiation component to a CDM state can easily be achieved. Recently, it has been shown that \sim eV sterile neutrino-like dark fermions, which have strong self-interactions mediated by a sub-eV scalar field, can be trapped into DM “nuggets” in the radiation-dominated era, a few *e-foldings* before the CMB [199]. The phase transition occurs when the attractive scalar fifth force overcomes free-streaming, which traps all of the \sim eV fermions within a Compton volume into degenerate DM nuggets. Collectively, these nuggets behave exactly like CDM and are produced with negligible thermal velocities due to

their \sim TeV mass, unlike other LFDM models with non-negligible peculiar velocities that evolve ballistically after the phase transition [121]. The stability of the nuggets is achieved by fermion degeneracy pressure, which balances the scalar fifth force, and the duration of the phase transition is negligible compared to the Hubble time for any transition redshift prior to the epoch of the CMB. Because of the heavy, composite nature of the nuggets resulting from their nonlinear formation process, the initial distribution function of the thermal dark fermions is not conserved. Thus, the nuggets avoid the Tremaine-Gunn phase-space bound derived from the internal dynamics of dwarf galaxies that applies to other light fermionic dark matter and WDM candidates [447, 83, 23]. This model therefore provides a concrete construction of a phase transition in which a fluid that initially behaves like dark radiation changes its equation of state almost instantaneously at a transition redshift $z_{T,SI}$.

Bosonic SI LFDM appears in theories of neutrino dark energy, in which neutrinos interact with multiple scalar fields and behave like a single thermalized fluid [124]. In these theories, the scalars generally have hybrid potentials reminiscent of hybrid inflationary potentials. As the neutrino temperature dilutes near the epoch of MRE, one of the scalar fields that was stuck in a metastable minimum becomes tachyonic and begins to oscillate around a new minimum. The coherently oscillating field then behaves exactly like CDM, similar to the transition axion dark matter undergoes when the Hubble rate drops below its oscillation frequency.

From a theoretical perspective, the epoch of the LFDM phase transition in neutrino dark energy theories is expected to be very late, and is therefore subject to constraints arising from linear perturbation theory. In particular, the relevant range of LFDM formation epochs can be estimated by assuming that the coupling of the particle model is of $\mathcal{O}(1)$, which yields $1 \text{ eV} \lesssim T(z_{T,SI}) \lesssim 10^3 \text{ eV}$ for the temperature of the Universe at the phase transition [124]. The wave numbers corresponding to horizon entry for this range of transition epochs are $2 \times 10^{-2}h \text{ Mpc}^{-1} \lesssim k_{T,SI} \lesssim 20h \text{ Mpc}^{-1}$. We reiterate that this is an order-of-magnitude estimate that only assumes natural values of the coupling constants.

8.3.2 Free-streaming LFDM

In the FS LFDM model, a noninteracting dark radiation component that free-streams until the DM phase transition starts to oscillate coherently and behave like CDM at redshift $z_{T,FS}$. It is shown in [124] that a thermal field theory correction can in principle make this phase transition possible. In particular, consider a scalar field ϕ with mass m and a zero-temperature potential

$$V(\phi) = V_0 - \frac{m^2 \phi^2}{2} - \epsilon \phi^3 + \frac{\lambda \phi^4}{4}, \quad (8.4)$$

where V_0 is the zero-point energy and ϵ, λ are coupling constants. This potential can pick up a correction due to the presence of other fermionic fields at finite temperature, resulting in fluctuations

$$\delta V = DT^2\phi^2 \quad (8.5)$$

where D depends on the spin, coupling, and number of degrees of freedom of the other fields.

Here we have assumed that ϕ is not in thermal equilibrium with other fields, which implies that ϕ is noninteracting in a cosmological sense. With such a potential, the field is trapped in a minimum at $\phi = 0$ for $T \geq m/\sqrt{2D}$ [124]. After the Universe cools below this temperature, the field becomes tachyonic about the origin and settles into the true minimum, after which it coherently oscillates and behaves like CDM. This model is therefore a concrete example of FS LFDM.

8.4 Linear Perturbations

8.4.1 Free-streaming LFDM

Despite the variety of particle models described above, the initial conditions for LFDM matter perturbations after its phase transition are identical to that of a dark radiation component at the transition epoch. If the dark radiation component has no self-interactions, then matter perturbations can be treated exactly as in the case of neutrinos, and the evolution of FS LFDM density perturbations is obtained by solving a series of coupled differential equations [307]:

$$\begin{aligned} \dot{\delta} &= -\frac{4}{3}\theta - \frac{2}{3}\dot{h}, \\ \dot{\theta} &= k^2 \left(\frac{\delta}{4} - \sigma \right), \\ 2\dot{\sigma} &= \frac{8}{15}\theta - \frac{3}{15}kF_3 + \frac{4}{15}\dot{h} + \frac{8}{5}\dot{\eta}, \text{ and} \\ \dot{F}_\ell &= \frac{k}{2\ell+1} (\ell F_{\ell-1} - (\ell+1)F_{\ell+1}), \end{aligned} \quad (8.6)$$

where δ is the LFDM overdensity field, θ is its velocity divergence, h and η are metric perturbations in synchronous gauge, σ is the shear stress, F_ℓ is the ℓ th Legendre component of the momentum-averaged LFDM distribution function, k is the cosmological wave number, and overdots denote derivatives with respect to conformal time [307]. The solution for δ is an exponentially damped oscillator at subhorizon scales; physically, this represents the free-streaming of highly relativistic neutrinos.

To compute the growth of linear matter perturbations for the FS LFDM model, we modify the Boltzmann solver CAMB to evolve matter fluctuations up to a redshift $z_{T,FS}$ without CDM, and we extract the transfer function for neutrino perturbations at this redshift according to Equation

[8.6](#). We then use these neutrino (dark radiation) perturbations as initial conditions for LFDM density fluctuations at the epoch of its formation, and we evolve LFDM perturbations identically to CDM thereafter to obtain the linear matter power spectrum at later times. Thus, oscillations at small scales in the linear matter power spectrum arise because LFDM obtained its initial density fluctuations from neutrino-like perturbations at $z_{T,FS}$, which were damped and oscillatory at scales smaller than the size of the horizon at that time.

8.4.2 Self-interacting LFDM

Equation [8.6](#) provides the initial conditions for a neutrino-like particle that transitions to CDM. For SI LFDM, the situation is simplified because a strongly self-interacting neutrino-like fluid can be treated in the tight-coupling approximation, in which the anisotropic stress and higher-order terms are neglected (analogous to the treatment of the photon-baryon fluid). The following equations then describe linear perturbations for the SI LFDM model:

$$\begin{aligned}\dot{\delta} &= -\frac{4}{3}\theta - \frac{2}{3}\dot{h}, \\ \dot{\theta} &= k^2 \left(\frac{\delta}{4} - \sigma \right).\end{aligned}\tag{8.7}$$

We note that the above perturbation equations for a tightly coupled dark matter-radiation fluid are only valid until then epoch of the phase transition, and that—once LFDM forms—it behaves identically to cold, collisionless CDM. In our modified CAMB implementation, we therefore evolve matter perturbations until the redshift of the phase transition, $z_{T,SI}$, according to Equation [8.7](#). We then use the solution as the initial condition for subsequent evolution, which is identical to CDM.

8.5 Transfer Functions

To compare linear matter power spectra in our LFDM models to CDM, we compute the *transfer function*

$$T^2(k) \equiv \frac{P_{\text{LFDM}}(k)}{P_{\text{CDM}}(k)},\tag{8.8}$$

where $P_{\text{LFDM}}(k)$ [$P_{\text{CDM}}(k)$] is the LFDM (CDM) linear matter power spectrum evaluated at $z = 0$. The *half-mode scale* k_{hm} is defined as the wave number at which $T^2(k) = 0.25$.

Linear matter power spectra and transfer functions for our SI and FS LFDM models with $z_T = 1.5 \times 10^6$ ($k_T = 7h \text{ Mpc}^{-1}$) are shown in Figure [8.1](#). We note that the transition redshift shown in Figure [8.1](#) is marginally consistent with Lyman- α forest and galaxy clustering data [\[407\]](#); however, as we demonstrate below, it is robustly ruled out for both LFDM models by our MW satellite population analysis.

The right panel of Figure 8.1 illustrates three main features of LFDM transfer functions that are common to both of our model variants:

1. There is a cutoff in power relative to CDM at the comoving wave number k_T , which corresponds to the size of the horizon at the epoch of the LFDM phase transition. In particular, power is significantly suppressed on scales smaller than those corresponding to

$$k_T = \frac{aH_T}{c} \approx \frac{H_0\sqrt{\Omega_{\text{rad}}z_T}}{c}, \quad (8.9)$$

where H_T is the Hubble rate at the LFDM transition, $H_0 = 100h \text{ km s}^{-1} \text{ Mpc}^{-1}$ is the present-day Hubble rate, and $\Omega_{\text{rad}} \approx 10^{-4}$ is the energy density in radiation.¹

2. There are damped DAOs at scales smaller than those corresponding to k_T , resulting from dark radiation perturbations prior to the LFDM phase transition.
3. Cutoffs in the transfer functions for both model variants exhibit *k-translation invariance*. Specifically, given two SI or FS LFDM models with transition redshifts $z_{T,1}$ and $z_{T,2}$ and transfer functions $T_1^2(k)$ and $T_2^2(k)$, we have

$$T_2^2(k) = T_1^2\left(\frac{z_{T,2}}{z_{T,1}}k\right) \quad (8.10)$$

along the initial cutoff. This symmetry follows from the linear relation between k_T and z_T in Equation 8.9 and from the scale invariance of Hubble expansion in the radiation-dominated epoch. We emphasize that Equation 8.10 only holds along the initial power spectrum cutoff; this is sufficient for our purposes because DAOs occur at extremely small scales for the typical transition redshift values we consider. Equation 8.10 is useful because it allows us to analytically compute LFDM transfer functions as a continuous function of z_T using the power spectra that were computed with CAMB for discrete transition redshifts.

8.5.1 Self-interacting LFDM

The SI LFDM transfer function exhibits a smooth cutoff that is remarkably similar to that in thermal relic WDM until the onset of DAOs. The tight correspondence between the cutoff in these transfer functions is reminiscent of the mapping between thermal relic WDM and velocity-independent DM-proton scattering found in Chapter 6 [335], and (to a lesser extent) a similar mapping identified for models with DM-radiation interactions [72, 167]. Despite different dark matter microphysics, the transfer function for our SI LFDM model is also similar to that for self-interacting dark matter models in which massive dark photon mediators decay to dark fermions [229]. More generally, [74, 71]

¹As discussed above, CMB constraints on ΔN_{eff} set a limit of $z_T \gtrsim 4 \times 10^5$. Later transitions also result in severe suppression of the matter power spectrum on quasilinear scales according to Equation 8.9.

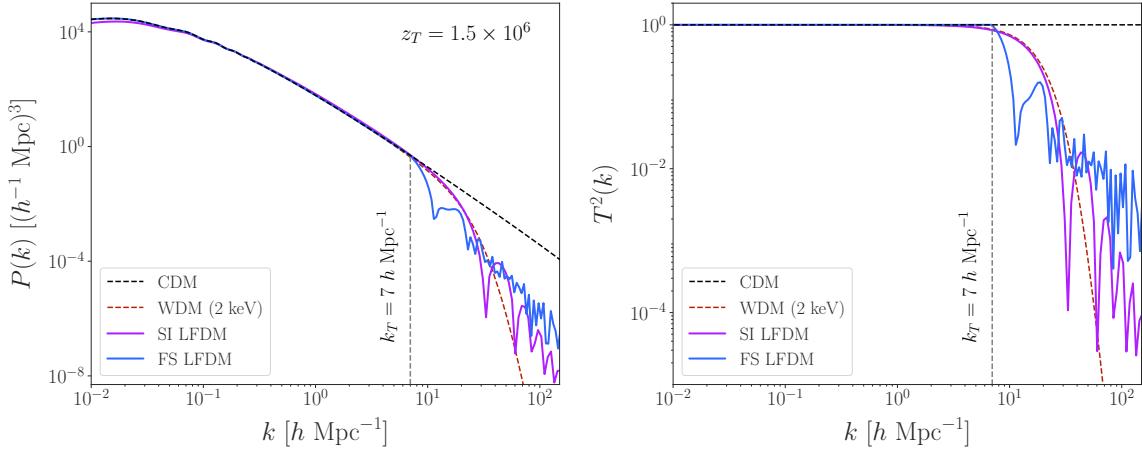


Figure 8.1: Linear matter power spectra (left) and transfer functions (right) for self-interacting (magenta) and free-streaming (cyan) late-forming dark matter models, compared to cold dark matter (dashed black) and thermal relic warm dark matter (dashed red). Both LFDM models are shown with a transition redshift of $z_T = 1.5 \times 10^6$, corresponding to a comoving wave number of $k_T = 7 h \text{ Mpc}^{-1}$. LFDM power spectra are suppressed relative to CDM at wave numbers greater than k_T , and they exhibit dark acoustic oscillations on even smaller scales, beginning at $\sim 6k_T$ ($\sim 2k_T$) for SI (FS) LFDM. The cutoff in the SI LFDM power spectrum is very similar to that in WDM, until the onset of DAOs.

have shown that interacting DM models often impact the linear matter power spectrum such that they are effectively “warm.” The existence of the mapping between SI LFDM and thermal relic WDM is therefore not surprising given its strong self-interactions prior to the phase transition.

To make this correspondence quantitative, we construct a relation between the SI LFDM and thermal relic WDM models following a half-mode scale matching procedure similar to [167, 335]. In particular, we derive the following relation from our CAMB output:

$$k_{\text{hm,SI}} \approx 2.8k_{T,\text{SI}} \approx 1.3 \left(\frac{z_{T,\text{SI}}}{10^5} \right) h \text{ Mpc}^{-1}. \quad (8.11)$$

Meanwhile, the half-mode scale in WDM is given by [460]

$$k_{\text{hm,WDM}} = \frac{2\pi}{\lambda_{\text{hm,WDM}}} = 9.2 \left(\frac{m_{\text{WDM}}}{1 \text{ keV}} \right)^{1.11} \left(\frac{\Omega_m}{0.25} \right)^{-0.11} \left(\frac{h}{0.7} \right)^{-1.22} h \text{ Mpc}^{-1}, \quad (8.12)$$

where m_{WDM} is the thermal relic WDM mass. Solving for the transition redshift that causes the half-mode scales of the WDM and SI LFDM transfer functions to match yields the relation

$$z_{T,\text{SI}} \approx 7 \times 10^5 \left(\frac{m_{\text{WDM}}}{1 \text{ keV}} \right)^{1.11} \left(\frac{\Omega_m}{0.25} \right)^{-0.11} \left(\frac{h}{0.7} \right)^{-1.22}. \quad (8.13)$$

We find that LFDM and WDM transfer functions matched in this way agree to better than $\sim 5\%$ along the initial cutoff over the entire SI LFDM parameter space of interest.

Examples of SI LFDM transfer functions along with matched WDM transfer functions are shown in the left panel of Figure 8.2. On this plot, we indicate the comoving wave number corresponding to the *minimum halo mass*, i.e., the lowest-mass halo inferred to host MW satellite galaxies. In particular, from an analysis of the MW satellite population using DES and PS1 data over nearly three-fourths of the sky, [333] found that the lowest peak virial halo mass corresponding to observed MW satellite galaxies is less than $\mathcal{M}_{\min} = 3.2 \times 10^8 \text{ M}_\odot$ at 95% confidence, corresponding to a comoving wave number of $k_{\text{crit}} \approx 36h \text{ Mpc}^{-1}$. We also indicate the WDM transfer function ruled out by these observations of the MW satellite population at 95% confidence, corresponding to a 6.5 keV thermal relic [332].

8.5.2 Free-streaming LFDM

The power spectrum cutoff in FS LFDM is significantly sharper than in SI LFDM, as expected due to its free-streaming behavior prior to the phase transition. Thus, it is difficult to directly map FS LFDM to WDM, which forces us to take a more conservative approach in order to derive constraints.

Nonetheless, we can still construct a relation between the half-mode scale and the transition redshift for FS LFDM based on our CAMB output. This yields

$$k_{\text{hm,FS}} \approx 1.4k_{T,\text{FS}} \approx 0.65 \left(\frac{z_{T,\text{FS}}}{10^5} \right) h \text{ Mpc}^{-1}. \quad (8.14)$$

For a fixed transition redshift, $k_{\text{hm,FS}} < k_{\text{hm,SI}}$, which makes sense given the sharper power spectrum cutoff in FS LFDM relative to SI LFDM. FS LFDM transfer functions are shown in the right panel of Figure 8.2.

8.6 Constraints from Milky Way Satellites

We use the relations derived above to translate thermal relic WDM limits from the MW satellite population into LFDM constraints. Given that halos with masses lower than $3.2 \times 10^8 \text{ M}_\odot$ are required to host currently observed MW satellite galaxies [333], there must be enough power to form bound DM halos on the corresponding comoving scales—i.e., down to a critical wave number of

$$k_{\text{crit}} = \frac{2\pi}{\lambda_{\min}} = \pi \left(\frac{4\pi\rho_m}{3\mathcal{M}_{\min}} \right)^{1/3} \approx 36h \text{ Mpc}^{-1}, \quad (8.15)$$

where ρ_m is the LFDM density today, \mathcal{M}_{\min} is the minimum halo mass, and λ_{\min} is the corresponding length scale in linear theory. Halos at this mass scale need not merely exist, but must be formed in enough *abundance* to match the observed MW satellite population population. Thus, we will

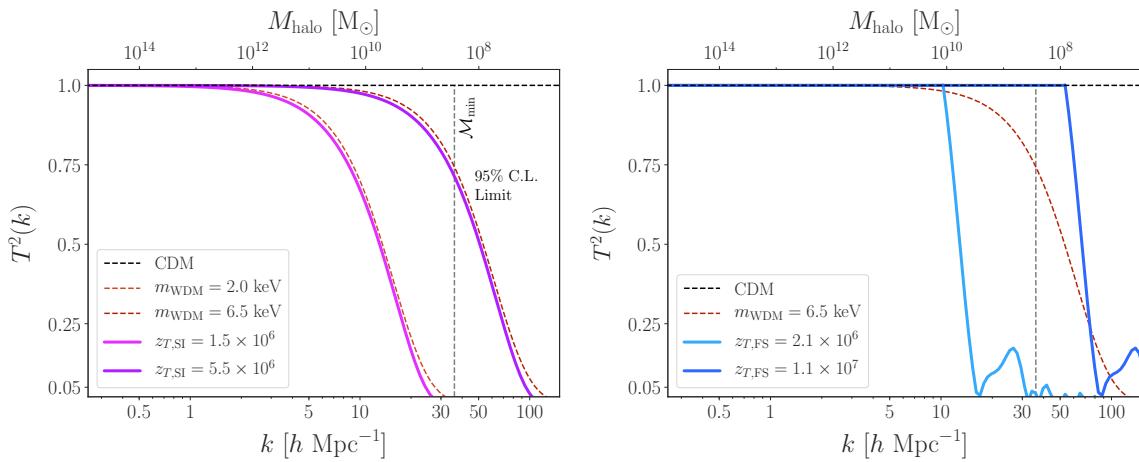


Figure 8.2: Transfer functions for self-interacting (left) and free-streaming (right) late-forming dark matter models, compared to cold dark matter (dashed black) and thermal relic warm dark matter (dashed red). SI LFDM models are shown for a range of transition redshifts, with the highest transition redshift corresponding to the SI LFDM model that is ruled out by the abundance of Milky Way satellites at 95% confidence: $z_{T,SI} > 5.5 \times 10^6$. The light-blue FS LFDM model corresponds to the transition redshift that is conservatively ruled out by our analysis: $z_{T,FS} > 2.1 \times 10^6$. Vertical dashed lines show the comoving scale that approximately corresponds to the mass of the smallest halo inferred to host observed MW satellite galaxies, $3.2 \times 10^8 M_\odot$ [333]. In the left panel, WDM transfer functions are slightly shifted horizontally for visual clarity.

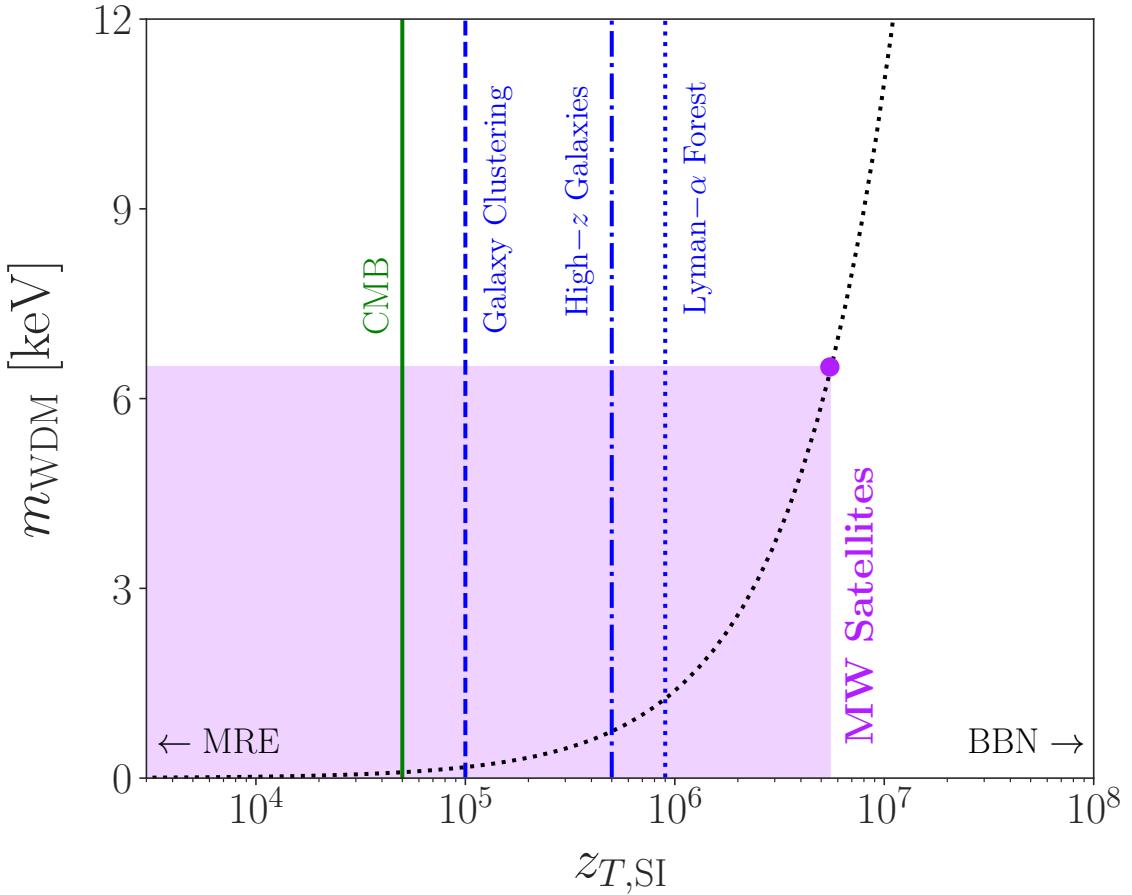


Figure 8.3: Constraints on the transition redshift for self-interacting late-forming dark matter, versus the corresponding thermal relic warm dark matter mass based on the half-mode mass relation in Equation 8.13. Our Milky Way satellite constraint on $z_{T,SI}$ and the lower limit on the thermal relic WDM mass of 6.5 keV from which we derive this limit [332] are shown by the shaded purple region. Limits on the SI LFDM transition redshift from the cosmic microwave background (green), Sloan Digital Sky Survey galaxy clustering (dashed blue [407]), the high-redshift galaxy luminosity function (dot-dashed blue [115]) and the Lyman- α forest (dotted blue [407]) are shown as vertical lines. Vertical lines indicate constraints derived specifically for LFDM, and do not indicate other recent WDM constraints from small-scale structure probes. LFDM must transition to CDM between matter-radiation equality ($z \approx 3 \times 10^3$) and big bang nucleosynthesis ($z \approx 10^{10}$), which are schematically indicated by arrows.

obtain a lower limit on the transition redshift in both LFDM models based on the lower limit on the thermal relic WDM mass.

8.6.1 Self-interacting LFDM

The LFDM-WDM mapping constructed above allows us to translate thermal relic WDM limits derived from the MW satellite population into LFDM constraints. High-resolution cosmological simulations have been performed in order to predict the WDM subhalo mass function in MW-mass halos [415, 299, 28, 80], and these have been used in conjunction with the observed MW satellite population to place stringent constraints on thermal relic WDM. Chapter 7 [332] reports $m_{\text{WDM}} > 6.5 \text{ keV}$ at 95% confidence, which we directly translate into a constraint on SI LFDM via Equation 8.13, yielding $z_{T,\text{SI}} > 5.5 \times 10^6$, also at 95% confidence. This limit implies that the dark radiation which transitions to LFDM causes $\Delta N_{\text{eff}} \lesssim 4 \times 10^{-3}$, assuming that LFDM constitutes the entire DM relic density [Equation 8.3]. Exploring the generality of this indirect constraint on ΔN_{eff} from small-scale structure measurements is a compelling avenue for future work.

Figure 8.3 compares this limit to constraints on $z_{T,\text{SI}}$ derived from the CMB (resulting from N_{eff} constraints), low-redshift galaxy clustering from the Sloan Digital Sky Survey [407], the high-redshift galaxy luminosity function [115], and the Lyman- α forest [407]. Our limit improves upon the Lyman- α forest result by a factor of ~ 6 , which can be understood in terms of the comoving scales probed by the MW satellite population. Specifically, the lowest-mass halo inferred to host an observed satellite is $\sim 3 \times 10^8 M_{\odot}$ [333], which roughly corresponds to a wave number of $k \sim 40h \text{ Mpc}^{-1}$, while the Lyman- α forest data used in [407] reaches $k \sim 5h \text{ Mpc}^{-1}$. We expect $z_{T,\text{SI}}$ to scale linearly with the wave number corresponding to the smallest scale probed in an observational analysis, and the improvement we observe relative to this Lyman- α constraint is consistent with this expectation.² Other small-scale probes that achieve comparable sensitivity to thermal relic WDM, including strong gravitational lensing [195, 226] and stellar streams [41], will yield similar LFDM constraints.

Our SI LFDM limit relies on an analytic mapping to thermal relic WDM and is therefore not directly validated using LFDM simulations. We note that [115] ran simulations of these models with similar half-mode scales and found that the high-redshift ($z > 4$) LFDM halo mass function is comparable to that in WDM. Those findings are further consistent with the suite of LFDM simulations from [13], which show that oscillatory features in the linear matter power spectrum are erased in the $z = 0$ halo mass function. Meanwhile, [75]—working in the Effective Theory of Structure Formation (ETHOS) framework [119]—found the peak heights of interest for our SI LFDM constraints lead to negligible differences in the high-redshift halo mass function relative to thermal relic WDM. Finally, [229] showed that the halo mass function for self-interacting dark matter models with similar transfer functions to our SI LFDM model are nearly indistinguishable from matched

²More recent Lyman- α forest analyses (e.g., [395, 394]) probe smaller scales and a wider range of redshifts, and will therefore improve upon the LFDM constraints in [407].

WDM models, and used this correspondence along with a conservative treatment of the subhalo population inferred from MW satellites to place constraints similar in spirit to ours. All of these results lend confidence to the robustness of our result when framed as a conservative limit.

8.6.2 Free-streaming LFDM

The right panel of Figure 8.2 demonstrates the reason that it would be dangerous to set a constraint on FS LFDM based on matching its half-mode scale to WDM. In particular, because the FS LFDM power spectrum cutoff is much steeper than in thermal relic WDM, the half mode-matched model is significantly *less* suppressed than the corresponding WDM model along the initial power spectrum cutoff. Thus, we bracket the range of allowed FS LFDM transition redshifts as follows:

1. We place a *fiducial* lower limit on $z_{T,FS}$ by finding the FS LFDM transfer function that yields strictly greater power suppression than the ruled-out thermal relic WDM model for all wave numbers $k > 10h \text{ Mpc}^{-1}$, roughly corresponding to halo masses below 10^{10} M_\odot .³ Below this wave number, small differences between the FS LFDM and WDM transfer functions are negligible for the FS LFDM models of interest. This yields a conservative limit of $z_{T,FS} > 2.1 \times 10^6$ and is shown by the light-blue transfer function in Figure 8.2.
2. We forecast an *optimistic* limit on $z_{T,FS}$ by matching it to the half-mode scale of the thermal relic WDM model that is ruled out at 95% confidence by the MW satellite population. This yields $z_{T,FS} > 1.1 \times 10^7$ and is shown by the dark-blue transfer function in Figure 8.2. This constraint is optimistic because the abundance of subhalos that host MW satellites are sensitive to a convolution of power on (nonlinear) scales, rather than a single mode at which the power spectrum is suppressed by a characteristic amount (e.g., k_{hm}); thus, transfer functions with different cutoff shapes cannot be matched in detail.

Because the FS LFDM model has not previously been considered in the context of small-scale structure measurements, we do not have a direct point of comparison for our constraints on its transition redshift. However, our *fiducial* FS LFDM is extremely conservative. It is therefore clear that $z_{T,FS}$ must be of the same order-of-magnitude $z_{T,SI}$, which is physically reasonable.

Like our SI LFDM constraint, our forecasted *optimistic* limit on $z_{T,FS}$ is analytic and therefore must be confirmed with measurements of the subhalo mass function in dedicated LFDM simulations of MW-like systems. This situation is reminiscent of that for fuzzy dark matter (FDM), which also features steeper power suppression (for a fixed half-mode scale) than thermal relic WDM. Half-mode matching predicts a stringent limit on the FDM mass (e.g., Chapter 6 [335]); however, constraints based directly on the FDM subhalo mass function are less strict [418, 332]. We are therefore confident that the correct limit on $z_{T,FS}$ lies between our *fiducial* and *optimistic* constraints.

³This procedure is similar to that used to constrain resonantly produced sterile neutrinos in [414, 332] and developed by [308] to constrain velocity-dependent DM-proton interactions.

8.7 Conclusion

In this study, we set novel constraints on the dark matter formation epoch using state-of-the-art limits on the suppression of the small-scale matter power spectrum from the Milky Way satellite population. Specifically, we focused on the theoretically motivated paradigm of late-forming dark matter, which transitions to collisionless, cold dark matter from a dark radiation state. We showed that the epoch of the LFDM transition determines the cutoff scale in the linear matter power spectrum, which is processed into a suppression of power throughout cosmic history. By exploiting the correspondence between the power spectrum cutoff in a LFDM model with strong self-interactions prior to the phase transition versus that in thermal relic warm dark matter, we used the latest WDM constraint from the MW satellite population to place a stringent lower limit on the LFDM transition redshift. This constraint improves upon previous results by nearly an order of magnitude. We also estimated lower limits on the transition redshift for free-streaming LFDM.

Crucially, several independent tracers of small-scale structure corroborate the dark matter constraints set by recent MW satellite studies; thus, our constraints are not highly dependent on the particular probe used to set the WDM limit we exploited in this paper. In particular, analyses of the Lyman- α forest flux power spectrum [459, 232], strongly lensed quasar flux ratio anomalies and magnifications [226, 195], and perturbations in Galactic stellar streams [41] have achieved similar sensitivity to thermal relic WDM relative to the MW satellite population, even though the observational and theoretical systematics of these probes differ. Thus, these other small-scale structure probes can also be used to constrain the dark matter transition redshift. This is particularly important because the dark acoustic oscillations imprinted prior to the LFDM phase transition can potentially have distinct consequences for different tracers of the matter power spectrum at various epochs (e.g., [122]).

Extending the sensitivity of dark matter formation epoch measurements to even earlier times requires probing the linear matter power spectrum on extremely small scales. For example, ruling out the possibility that LFDM forms after BBN requires sensitivity to linear modes with $k \sim 10^5 h \text{ Mpc}^{-1}$, or halos with masses of $\sim 10^{-2} M_\odot$. These tiny, baryon-free halos are only detectable through their gravitational effects, which next-generation pulsar timing arrays [375] and gravitational wave lensing measurements [348] can potentially discover.

Chapter 9

Dark Matter Self-interactions in Milky Way-mass Halos

Abstract

We explore the impact of elastic, anisotropic, velocity-dependent dark matter (DM) self-interactions on the host halo and subhalos of Milky Way (MW)-mass systems. We consider a generic self-interacting dark matter (SIDM) model parameterized by the masses of a light mediator and the DM particle. The ratio of these masses, w , sets the velocity scale above which momentum transfer due to DM self-interactions becomes inefficient. We perform high-resolution zoom-in simulations of an MW-mass halo for values of w that span scenarios in which self-interactions either between the host and its subhalos or only within subhalos efficiently transfer momentum, and we study the effects of self-interactions on the host halo and on the abundance, radial distribution, orbital dynamics, and density profiles of subhalos in each case. The abundance and properties of surviving subhalos are consistent with being determined primarily by subhalo–host halo interactions. In particular, subhalos on radial orbits in models with larger values of the cross section at the host halo velocity scale are more susceptible to tidal disruption owing to mass loss from ram pressure stripping caused by self-interactions with the host. This mechanism suppresses the abundance of surviving subhalos relative to collisionless DM simulations, with stronger suppression for larger values of w . Thus, probes of subhalo abundance around MW-mass hosts can be used to place upper limits on the self-interaction cross section at velocity scales of $\sim 200 \text{ km s}^{-1}$, and combining these measurements with the orbital properties and internal dynamics of subhalos may break degeneracies among velocity-dependent SIDM models.

9.1 Paper Status and External Contributions

This chapter is published in slightly modified form in the *Astrophysical Journal*, Volume 896, Issue 2, p.112-128 with the title, “Signatures of Velocity-dependent Dark Matter Self-interactions in Milky Way-mass Halos,” on which I am the corresponding author. It is the result of a collaboration with Arka Banerjee, Susmita Adhikari, Yao-Yuan Mao, and Risa Wechsler. Arka and Susmita led the development of the self-interacting dark matter cosmological simulation algorithm and all of us worked together to run and analyze the simulations presented below. In addition, all coauthors made editorial contributions to the text. This work was stimulated by workshops hosted by the Dark Energy Science Collaboration’s Dark Matter Working Group, and I am grateful to all of the members of this group for their support and insights.

9.2 Introduction

Self-interacting dark matter (SIDM) has long been an attractive alternative to purely cold, collisionless dark matter (CDM) owing to several potential “small-scale problems” attributed to CDM (see [97] for a review of these problems, and see [448] for a review of their potential resolutions in SIDM). Historically, SIDM was motivated by the core–cusp problem, which concerns the apparent discrepancy between the steep, cuspy NFW profiles ubiquitous among DM halos in CDM simulations and the flatter, cored profiles inferred from the dynamics of various tracers in dwarf galaxies [174, 432, 125, 113]. The core–cusp problem is sensitive to the impact of baryonic feedback, and it remains an active area of study (e.g., [157]). In particular, it has recently been cast in terms of the diversity of inner DM density profiles for field [350, 251, 254, 380, 495, 406] or satellite [450, 247] galaxies at fixed halo properties. For Milky Way (MW) satellite galaxies, self-interactions in the presence of the tidal field of the Galactic disk may accelerate gravothermal core collapse, and this process has been proposed as a unique signature of SIDM ([256, 345, 403]).

Self-interactions can also affect the abundance of DM substructure in a statistical fashion. In particular, many authors have studied the subhalo populations of MW-mass hosts in the context of SIDM (e.g., [125, 113, 142, 461, 392, 493, 147, 462, 391]), and a few cosmological SIDM-plus-hydrodynamic simulations of isolated dwarfs [182, 215, 177] and MW-mass systems [390] have been performed.

The consensus of these studies has been that SIDM models that do not drastically reduce the number counts of surviving subhalos have little effect on subhalo abundances in MW-mass systems (see, e.g., the discussion in [448]). However, this is a model-dependent statement that warrants scrutiny in light of increasingly precise constraints on the abundance of subhalos in MW-mass systems enabled by recent measurements of the luminosity function of satellite galaxies around the MW (e.g., [149]) and around MW analogs (e.g., [193, 270, 55]). In addition, theoretical predictions for gaps and perturbations in nearby stellar streams (e.g., [41, 76]), direct and indirect detection

experiments (e.g., [230, 235]), and flux ratios in strongly lensed systems (e.g., [195, 226]) are all sensitive to the abundance and internal properties of DM substructure in distinct ways.

Many of the aforementioned SIDM studies only consider isotropic, velocity-independent self-interactions, at least on the velocity scales relevant for the MW host halo and its subhalos. Nonetheless, there are theoretical and observational arguments in favor of a velocity-*dependent* SIDM cross section, which is a generic consequence of interactions in a nonminimal dark sector (see, e.g., [171, 294, 449, 66, 255, 448]). While velocity-dependent SIDM models have been explored on galaxy cluster scales (e.g., [389, 39]), a thorough study of the corresponding predictions on Galactic scales is now crucial. This is particularly relevant because the increasingly well-measured abundance and internal properties of MW satellites can be affected by self-interactions at both the host halo velocity scale (which depend on satellites' typical velocities relative to the host) and subhalo velocity scales (which depend on the internal velocity dispersion of individual subhalos).

Herein we investigate the effects of an SIDM model with a generic, velocity-dependent self-interaction cross section on the DM host halo and subhalos of MW-mass systems. We focus on the physical mechanisms that shape the abundance and properties of surviving and disrupted subhalos, setting the stage for future analyses to constrain the SIDM cross section at the velocity scale of MW-mass systems. This paper is organized as follows: We describe our SIDM model in Section 9.3 and our simulations in Section 9.4. We then present our results, focusing on the properties of our host halos in Section 9.5 and their subhalo populations in Section 9.6. We discuss challenges and caveats in Section 9.7, we compare to previous studies in Section 9.8, we consider prospects for SIDM constraints in Section 9.9, and we conclude in Section 9.10.

9.3 SIDM Model

We consider an SIDM model in which a DM particle χ interacts under the exchange of a light mediator, which can be either a scalar particle ϕ or a vector particle ϕ^μ . These scenarios are respectively described by the interaction Lagrangians [448]

$$\mathcal{L}_{\text{int}} = \begin{cases} g_\chi \bar{\chi} \chi \phi, & \text{scalar mediator} \\ g_\chi \bar{\chi} \gamma^\mu \chi \phi_\mu, & \text{vector mediator,} \end{cases} \quad (9.1)$$

where g_χ is a coupling constant and γ^μ are Dirac matrices.

Self-interactions governed by Equation 9.1 can be described by a Yukawa potential

$$V(r) = \pm \frac{\alpha_\chi}{r} e^{-m_\phi r}, \quad (9.2)$$

where r is the separation between DM particles, $\alpha_\chi \equiv g_\chi^2/4\pi$ is the analog of the fine-structure constant in the dark sector, and m_ϕ is the mediator mass.

We focus on t -channel scattering, which leads to an effective differential scattering cross section [231, 281]

$$\frac{d\sigma}{d\Omega} = \frac{\sigma_0}{2 \left[1 + \frac{v^2}{w^2} \sin^2 \left(\frac{\theta}{2} \right) \right]^2}, \quad (9.3)$$

where v is the relative velocity between interacting DM particles with mass m_χ , $w \equiv m_\phi/m_\chi$ is a characteristic velocity scale, $\sigma_0 \equiv 4\pi\alpha_\chi^2 m_\chi^2/m_\phi^4$ is the amplitude of the cross section, and θ is the scattering angle in the center-of-mass frame. For $w \ll v$ (in natural units), typically corresponding to an MeV-scale mediator for a GeV-scale DM particle mass, Equation 9.3 reduces to Rutherford-like scattering; for a heavy mediator, it reduces to velocity-independent isotropic scattering.

The corresponding momentum transfer cross section for identical particles is given by [248]

$$\sigma_T = \int \frac{d\sigma}{d\Omega} (1 - |\cos \theta|) d\Omega. \quad (9.4)$$

Note that σ_T is a function of v , w , and σ_0 . Finally, the *total* cross section

$$\sigma = \int \frac{d\sigma}{d\Omega} d\Omega \quad (9.5)$$

determines the probability of DM self-interactions in our numerical implementation (see Section 9.4.2). For the case of isotropic, velocity-independent self-interactions, the total cross section is related to σ_T and σ_0 via $\sigma = 2\sigma_T = 2\pi\sigma_0$.

At velocity scales $v > w$, the momentum transfer cross section falls off as v^{-4} ; meanwhile, for $v \ll w$, it flattens toward its asymptotic value. Thus, for values of w that are small relative to the typical velocities in a virialized system, interactions at low relative velocities are more effective at transferring momentum than interactions at high relative velocities. In the context of MW-mass systems, this implies that interactions among host halo particles and between host halo and subhalo particles are less effective than interactions among subhalo particles if w is small relative to the characteristic velocity scale of the host (~ 200 km s $^{-1}$). On the other hand, as w increases toward the velocity scale of the host halo, subhalo–host halo and host halo–host halo interactions become more significant.

The characteristic velocity scale w is an easily interpretable quantity that captures much of the key physics that shapes MW-mass systems in our two-parameter SIDM model. Thus, we parameterize and refer to our simulations in terms of w , and we choose σ_0/m_χ for each model so that the momentum transfer cross section at the velocity scales of interest yields enough self-interactions to produce observable effects, but not so many that the models are likely ruled out already. In particular, we study the four following SIDM model variants:

1. w_{10} : a model with $w = 10$ km s $^{-1}$, $\sigma_0/m_\chi = 8/\pi$ cm 2 g $^{-1}$, for which self-interactions within subhalos are significant, but self-interactions at the host halo velocity scale are negligible;

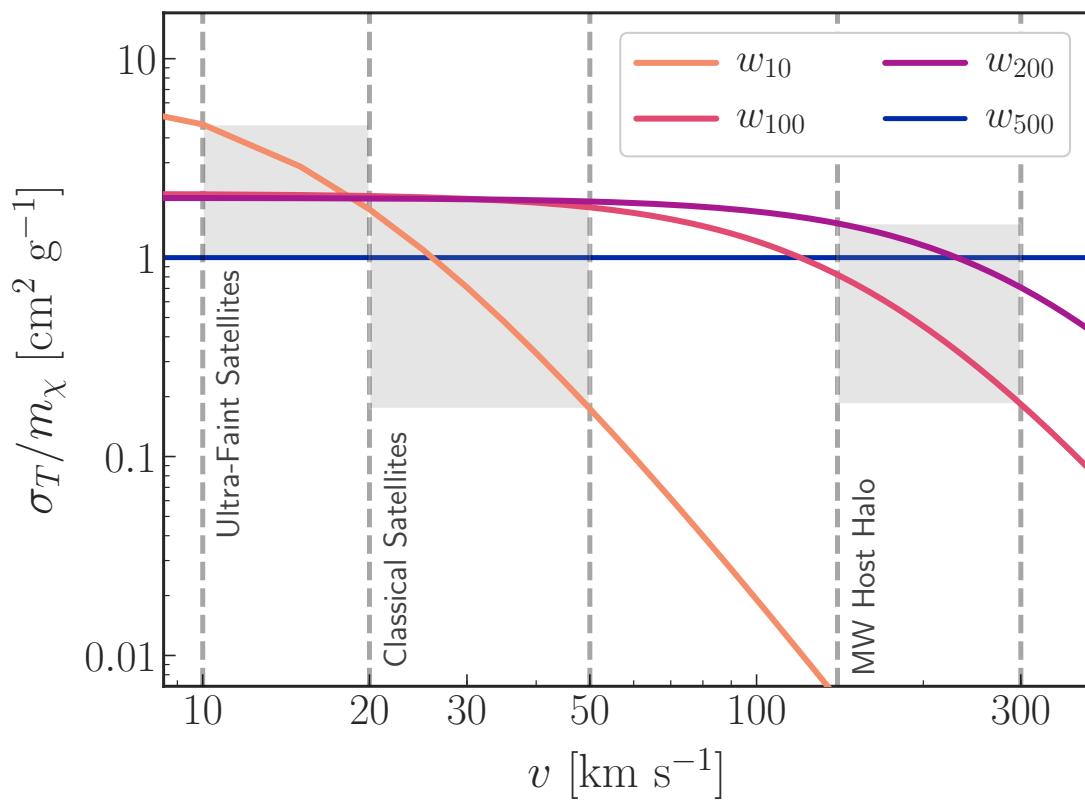


Figure 9.1: Momentum transfer cross sections for our SIDM model variants as a function of relative scattering velocity. Each model variant is labeled by w , the velocity scale above which the SIDM cross section falls off as v^{-4} . The velocity scale relevant for interactions among host halo particles and between host halo and subhalo particles is indicated by the “MW Host Halo” band. Shaded bands indicate characteristic velocities for DM particles within the subhalos expected to host classical and ultrafaint MW satellite galaxies.

2. w_{100} : a model with $w = 100 \text{ km s}^{-1}$, $\sigma_0/m_\chi = 2/\pi \text{ cm}^2 \text{ g}^{-1}$, for which self-interactions within low-mass subhalos are less significant than in w_{10} , but self-interactions at the host halo velocity scale are more significant;
3. w_{200} : a model with $w = 200 \text{ km s}^{-1}$, $\sigma_0/m_\chi = 2/\pi \text{ cm}^2 \text{ g}^{-1}$, for which self-interactions within low-mass subhalos are identical to w_{100} , but self-interactions at the host halo velocity scale are more significant;
4. w_{500} : a model with $w = 500 \text{ km s}^{-1}$, $\sigma_0/m_\chi = 1/\pi \text{ cm}^2 \text{ g}^{-1}$, for which self-interactions within low-mass subhalos are the least significant among our model variants, while interactions at the host halo velocity scale are similar to w_{200} (though slightly more effective at high velocities). Scattering in this model is isotropic and velocity independent on the scales relevant for MW-mass systems, meaning that self-interactions with large scattering angles at high relative velocities are potentially significant.

Figure 9.1 shows the momentum transfer cross section corresponding to each model variant, and their main properties are summarized in Table 9.1.

9.4 Simulations

9.4.1 General Description

For each SIDM model variant, we resimulate the same MW-mass halo using fixed initial conditions. This host halo is chosen from the suite of CDM-only MW-mass zoom-in simulations presented in [315]. The highest-resolution particles in our zoom-in simulations have a mass of $3 \times 10^5 M_\odot h^{-1}$, and the softening length in the highest-resolution regions is $170 \text{ pc } h^{-1}$. Halo catalogs and merger trees were generated using the ROCKSTAR halo finder and the CONSISTENT-TREES merger tree code [47, 48]. Throughout we define virial quantities using the [93] critical overdensity $\Delta_{\text{vir}} \simeq 99.2$, as appropriate for the cosmological parameters in our simulations: $h = 0.7$, $\Omega_m = 0.286$, $\Omega_b = 0.047$, and $\Omega_\Lambda = 0.714$.

9.4.2 SIDM Implementation

To implement DM self-interactions in our zoom-in simulations, we follow the prescription in [39]. Briefly, self-interactions are implemented using a modified version of the GADGET-2 N -body code that allows DM particles to transfer momentum and energy with an interaction probability set by the total SIDM cross section in Equation 9.5 and a scattering angle drawn from the distribution given by Equation 9.3. We refer the reader to [39] for the details of our SIDM implementation.

Simulation	w	σ_0/m_χ	Cored Host?	Cored Subs?	RPS?	$N_{\text{SIDM}}/N_{\text{CDM}}$
CDM	X	X	X	1.0
w_{10}	10	$8/\pi$	X	✓	X	0.82
w_{100}	100	$2/\pi$	✓	✓	✓	0.64
w_{200}	200	$2/\pi$	✓	✓	✓	0.65
w_{500}	500	$1/\pi$	✓	✓	✓	0.44

Table 9.1: Summary of SIDM model variants and simulation results. SIDM model variants considered in this work and the main qualitative results of our zoom-in simulation of an MW-mass system for each case. The first column lists the DM model, the second and third columns list the characteristic velocity scale (in km s^{-1}) and amplitude of the self-interaction cross section (in $\text{cm}^2 \text{ g}^{-1}$) for our SIDM model variants, the fourth (fifth) column lists whether the host halo (subhalos) are cored by self-interactions, the sixth column lists whether subhalos are affected by ram pressure stripping due to self-interactions with the host, and the seventh column lists the fraction of subhalos among our matched subhalo populations that survive to $z = 0$ relative to the number of surviving subhalos with $V_{\text{peak}} > 20 \text{ km s}^{-1}$ in our CDM simulation.

9.4.3 Subhalo Definitions and Resolution Cuts

Throughout we employ the following definitions when referring to “subhalos”:

1. Surviving subhalos: DM systems identified by ROCKSTAR as distinct bound objects within the virial radius of the MW host halo at $z = 0$;
2. Disrupted subhalos: DM systems that have crossed within the virial radius of the MW halo at any simulation snapshot (where the virial radius of the MW halo is evaluated as a function of time) but are no longer identified as distinct bound objects by ROCKSTAR at $z = 0$ because they have deposited the majority of their particles onto the main-branch progenitor of the host at any earlier snapshot. Operationally, we require the descendant of each disrupted subhalo to be on the main branch of the host, following Chapter 2 [337].¹

Subhalos in our CDM simulation are well resolved down to a maximum peak circular velocity of $V_{\text{peak}} \approx 10 \text{ km s}^{-1}$, where V_{peak} is the largest maximum circular velocity a halo attains over its entire history ([315, 336]). However, we caution that ROCKSTAR is optimized for halo finding in collisionless DM simulations and therefore might be unable to reliably identify the more diffuse halos present in SIDM simulations, particularly near the resolution limit (X. Du, A. Peter, & C. Zeng 2019, private communication). Thus, we study subhalos above a very conservative resolution limit of $V_{\text{peak}} > 20 \text{ km s}^{-1}$ in our CDM simulation. This cut is applied to both surviving and disrupted subhalos, and it corresponds to subhalos resolved with more than ~ 1000 particles at the time V_{peak} is achieved.

¹For concreteness, we find that the average virial mass of disrupted subhalos in our CDM simulation is $\sim 10^8 M_\odot$ at the time of disruption, and that $\sim 90\%$ of these disrupting subhalos have virial masses below $10^9 M_\odot$. These values do not change significantly in our SIDM simulations.

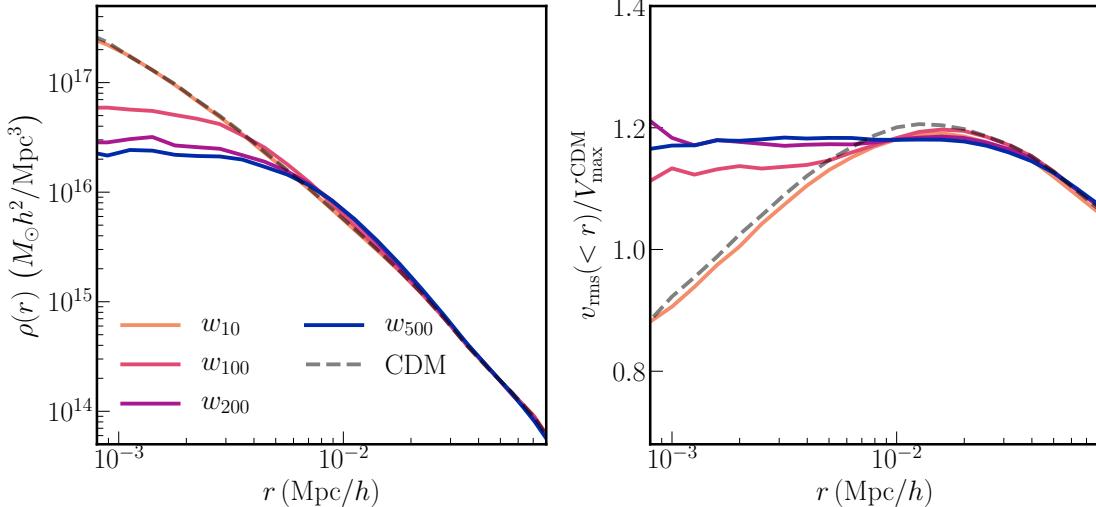


Figure 9.2: SIDM effects on the host halo. Left panel: host halo density profiles for our SIDM model variants. Right panel: corresponding velocity dispersion profiles. As the SIDM cross section at the characteristic velocity scale of the MW host halo increases, the inner regions of the host become increasingly cored and thermalized. Note that the host halo density and velocity dispersion profiles are nearly indistinguishable in CDM and w_{10} .

Due to this conservative resolution cut, we expect that artificial subhalo disruption (e.g., [454, 453]) is not a large source of error relative to accurate halo finding in the analysis of our SIDM simulations. It is important to note that our working definition of “subhalo disruption” cannot distinguish truly unbound systems from those that fall below the resolution limit of our simulations or the detection capabilities of our halo finder. Nonetheless, convergence tests based on higher-resolution resimulations of CDM and w_{500} presented in Appendix E.1 confirm that the trends reported in this paper, and particularly the correlation between w and the severity of subhalo disruption, are robust to the resolution of our simulations.

9.4.4 Subhalo Matching Procedure

For a fixed V_{peak} threshold, different numbers of surviving and disrupted subhalos may exist in our CDM and SIDM simulations at $z = 0$. Thus, to ensure that we analyze the same population of subhalos in CDM and SIDM, we adopt a variable V_{peak} threshold for our SIDM simulations, denoted by V_{thresh} . We define V_{thresh} such that the number of surviving-plus-disrupted subhalos with $V_{\text{peak}} > V_{\text{thresh}}$ is equal to the number of such subhalos in CDM with $V_{\text{peak}} > 20 \text{ km s}^{-1}$.² This yields thresholds of $V_{\text{thresh}} = [19.5, 19.05, 18.75, 18.7] \text{ km s}^{-1}$ for $[w_{10}, w_{100}, w_{200}, w_{500}]$, respectively.

²There are 154 surviving and disrupted subhalos in our CDM simulation with $V_{\text{peak}} > 20 \text{ km s}^{-1}$.

We have verified that these resolution cuts result in the same population of subhalos on an object-by-object basis by inspecting the initial phase-space region corresponding to each subhalo. We explore less conservative subhalo resolution thresholds in Appendix E.2, where we show that our main findings are robust to the chosen V_{peak} threshold.

We reiterate that authors have used various halo-finding algorithms and SIDM implementations to study the effects of self-interactions in MW-mass systems. Thus, direct comparisons of different simulations should be performed with care, particularly near the resolution limit, where halo finder shortcomings and other spurious numerical effects are expected to be most severe. We compare our findings to previous results in Section 9.8.

9.5 SIDM Effects on the MW Host Halo

We now present our results, focusing on the properties of the MW host halo and its subhalo population in each SIDM model variant described above. Table 9.1 lists the main qualitative results of each simulation.

The virial mass of the MW host halo is virtually identical in all of our simulations. However, the DM profile in the inner regions of the host varies significantly as a function of w . This is illustrated in Figure 9.2, which shows the density and velocity dispersion profiles of the host in CDM and in each of our SIDM model variants. As expected, SIDM models with larger self-interaction cross sections at the velocity scale set by the host's velocity dispersion ($\sim 200 \text{ km s}^{-1}$) exhibit cored density profiles, and the DM distribution in the inner regions of these hosts is roughly isothermal, consistent with previous findings (e.g., [125, 113, 461, 462, 392, 391]). In contrast, the host halo in w_{10} is very similar to that in CDM because self-interactions at the host's velocity scale are negligible in this case.³

As demonstrated by several authors (e.g., [253, 402, 391]), these results are expected to change in the presence of a central baryonic component such as the Galactic disk. In particular, because DM dynamically responds to the *total* (i.e., DM-plus-baryonic) gravitational potential, we expect our host halos to exhibit much smaller cores or even cusps in the presence of baryons. Importantly, this implies that the findings discussed below on the disruption of SIDM subhalos are strictly lower limits, since tidal disruption and ram pressure stripping would be more severe for the denser, hotter host halo expected in the presence of baryons.

To visualize the present-day DM structure in our simulations, Figure 9.3 shows the DM particle density in the phase space of radial distance versus radial velocity with respect to the center of the host. We observe that the DM profiles of both the host halo and its subhalos are less concentrated

³We have also examined the ellipticity profile of the host halo in each simulation. We find that hosts in model variants with larger values of w are more spherical, with typical ellipticities near unity, while the ellipticity profiles in w_{10} and CDM rise from ~ 0.9 in the inner regions to ~ 1 at the virial radius; these findings are consistent with many previous studies (see, e.g., [448]).

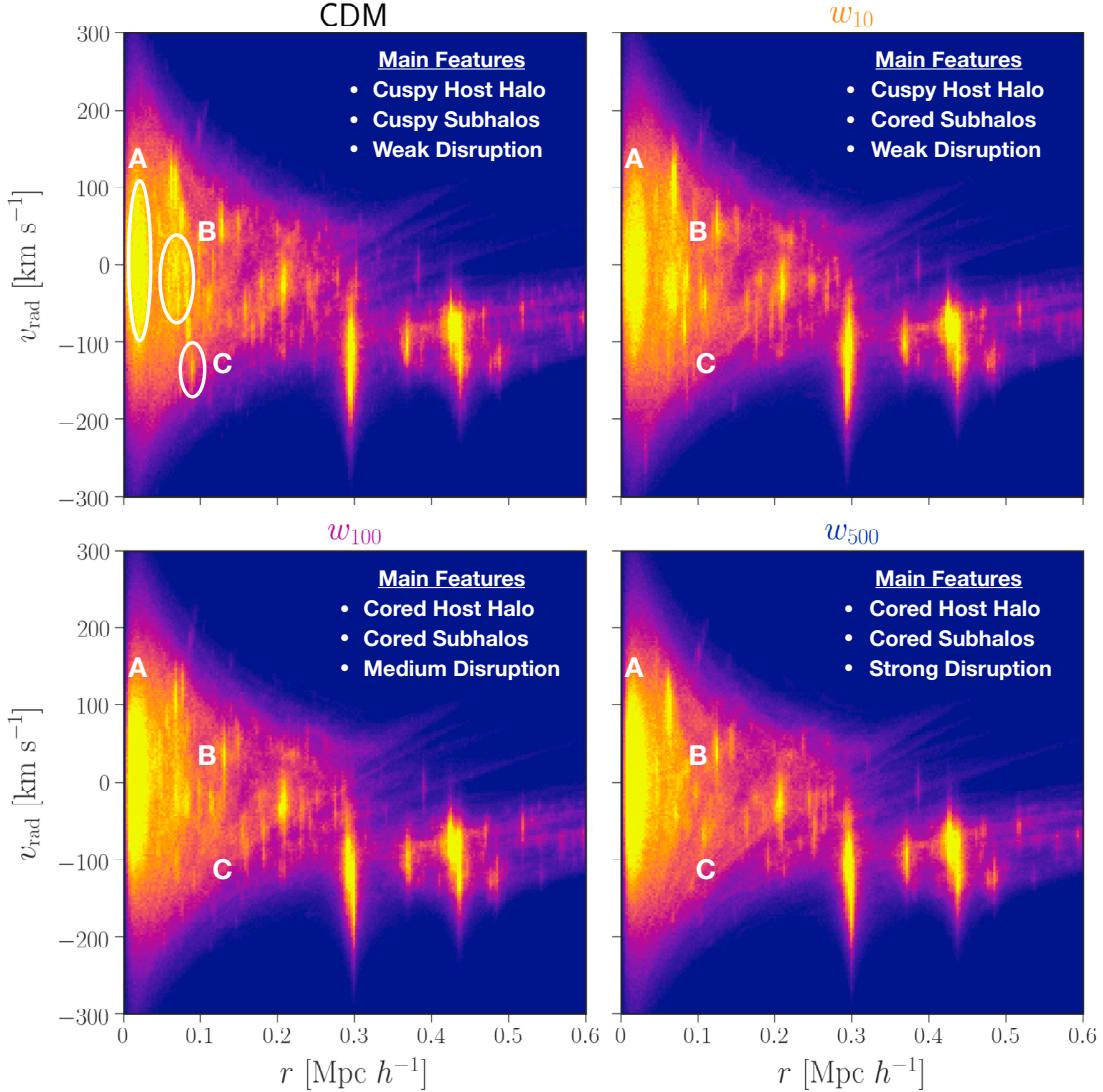


Figure 9.3: DM phase-space distributions for our zoom-in simulations of an MW-mass host halo in CDM and SIDM. The density of DM particles in bins of radial velocity and radial distance from the center of the host is plotted for CDM (top left) and for three of our SIDM model variants: w_{10} (top right), w_{100} (bottom left), and w_{500} (bottom right). These distributions qualitatively illustrate several of our main findings. For example, the host halo (labeled A) has a very similar phase space distribution in CDM and w_{10} , while subhalos in w_{10} (e.g., subhalo B) are somewhat less dense because of the large self-interaction cross section at low relative velocities in this case (see Figure 9.1). On the other hand, particles near the center of the host in w_{200} and w_{500} are preferentially scattered onto tangential orbits, and many of the low-mass subhalos that survive in CDM (e.g., subhalo C) are disrupted in these SIDM model variants owing to a combination of ram pressure stripping caused by self-interactions with the host and tidal stripping.

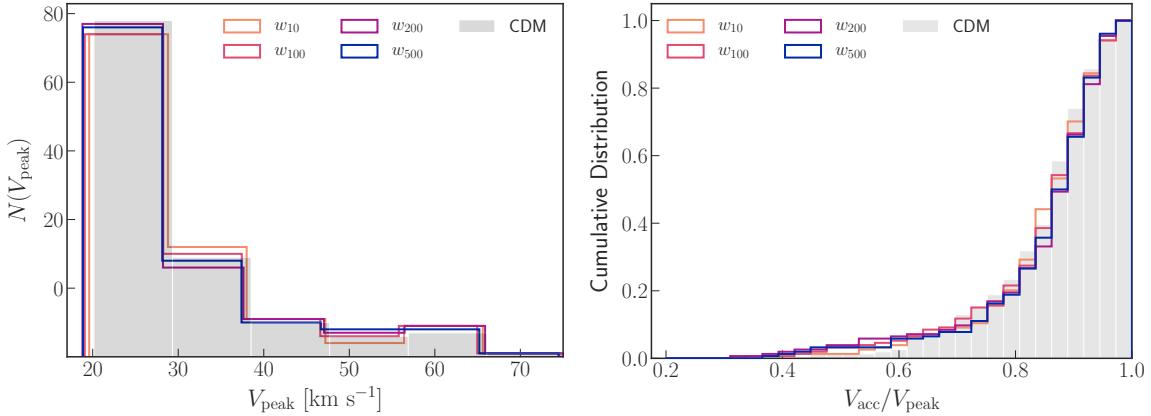


Figure 9.4: SIDM effects on subhalos before infall. Left panel: distributions of peak maximum circular velocity for surviving and disrupted subhalos in each SIDM model variant (unfilled histograms) vs. CDM (filled histogram). Right panel: cumulative distributions of the ratio of maximum circular velocity evaluated at the time of each subhalo's accretion onto the host divided by the peak maximum circular velocity along the main branch of the subhalo. Although subhalo assembly is statistically identical in CDM and SIDM, subhalos are mildly stripped by self-interactions prior to accretion onto the host halo in our SIDM simulations.

in our SIDM simulations. This effect is more significant for larger values of w ; for example, visual inspection of Figure 9.3 suggests that many prominent substructures that survive in CDM are completely disrupted in w_{500} . We also find that the phase-space density at small radii and low radial velocities with respect to the host center increases with w , implying that particles on radial orbits (i.e., those with high radial velocities in our CDM simulation) are scattered onto tangential orbits owing to self-interactions. Finally, we note that specific substructures in w_{10} appear more diffuse than their counterparts in CDM, which arises as a result of the large momentum transfer cross section at low relative velocities for w_{10} (see Figure 9.1).

The differences in the host halo's density and velocity dispersion profiles discussed above impact the post-infall evolution of subhalos. In particular, it is expected that the present-day ($z = 0$) abundance and properties of subhalos are affected by ram-pressure-like stripping due to self-interactions between subhalo and host halo particles, as well as by tidal effects in the host halo's gravitational field. We investigate these effects in Section 9.6.2.

Thus, the host halo is cored and thermalized in the absence of baryons if self-interactions are significant at the velocity scale set by the host's velocity dispersion. These effects will be weakened in the presence of baryons.

9.6 SIDM Effects on MW Subhalos

Next, we examine the pre- and post-infall evolution of all subhalos that fall into the host in each of our simulations. We then compare the corresponding *surviving* subhalo populations at $z = 0$.

9.6.1 Pre-infall Subhalo Evolution

In this subsection, we explore the formation and evolution of subhalos before they fall into the host.

Subhalo Assembly

First, we investigate the formation and initial properties of subhalos in each of our simulations. To do so, we plot the distribution of V_{peak} for each SIDM model variant in the left panel of Figure 9.4, and we compare these to CDM. The V_{peak} distributions are calculated using the matched subhalo populations defined in Section 9.4.4, and they include both subhalos that survive to $z = 0$ and those that fall into the host halo before disrupting. We find that the V_{peak} distributions are unchanged relative to CDM. In addition, the distributions of the time at which V_{peak} is achieved are nearly identical among the simulations, although there is a small amount of scatter on a subhalo-by-subhalo basis.

Thus, the formation times and initial properties of subhalos in all of our SIDM simulations—defined in terms of V_{peak} and the time at which V_{peak} occurs—are statistically identical to the corresponding quantities in our CDM simulation.

Effects of Early Self-interactions

Next, we assess the effects of self-interactions on subhalos before infall into the host. In particular, the right panel of Figure 9.4 shows the cumulative distribution of maximum circular velocity evaluated at the time of accretion onto the host, V_{acc} , divided by V_{peak} .⁴ This quantity captures the impact of self-interactions between the early time at which V_{peak} is usually achieved ($z_{\text{peak}} \sim 3$ for surviving subhalos, on average) and the time of accretion ($z_{\text{acc}} \sim 1$ for surviving subhalos, on average). These characteristic peak and infall times are very similar among our CDM and SIDM simulations.

We observe a subtle but systematic trend in $V_{\text{acc}}/V_{\text{peak}}$ as a function of the characteristic self-interaction velocity scale w . In particular, the low- $V_{\text{acc}}/V_{\text{peak}}$ tail of this distribution, which is present in CDM and results from pre-infall tidal stripping (e.g., [46, 480]), is more prominent in SIDM. This difference becomes increasingly pronounced for SIDM model variants with larger values of w , indicating that self-interactions at the typical relative velocity scale affect subhalos during pre-infall tidal stripping.

⁴Accretion is defined as the snapshot at which the center of a subhalo crosses into the virial radius of the host. Note that the virial radius is measured as a function of time using the ROCKSTAR output.

This finding is consistent with the fact that regions outside the conventional virial radius of the host halo can be quite dense (note that the self-interaction rate depends on the physical density, rather than the comoving density). In particular, typical splashback boundaries for MW-mass halos extend to ~ 1.5 times the conventional virial radius (e.g., [12, 139, 326]). Accordingly, we find that there is a noticeably smaller difference in the low- $V_{\text{acc}}/V_{\text{peak}}$ tail relative to CDM if V_{acc} is evaluated when subhalos cross within twice the virial radius of the host.

In Section 9.6.3, we show that the density profiles of subhalos in our SIDM model variants that exhibit appreciable changes to their $V_{\text{acc}}/V_{\text{peak}}$ distribution relative to CDM (i.e., all model variants other than w_{10}) already exhibit cored density profiles at infall. We provide a physical interpretation of the scaling with w in Section 9.6.2.

Thus, subhalos in our SIDM simulations are affected by self-interactions before accretion onto the host, leading to lower values of $V_{\text{acc}}/V_{\text{peak}}$ relative to subhalos in CDM.

9.6.2 Post-infall Subhalo Evolution

We now explore the evolution of subhalos in our SIDM model variants after infall into the host. In particular, we describe the physical effects that influence subhalos after infall, and we investigate subhalo disruption in our simulations.

Physical Effects

After infall, two main effects shape the subhalo populations in our SIDM simulations:

1. *Tidal stripping*: The gravitational field of the host halo tidally strips material from subhalos.

Note that this effect is velocity independent; in particular, assuming spherical symmetry and considering tidal interactions with the host halo only, the mass-loss rate due to tidal stripping can be written as [454]

$$\left(\frac{\dot{M}_{\text{sub}}}{M_{\text{sub}}} \right)_{\text{tidal}} = -\frac{1}{\alpha} \frac{1}{\tau_{\text{orb}}} \frac{M_{\text{sub}}(> r_t)}{M_{\text{sub}}}, \quad (9.6)$$

where $M(> r_t)$ is the mass of a subhalo contained outside of its tidal radius, τ_{orb} is the orbital timescale, and α is an order-unity constant. Note that the tidal radius (i.e., the distance from the center of the subhalo at which the host's tidal force is balanced by the subhalo's own gravity) depends on both the host halo and subhalo density profiles and the distance of the subhalo from the center of the host. Although tidal stripping occurs in CDM-only simulations, differences in the strength of this effect may arise in SIDM owing to changes in the density profiles of both the host halo and its subhalos caused by self-interactions. In particular, we expect that the cored density profile of the host halo in SIDM reduces the efficacy of tidal stripping relative to CDM for a fixed subhalo profile. However, as discussed in, e.g., [247],

the cores induced in SIDM subhalos make them more susceptible to disruption in a fixed tidal field.

2. *Ram pressure stripping*: Self-interactions with host halo particles drive material out of subhalos. Taken to the extreme, this process can completely dissociate subhalos in a phenomenon known as subhalo evaporation (see, e.g., [462]). Unlike tidal stripping, which strips mass from the outskirts of subhalos, ram pressure interactions remove material throughout their extent. In particular, subhalos experience an effective pressure owing to interactions with the host; the mass-loss rate due to these ram-pressure-like interactions can be written as [281]

$$\left(\frac{\dot{M}_{\text{sub}}}{M_{\text{sub}}} \right)_{\text{ram-pressure}} = -\chi_e(v_{\text{sub}}, v_{\text{esc,sub}}) \rho_{\text{host}}(\mathbf{r}) v_{\text{sub}}(\mathbf{r}) \frac{\sigma(v_{\text{sub}})}{m_\chi}, \quad (9.7)$$

where χ_e is an order-unity factor, $v_{\text{sub}}(\mathbf{r})$ is the velocity of the subhalo relative to the host evaluated along its orbit, $v_{\text{esc,sub}}$ is the escape velocity from the subhalo (which depends on the subhalo's density profile), $\rho_{\text{host}}(\mathbf{r})$ is the density of the host halo evaluated along the subhalo's orbit, and $\sigma(v_{\text{sub}})/m_\chi$ is the total SIDM cross section evaluated at v_{sub} .

The mass-loss rate due to ram pressure stripping depends on the SIDM cross section at the typical relative velocity scale between the host halo and its subhalos, which is set by the gravitational potential of the host. Because SIDM models with larger values of w have larger self-interaction cross sections at this velocity scale, the strength of ram pressure stripping increases with w .

By integrating Equation 9.7 over the course of a typical orbit, we find that subhalos with close pericentric passages to the center of the host ($d_{\text{peri}} \lesssim 70$ kpc) can lose $\sim 10\%$ of their infall mass owing to ram pressure stripping alone, for an isotropic cross section of $\sigma/m_\chi = 2 \text{ cm}^2 \text{ g}^{-1}$. Combined with tidal effects, and particularly the fact that subhalos with cored density profiles are more susceptible to tidal disruption, our calculations suggest that self-interactions at the relevant velocity scales are sufficient to severely strip subhalos in SIDM model variants with large values of w . We have confirmed that this behavior—i.e., increased mass loss relative to CDM followed by disruption due to tidal forces at pericenter—is indeed the dominant disruption mechanism in our simulations by inspecting the histories of matched subhalos that survive in CDM but disrupt in our SIDM model variants.

Thus, the evolution of subhalos after infall in our SIDM simulations is driven by a combination of tidal forces and ram-pressure stripping. Ram-pressure stripping is more effective for models with larger values of w , which have larger self-interaction cross sections at the host halo velocity scale. This mass-loss mechanism makes subhalos more susceptible to tidal disruption, particularly during pericentric passages.

Subhalo Disruption

We now analyze our simulation results to quantitatively assess the impact of the effects described in the previous section. In our CDM simulation, roughly half of all subhalos with $V_{\text{peak}} > 20 \text{ km s}^{-1}$ that fell into the host disrupt by $z = 0$. We find that even larger fractions of subhalos disrupt in our SIDM simulations, as indicated by the surviving subhalo fractions in Table 9.1. Based on the arguments in Section 9.6.2, we expect that extra subhalo disruption relative to CDM should be correlated with the SIDM cross section at the typical relative velocity scale set by the host’s gravitational potential, such that subhalos are disrupted more efficiently in SIDM model variants with larger values of w . As demonstrated below, our results are consistent with this hypothesis.

To build intuition about subhalo disruption in our SIDM model, we first investigate the distribution of pericentric distance, d_{peri} , for disrupted subhalos in each simulation. We define d_{peri} as the distance of closest approach to the center of the host halo during each subhalo’s first orbit around the host after infall.⁵ The top left and top right panels of Figure 9.5 show the distributions of d_{peri} for surviving and disrupted subhalos in each of our simulations. We observe that the amount of subhalo disruption at *large* pericentric distances increases strongly as a function of w , and we have verified that most of this disruption occurs during early pericentric passages. Subhalos in our SIDM simulations therefore disrupt at distances at which subhalos in our CDM simulation almost *never* disrupt, indicating that disruption is driven by tidal effects during early pericentric passages that are enhanced by ram-pressure-like interactions with the host.

To explore the data further, the bottom left and bottom right panels of Figure 9.5 show the distributions of maximum relative velocity with respect to the host evaluated along the orbit of each surviving and disrupted subhalo. The maximum relative velocity distributions for disrupted subhalos in our SIDM model variants peak strongly relative to CDM above $\sim 200 \text{ km s}^{-1}$, and the strength of this effect is correlated with w . This demonstrates that extra subhalo disruption relative to CDM is mainly driven by self-interactions at relative velocities above $\sim 200 \text{ km s}^{-1}$, and that subhalos are more easily disrupted in SIDM model variants with larger momentum transfer cross sections at these velocity scales.⁶

This picture is consistent with the fact that w_{500} —which has the largest momentum transfer cross section at large relative velocity scales—exhibits the most subhalo disruption among our SIDM model variants (see Section 9.6.3). Thus, our findings support the hypothesis that a combination of ram pressure and tidal stripping increases the amount of subhalo disruption relative to CDM.

Thus, a significant fraction of the subhalos that survive in our CDM simulation are disrupted in our SIDM simulations due to extra mass loss caused by ram-pressure stripping from self-interactions with the host at large relative velocities, which makes these subhalos more susceptible to tidal

⁵For infalling subhalos that have not completed a pericentric passage, we set d_{peri} equal to the distance from the center of the host at $z = 0$ (for surviving subhalos) or at the time of disruption (for disrupted subhalos).

⁶The bottom left panel of Figure 9.5 also indicates that a small fraction of surviving subhalos are impacted by ram pressure stripping owing to high-velocity interactions; we return to this point in Section 9.6.3.

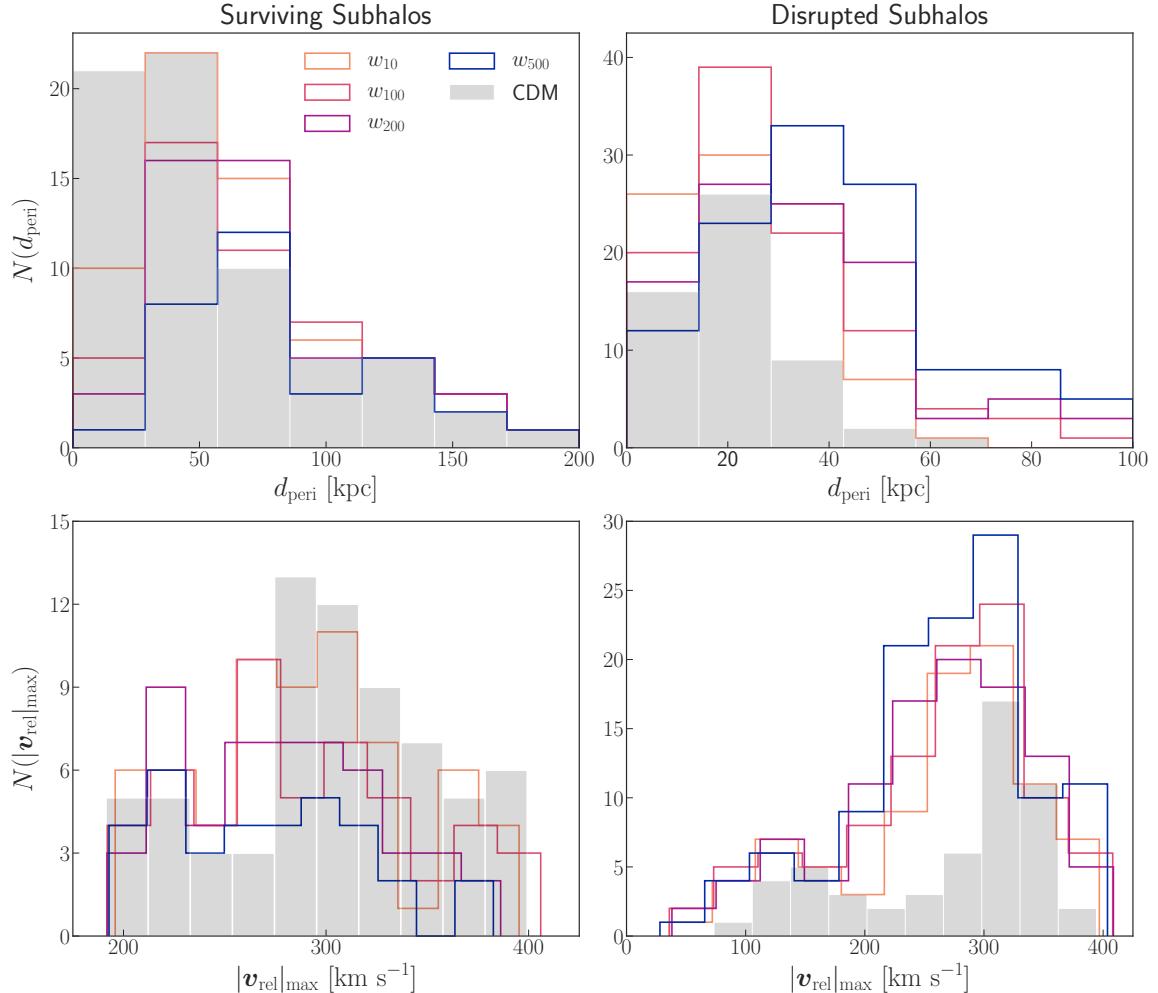


Figure 9.5: Properties of surviving and disrupted subhalos. Top left panel: distribution of the distance of closest approach to the host at first pericentric passage for surviving subhalos in each of our SIDM model variants (open histograms) vs. CDM (filled histogram). Top right panel: same as the top left panel, but for disrupted subhalos (surviving and disrupted subhalos are defined in Section 9.4.3). Bottom left panel: distribution of the maximum relative velocity with respect to the host halo evaluated along the orbit of each surviving subhalo. Bottom right panel: same as the bottom left panel, but for disrupted subhalos. Many subhalos that survive in our CDM simulation disrupt during early pericentric passages in our SIDM simulations, because ram pressure stripping caused by self-interactions with the host at large relative velocities makes subhalos more susceptible to tidal disruption.

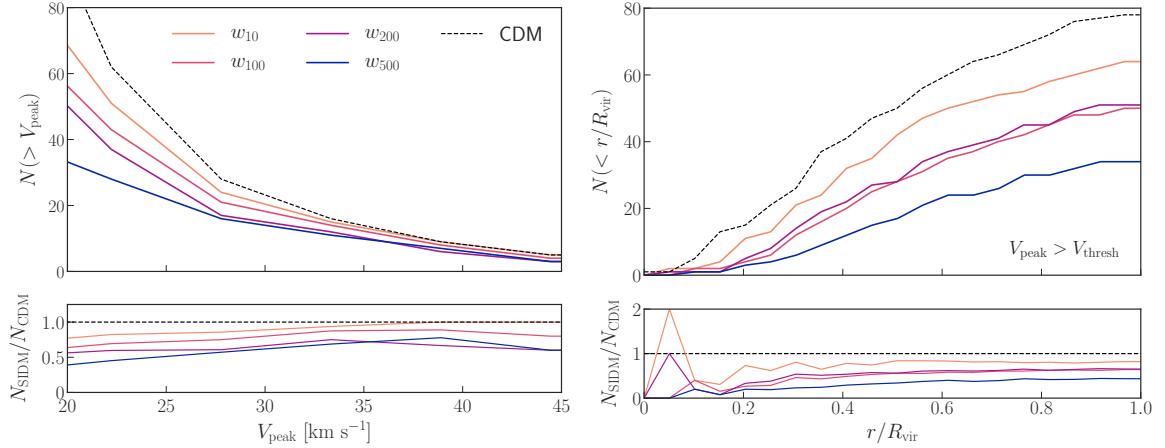


Figure 9.6: Surviving subhalo populations. Left panel: peak velocity function of subhalos in our CDM simulation and in each of our SIDM model variants. Right panel: corresponding radial subhalo distributions in units of the host halo virial radius in each simulation. The abundance of surviving subhalos is reduced in SIDM, and the strength of this effect increases with w owing to more significant ram pressure stripping caused by self-interactions with the host. Subhalo disruption in our SIDM simulations is particularly severe in the inner regions of the host halo.

disruption. This effect is more severe for models with larger values of w .

9.6.3 Surviving Subhalo Populations

We now examine the population statistics and properties of surviving subhalos in our SIDM simulations.

Population Statistics

In the left panel of Figure 9.6, we plot the cumulative number of surviving subhalos as a function of V_{peak} in each of our simulations. We find that the abundance of surviving subhalos monotonically decreases as a function of w . In particular, subhalo abundances are re-scaled in an approximately V_{peak} -independent fashion, and the number of surviving subhalos in SIDM divided by the number of surviving subhalos in CDM ranges from ~ 0.8 in w_{10} to ~ 0.4 in w_{500} . We have chosen to present these results in terms of V_{peak} because this quantity is expected to correlate more directly with satellite luminosity than, e.g., present-day virial mass or maximum circular velocity (e.g., [379, 288, 336, 333]).

We have verified that similar trends hold for the present-day M_{vir} and V_{max} functions; however, because these distributions mix the pre-infall properties and post-infall evolution of surviving subhalos, their shapes differ in detail from the corresponding peak functions. For example, we observe an *enhancement* in subhalo abundance at large values of V_{max} for all of our SIDM model variants

relative to CDM. We speculate that this is a consequence of less effective tidal stripping for surviving subhalos in our SIDM simulations, which occupy tangential orbits around a cored host halo. Previous authors have also found hints of this trend (e.g., [392]; also see Section 9.8).

In the right panel of Figure 9.6, we plot the radial distribution of surviving subhalos in each simulation. Similar to the V_{peak} functions, we find that the radial distributions are approximately rescaled in the outer regions. However, disruption becomes increasingly severe near the center of the host (i.e., $r/R_{\text{vir}} \lesssim 0.5$), causing $N_{\text{SIDM}}/N_{\text{CDM}}$ to decrease sharply in that region, except for the “spikes” at small radii observed for w_{10} and w_{200} . These correspond to either one or two additional subhalos near the center of these hosts; we investigate these subhalos further in Appendix E.3.

Thus, subhalo disruption due to self-interactions approximately re-scales the number of surviving subhalos relative to CDM as a function of V_{peak} . Subhalo disruption is particularly effective in the inner regions of the host halo, and it is more severe for models with larger values of w .

Orbital Anisotropy Profile

An immediate consequence of the ram-pressure-plus-tidal-stripping disruption mechanism described above is that subhalos on radial orbits are preferentially disrupted. We therefore expect surviving subhalos in SIDM to preferentially occupy tangential orbits relative to surviving subhalos in CDM. We test this by measuring the orbital anisotropy profile

$$\beta(r) \equiv 1 - \frac{\sigma_t(r)^2}{2\sigma_r(r)^2}, \quad (9.8)$$

where σ_r and σ_t denote the radial and tangential velocity dispersion of subhalos, respectively. Note that $\beta < 0$ corresponds to a tangentially biased orbital distribution, $\beta = 0$ corresponds to an isotropic orbital distribution, and the maximum allowed value of $\beta = 1$ corresponds to a purely radial orbital distribution.

We measure $\beta(r)$ by calculating the radial and tangential velocity distributions of subhalos in the frame of the host halo, binning subhalos in distance from the center of the host. The result is shown in Figure 9.7. We find that the orbital anisotropy profile rises toward the outskirts of the host halo in our SIDM simulations, while it is roughly flat in CDM. As expected based on the results in Section 9.6.2, surviving subhalos preferentially occupy tangential orbits in our SIDM simulations, and this effect is more significant for model variants with larger values of w , for which subhalo disruption via ram pressure plus tidal stripping is more severe.⁷ We observe a similar trend for anisotropy profiles computed using the DM particles belonging to the host halo in each simulation, although the magnitude of the effect is less severe in this case. An interesting consequence of this effect is that, at fixed V_{peak} , surviving subhalos in our SIDM simulations tend to be *more* massive on average than

⁷The amplitudes of the $\beta(r)$ profiles at small radii are sensitive to the choice of radial binning because there are few resolved subhalos in the inner regions. However, the overall shape of $\beta(r)$ and the trends with w are robust.

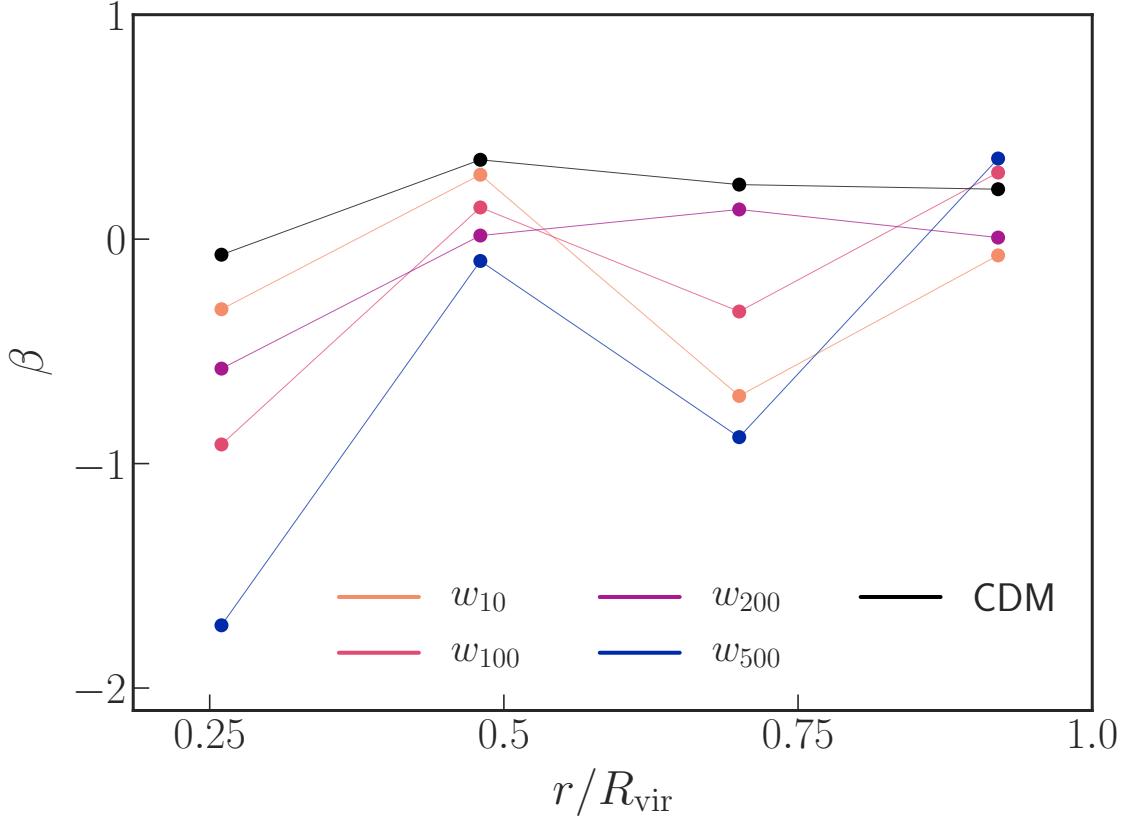


Figure 9.7: Orbital anisotropy profile of subhalos in our SIDM simulations. Surviving subhalos in SIDM model variants with larger values of w occupy tangentially biased orbits relative to CDM. This occurs because a combination of ram pressure and tidal stripping preferentially disrupts subhalos on radial orbits.

their counterparts in CDM, since they are less susceptible to tidal stripping on tangential orbits; however, this difference largely vanishes if subhalos are matched based on their orbital properties.

These results are interesting in light of recent high-precision measurements of the orbital properties of MW satellites enabled by Gaia (e.g., [181, 183, 424]). In particular, recent studies of the orbital anisotropy profile find that the typically flat $\beta(r)$ profiles found in DM-only simulations are in tension with the observed velocity anisotropy profile of MW satellites [385]. Baryonic effects, and particularly the tidal influence of the Galactic disk, affect the $\beta(r)$ profile in a similar manner to ram pressure stripping in our SIDM simulations by disrupting subhalos on radial orbits near the center of the host. However, as noted by [385], the changes to $\beta(r)$ due to baryonic physics only resolve the discrepancy with the observed profile if the Galactic disk is sufficiently massive, which in turn depends on the prescription for baryonic feedback and the mass accretion history of the MW system. Adding baryonic subhalo disruption to an orbital distribution that is *already* tangentially

biased owing to self-interactions may result in even more tangential bias than observed in the inner regions of the MW, potentially yielding an upper limit on the SIDM cross section at $\sim 200 \text{ km s}^{-1}$.

Thus, surviving subhalos in SIDM occupy tangentially biased orbits relative to subhalos in CDM, leaving a systematic imprint on the orbital anisotropy profile.

Subhalo Profiles

Because the momentum transfer cross section in our SIDM model scales as v^{-4} above the characteristic velocity scale w , self-interactions at low relative velocities are strictly more efficient at transferring momentum than interactions at high relative velocities. Thus, we expect the density profiles of low-mass subhalos in our SIDM simulations to be impacted by self-interactions. In contrast to subhalo disruption, it is not clear *a priori* which SIDM model variants will most significantly affect subhalo profiles, since both self-interactions within subhalos and self-interactions with the host halo potentially impact subhalo density profiles.

To investigate the impact of self-interactions on subhalo density profiles, we select representative subhalos in our CDM simulation, and we find the corresponding subhalos in our SIDM simulations by matching on V_{peak} , accretion time, and pericentric distance. We then select the particles in the initial phase-space region associated with these matched subhalos and we track the evolution of their density profiles until $z = 0$. Figure 9.8 shows the results of this procedure for a particular set of subhalos matched to a CDM subhalo with $V_{\text{peak}} \approx 40 \text{ km s}^{-1}$ and accretion time $z_{\text{acc}} \approx 1$ that passes close to the center of the host ($d_{\text{peri}} \approx 70 \text{ kpc}$). We show the density profiles of these matched subhalos at the time of accretion onto the host and at $z = 0$; we also show the density profile of a subhalo in w_{500} with similar V_{peak} and z_{acc} that does *not* pass close to the center of its host ($d_{\text{peri}} \approx 160 \text{ kpc}$).

We find that subhalos in SIDM model variants with larger values of w have lower inner densities and are more cored, i.e., their inner density profiles are flatter as a function of distance from the center of the subhalo than in CDM. We interpret this as a consequence of ram pressure stripping; even though self-interactions within subhalos are more effective for SIDM models with *smaller* values of w , interactions with the host at large relative velocities significantly alter the inner profiles of subhalos in model variants with *larger* values of w . As indicated by Figure 9.8, we find that—at fixed V_{peak} and z_{acc} —subhalos with closer pericentric passages to the center of the host are significantly more cored. Detailed orbital modeling of satellites in MW-mass systems is therefore crucial in order to interpret the inferred DM density profiles of their subhalos.

We find that the circular velocity profiles of subhalos with close pericentric passages are significantly altered relative to CDM, as expected based on the changes to their density profiles and consistent with previous findings [461, 493, 390, 391, 177, 403]. This has implications for SIDM solutions to the diversity and “too big to fail” problems concerning MW satellites [247, 495], although baryonic feedback mechanisms, including heating from supernova feedback, can also reduce

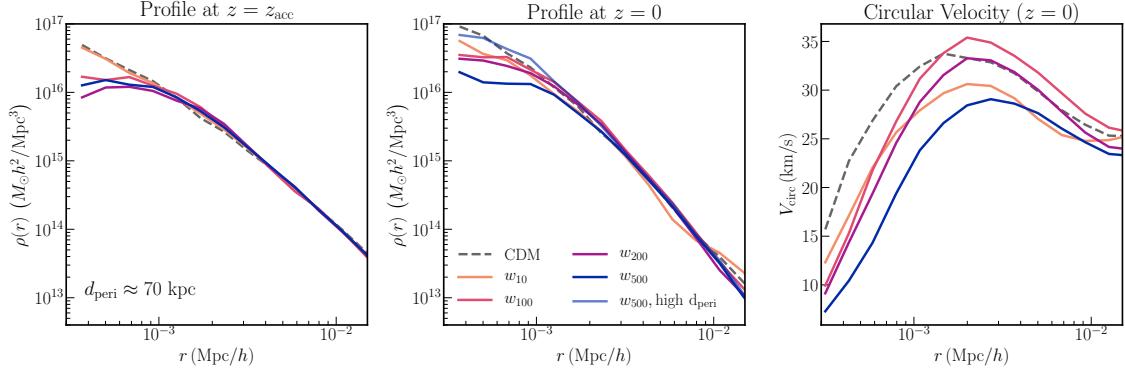


Figure 9.8: DM profiles of a matched set of surviving subhalos in our CDM and SIDM simulations. Density profiles defined by the initial set of bound particles are shown for the same subhalo with $V_{\text{peak}} \approx 40 \text{ km s}^{-1}$ and $z_{\text{acc}} \approx 1$ at the time of accretion onto the host (left panel) and at $z = 0$ (middle panel). Subhalos in model variants with larger values of w have lower-amplitude, flatter inner density profiles owing to self-interactions with the host halo. The light-blue line in the middle panel shows the density profile for a subhalo with similar V_{peak} and z_{acc} but with a *large* pericentric distance ($d_{\text{peri}} \approx 160$ kpc). The discrepancy between this profile and that of the corresponding low-pericenter subhalo demonstrates that the impact of self-interactions on subhalo density profiles depends sensitively on their orbital properties. The right panel shows the corresponding circular velocity profile for each subhalo at $z = 0$.

and flatten subhalos' central density profiles [369, 91, 117, 405, 378]. However, if cores are created by stellar feedback, then subhalos' inner density profiles are expected to be correlated with their galaxies' star formation histories (e.g., [378]). Meanwhile, in SIDM, we expect the inner amplitude and flatness of subhalos' density profiles to correlate most directly with their orbital histories, which can be encapsulated by their infall times and orbital eccentricities.

Thus, surviving subhalos in SIDM models with larger values of w have lower-amplitude, flatter density profiles relative to their CDM counterparts. The magnitude of this effect depends on subhalos' orbital properties such as their infall time and distance of closest approach to the center of the host, and this effect is more significant for larger values of w .

9.7 Challenges

We now discuss several caveats and systematics associated with our analysis and main results.

9.7.1 Numerical Effects

Various halo-finding algorithms, including SUBFIND, AHF, and ROCKSTAR have been used to identify and track halos in cosmological SIDM simulations. Several authors have thoroughly compared

these algorithms in the context of CDM simulations (e.g., [267, 434]). However, the pros and cons of different halo finders and merger tree algorithms are largely unknown for SIDM simulations, and a comprehensive comparison study is crucial in order to go beyond the statistical halo matching technique employed in this paper. Such a study would also be relevant for robust halo finding in hydrodynamic simulations, where baryonic effects like supernova feedback can substantially alter density profiles relative to CDM.

Meanwhile, resolution effects (including artificial disruption; [454, 453]) are always important to mitigate when analyzing substructure in cosmological simulations, particularly if the abundance of objects near the resolution limit is important. Thus, a detailed study of the underlying mechanisms and numerical stability of subhalo disruption in the presence of self-interactions using extremely high-resolution simulations is another important avenue for future work.

9.7.2 Sample Variance

Because we have simulated a fixed realization of an MW system for various SIDM model variants, we have clearly not sampled a cosmologically representative range of host halo mass accretion histories. Doing so will add scatter to the prediction for the amount of subhalo disruption in SIDM relative to CDM because the number of infalling subhalos and their properties depend on the mass accretion history of the host (e.g., [315]). However, the main physical trends we report, and particularly the correlation between the severity of subhalo disruption and the SIDM cross section evaluated at the host halo velocity scale, should not be affected by marginalizing over mass accretion histories. We plan to test this explicitly by running a suite of zoom-in SIDM simulations. When analyzing MW satellites specifically, certain features of the MW system, such as the properties of the Large Magellanic Cloud system and major accretion events inferred from Gaia data, may lessen the impact of this uncertainty by constraining the allowed range of assembly histories (e.g., [333]).

9.7.3 Baryonic Effects

Our analysis is meant to provide insights into the basic physical mechanisms that shape the properties of the MW system in the presence of DM self-interactions. Although we do not include baryonic effects in our simulations, a number of authors have studied SIDM effects on MW and dwarf galaxy scales in the presence of baryons [253, 182, 390, 391, 158, 402, 177]. These studies have revealed that including baryons impacts both the host halo and subhalos in SIDM-only simulations. We now discuss these effects in light of our findings.

Changes to Density Profiles

One of the key predictions of the above studies is that the presence of baryons changes SIDM host halo profiles. This occurs because SIDM dynamically responds to the total (i.e., DM-plus-baryonic)

gravitational potential (e.g., [253, 402, 391]). This response results in a cuspier host that is nearly indistinguishable from the host halo in a CDM-plus-baryon simulation. Because this process makes the host cuspier, the subhalo disruption effects reported in this paper should be viewed as lower limits, since both tidal effects and ram pressure stripping will be enhanced in the presence of baryons owing to the host’s cuspier inner density profile and increased velocity dispersion. As discussed by [247], we do not expect the change to the host halo’s density profile to significantly alter surviving subhalo populations, since the host’s profile only changes within the few inner kiloparsecs, where very few subhalos survive in our SIDM simulations (see, e.g., Figures 9.5 and 9.6). However, including the Galactic disk significantly enhances the likelihood of disruption for subhalos with $d_{\text{peri}} \lesssim 20$ kpc (e.g., [190]).

Unlike the host halo, we do not expect the density profiles of subhalos in our SIDM simulations to change appreciably in the presence of baryons, although they may become more cored owing to tidal interactions with the cuspier host halo and with the Galactic disk. In particular, faint satellites have extremely low stellar masses, which makes both heating from supernova feedback and adiabatic contraction of the DM profile negligible (e.g., [157, 158]). In other words, the presence of a central dwarf galaxy has little effect on the DM profile of its subhalo.

Accelerated Core Collapse

By choosing SIDM model variants with reasonably small self-interaction cross sections at all but the lowest velocity scales, we have explicitly avoided studying gravothermal core collapse [38, 20, 269]. Of particular relevance for subhalos orbiting the MW, recent analyses demonstrate that the timescale for core collapse can be significantly shortened owing to tidal stripping by the Galactic disk [247, 345]. This accelerated core collapse mechanism has been used to argue that the survival of dense satellites near the center of the MW, and particularly the apparent anticorrelation between satellites’ pericentric distances and their inferred central DM densities, is a unique signature of SIDM [256].

We have chosen to study models in which core collapse is not likely to occur over the timescale of our simulations because this generally requires a very large SIDM cross section, at least at low velocities (e.g., $\sigma_T/m_\chi \gtrsim 3 \text{ cm}^2 \text{ g}^{-1}$ at $v \sim 30 \text{ km s}^{-1}$; [247, 345]). This makes the effect somewhat model dependent because such large cross section amplitudes are likely ruled out by MW satellite abundances for velocity-independent scattering. Furthermore, it is important to study the region of parameter space for velocity-dependent SIDM models that does not result in core-collapsed MW subhalos, since it is not clear whether core collapse is required by the current data.

9.8 Comparison to Previous Studies

Comparing our results to previous studies of SIDM effects on MW-mass systems is not straightforward. For clarity, we limit this discussion to a comparison of results for subhalo abundances. The following theoretical systematics should be kept in mind throughout the discussion in this section:

1. *SIDM cross section (amplitude and velocity dependence)*: Authors have studied a variety of velocity-dependent SIDM models with different cross section amplitudes. We comment on case-by-case comparisons with our model below.
2. *SIDM implementation*: Most SIDM implementations in the literature are functionally identical to ours, with the exception of [461] (and, by extension, [493] and [147]). In particular, we determine whether a given DM particle scatters by calculating the interaction probability with all neighboring particles within a sphere of radius equal to the softening length, while [461] calculate interaction probabilities using the $k = 38$ nearest neighbors of each particle.
3. *Halo finding*: Various halo-finding algorithms have been used to analyze SIDM simulations of MW-mass systems.

Our findings are in reasonable agreement with [461]. These authors claim that a velocity-independent SIDM model with $\sigma_T/m_\chi = 10 \text{ cm}^2 \text{ g}^{-1}$ (i.e., our w_{500} model with a 10 times higher cross section) is ruled out by the abundance and density profiles of bright MW satellites. These authors find that subhalo disruption in this velocity-independent model is driven by evaporation due to subhalo–host halo interactions, similar to our conclusion that ram pressure stripping plus tidal effects drive subhalo disruption in w_{500} . [461] also find that two velocity-dependent SIDM models, both of which are similar to our w_{10} model but with slightly larger momentum transfer cross sections at the host halo velocity scale, are allowed by the MW satellite data. By visual inspection, it appears that subhalo abundances are suppressed by $\sim 20\%$ relative to CDM in these models, with the largest differences at low subhalo masses. The amount of subhalo disruption in our w_{10} model is consistent with this result (see, e.g., Figure 9.6), although we find that the suppression of subhalo abundance relative to CDM is a much weaker function of subhalo mass, possibly due to our use of a matched subhalo sample.

[392] find that the abundance of MW subhalos in SIDM models with velocity-independent cross sections of $\sigma/m_\chi = 0.1$ and $1 \text{ cm}^2 \text{ g}^{-1}$ is very similar to that in CDM. However, they find that $\sim 20\%$ fewer subhalos survive near the inner regions of their host halo ($r \lesssim 0.5R_{\text{vir}}$) for $\sigma/m_\chi = 1 \text{ cm}^2 \text{ g}^{-1}$, corresponding to $\sigma_T/m_\chi = 0.5 \text{ cm}^2 \text{ g}^{-1}$. Our w_{100} model, in which $\sim 35\%$ of subhalos are disrupted relative to CDM, has a similar momentum transfer cross section at the MW-mass host halo velocity scale. Thus, our results are roughly consistent with [392], although further investigation into halo finder differences is needed to assess whether the remaining $\sim 15\%$ discrepancy is due to the velocity dependence of our w_{100} model. Interestingly, [392] find that subhalo abundances

for $\sigma/m_\chi = 0.1 \text{ cm}^2 \text{ g}^{-1}$ are identical to CDM at low V_{\max} , and that there are actually *more* high- V_{\max} subhalos than CDM in this case. Although we cannot compare our results to their $0.1 \text{ cm}^2 \text{ g}^{-1}$ model directly, we similarly find that the high- V_{\max} tail of the surviving subhalo distribution is enhanced in our SIDM simulations relative to CDM, which might be due to the fact that surviving subhalos occupy tangential orbits on which tidal stripping is less severe.

[493] simulate velocity-independent SIDM models with $\sigma_T/m_\chi = 0.1, 1, \text{ and } 10 \text{ cm}^2 \text{ g}^{-1}$, as well as two velocity-dependent models that lie between our w_{10} and w_{100} models at the MW-mass host halo velocity scale, but which rise more steeply at low relative velocities. They find that the only model that leads to a difference in substructure abundance relative to CDM is $\sigma_T/m_\chi = 10 \text{ cm}^2 \text{ g}^{-1}$, in stark contrast to our finding that subhalo disruption is significant in all of our model variants. This may be explained by the difference in SIDM implementation, since interactions between subhalo and host halo particles are less likely in the [493] implementation than in ours. In particular, our implementation allows interactions between any DM particles within a softening-length-sized sphere, while the k -nearest neighbors technique used in [493] does not. Although [493] use the same SIDM implementation as [461], we speculate that there is less of a discrepancy between our results and those of [461] because these authors only study models in which the cross section is either negligible or extremely large at the host halo velocity scale.

We note that [147] study stellar stripping using the same set of SIDM models and implementation as [493]. These authors reiterate that substructure is only affected in the most extreme SIDM model, although they remark that subhalos with close pericentric passages in any of the models typically evaporate within ~ 6 Gyr, which is qualitatively consistent with our findings.

Finally, although they focus on the impact of adding a Galactic disk to SIDM simulations, [391] simulate MW host halos without disks in CDM and SIDM. The subhalo V_{\max} functions and radial distributions shown in [391] suggest that $\sim 60\%$ of subhalos survive in their $\sigma_T/m_\chi = 1 \text{ cm}^2 \text{ g}^{-1}$ model relative to CDM, which is in reasonable agreement with our findings. Again, this effect is only noticeable at low V_{\max} . Like the results in [392], and consistent with our findings, the SIDM-only model in [391] shows an *enhancement* in the abundance of high- V_{\max} subhalos relative to CDM.

9.9 Prospects for SIDM Constraints

The effects reported in this work can be incorporated in models of DM substructure in MW-mass systems in order to constrain the SIDM cross section at relatively low velocity scales. For example, several authors have recently proposed forward-modeling frameworks in which various DM properties can be constrained based on the abundance of observed MW satellites (e.g., [241, 335, 336]). These constraints are set by the fact that deviations from CDM yield a suppression in the abundance of the low-mass subhalos that are inferred to host faint satellite galaxies. Incorporating our finding that subhalo disruption is driven by the self-interaction cross section at the host halo velocity scale

will therefore yield an upper limit on the SIDM cross section at $\sim 200 \text{ km s}^{-1}$. The precise value of this limit is difficult to forecast given the theoretical uncertainties discussed above; thus, we plan to carry out this study comprehensively in future work. Meanwhile, gaps and perturbations in stellar streams—which are also sensitive to the abundance, radial distribution, and density profiles of both surviving and disrupted subhalos in the MW—can be used to place complementary constraints (e.g., [81]).

Our results suggest that placing an upper limit on the SIDM cross section at velocity scales *below* $\sim 200 \text{ km s}^{-1}$ by probing subhalo abundances is challenging, or at least highly model dependent, since the abundance of surviving subhalos is not very sensitive to the SIDM cross section below the velocity scale of the host. However, if observations of the inner density profiles of the faintest MW satellites unambiguously favor cored or cuspy inner DM density profiles, it may be possible to place a stringent limit on σ_T/m_χ at low relative velocities (e.g., $\sim 10 \text{ km s}^{-1}$). [377] claim that the cuspy halo profile inferred for the MW satellite Draco places an upper bound on the velocity-independent SIDM cross section of $\sigma/m_\chi \lesssim 0.6 \text{ cm}^2 \text{ g}^{-1}$ ($\sigma_T/m_\chi \lesssim 0.3 \text{ cm}^2 \text{ g}^{-1}$) at 99% confidence. However, as we have argued, modulations to subhalos’ density profiles in the presence of self-interactions are highly dependent on their orbital histories; in addition, it is not clear how the constraint in [377] translates to the velocity-dependent SIDM model considered here. We also note that the stellar profiles of surviving satellites can evolve differently depending on whether they occupy a cored or cuspy halo (e.g., [165]).

High-resolution spectroscopy on giant segmented mirror telescopes (e.g., [423]), along with improvements in mass profile modeling techniques (e.g., [194, 378, 286]), will increase the precision of density profile measurements and constraints. However, observational systematics associated with inferring the underlying DM density profiles (e.g., [366, 351]), and degeneracies with baryonic coring mechanisms (e.g., [378]), will likely make a *detection* of DM self-interactions at low relative velocities challenging. On a positive note, our results demonstrate that observables that cover a range of velocity scales related to MW-mass systems—such as the inferred abundances and density profiles of low-mass subhalos—are highly complementary. In particular, combining such observables informs the velocity dependence of the SIDM cross section, which is a key facet of many well-motivated SIDM models.

9.10 Conclusions

Given increasingly precise constraints on the abundance and internal properties of DM substructure enabled by observations of satellite galaxies, strongly lensed systems, and stellar streams, a detailed examination of DM models at the edge of allowed parameter space is crucial. SIDM is particularly interesting in this context, since constraints on the velocity-dependent self-interaction cross section at relative velocities that are small compared to those typical for galaxy clusters provide an important

handle on particle physics in the dark sector. In this paper, we studied the phenomenology of a generic, velocity-dependent SIDM model in the context of the DM host halo and subhalos of MW-mass systems. Due to its velocity dependence, the effects of our SIDM model are inherently scale dependent. Thus, observations of MW-mass halos and their subhalos provide a range of velocity scales with which to probe DM self-interactions.

We have demonstrated that the characteristic velocity above which momentum transfer due to self-interactions becomes inefficient, which is related to the DM and mediator masses in our SIDM model, has a variety of phenomenological consequences for MW-mass host halos and their subhalos. Our main findings are as follows:

1. In the absence of baryons, the DM distribution in the inner regions of the host is cored and thermalized owing to self-interactions (Figure 9.2). This effect is stronger for larger values of the cross section at the host halo velocity scale, corresponding to our model variants with larger values of w .
2. The initial assembly of subhalos in our SIDM simulations, quantified by V_{peak} and the time at which V_{peak} is achieved, is statistically similar to that in CDM (Figure 9.4).
3. Subhalos are mildly affected by self-interactions before infall into the host, leading to lower maximum circular velocity values and moderately cored density profiles at infall relative to subhalos in CDM (Figures 9.4 and 9.8).
4. A significant fraction of the subhalos that survive in CDM are not found in SIDM owing to extra mass loss from ram pressure stripping caused by self-interactions with the host halo, which makes subhalos more susceptible to tidal disruption. This effect is more severe for models with larger values of w (Table 9.1, Figures 9.5 and 9.6).
5. Surviving subhalos in SIDM occupy tangentially biased orbits relative to those in CDM, causing a systematic trend in the orbital anisotropy profile that is more significant for models with larger values of w (Figure 9.7).
6. Surviving subhalos in SIDM models with larger values of w are less dense and more cored than surviving subhalos in CDM, and the magnitude of this effect depends sensitively on orbital properties such as infall time and pericentric distance, with more severe coring for subhalos that pass closer to the center of their host (Figure 9.8).

Given these findings, we plan to carry out a comprehensive study of these effects for a simulated MW-like host halo that includes a realistic Large Magellanic Cloud analog system, which has recently been used to fit the full-sky MW satellite luminosity function [333]. This analysis will allow us to constrain the SIDM cross section at the MW host halo velocity scale ($\sim 200 \text{ km s}^{-1}$) using the abundance, surface brightness distribution, and radial distribution of observed MW satellites.

We emphasize that self-interactions also affect the orbital distribution and density profiles of subhalos in MW-mass systems. Comparing these quantities to data in a forward-modeling approach will help break degeneracies among SIDM models that are not ruled out by subhalo abundances alone, and will therefore inform the velocity dependence of the SIDM cross section. For example, surviving subhalos in our SIDM simulations preferentially occupy tangential orbits, and this prediction can be compared to the measured orbital distribution of MW satellites using proper-motion measurements from Gaia and its future data releases. Meanwhile, surviving subhalos in SIDM exhibit cored density profiles, and the strength of this effect is a function of both microphysical SIDM parameters and the orbital history of each subhalo.

We expect that combining high-precision spectroscopic and proper-motion measurements of satellite galaxies with analyses of strongly lensed systems and perturbations in stellar streams will test these predictions, providing a new window into velocity-dependent DM self-interactions.

Chapter 10

Disentangling Dark Matter and Baryonic Physics

Abstract

We compare the signatures of different dark matter models and baryonic physics on properties of the MW subhalo population. In particular, we consider the effects of self-interacting dark matter (SIDM), warm dark matter (WDM), and the Galactic disk on the peak velocity function, radial distribution, and spatial distribution of MW subhalos using cosmological zoom-in simulations of MW-like systems that include realistic Large Magellanic Cloud (LMC) analogs. For a fixed abundance of subhalos expected to host dwarf galaxies ($V_{\text{peak}} \gtrsim 20 \text{ km s}^{-1}$), SIDM and WDM can produce a similar mass-dependent suppression of the subhalo velocity function, while disruption due to the MW disk is mass-independent. Subhalos in the inner regions of the MW are preferentially disrupted by both self-interactions and the Galactic disk, while WDM suppression is radially independent. Spatial anisotropy in the MW subhalo population due to the LMC system persists in our disk disruption model and is not strongly affected in WDM. Meanwhile, enhanced tidal disruption in SIDM caused by self-interactions within the LMC halo at early times and during LMC infall can erase this anisotropy. These effects provide avenues to distinguish dark matter models and baryonic physics by combining multiple properties of the MW subhalo population as traced by satellite galaxies and stellar streams.

10.1 Paper Status and External Contributions

This chapter is in preparation, to be submitted to the *Astrophysical Journal* with the title, “Disentangling the Effects of SIDM, WDM, and Baryons on the Milky Way Subhalo Population,” on which I will be the corresponding author. It is the result of a collaboration with Arka Banerjee,

Susmita Adhikari, Yao-Yuan Mao, and Risa Wechsler. Arka and Susmita led the development of the self-interacting dark matter cosmological simulation algorithm, and all of us worked together to run and analyze the simulations presented below. In addition, all coauthors made editorial contributions to the text.

10.2 Introduction

Dark matter subhalos within the Milky Way (MW) provide a powerful means to test theories of structure formation and dark matter physics on the smallest observationally accessible cosmic scales. For decades, the “missing satellites problem” concerning MW satellite galaxies [266, 325] inspired dark matter model-building efforts focused on reducing the abundance of low-mass subhalos. This can be achieved by suppressing the primordial matter power spectrum—as in models of warm [11], interacting [74, 335], and fuzzy [227, 228] dark matter—or dynamically, as in models of self-interacting dark matter (SIDM; [432, 448]).

Our understanding of the MW subhalo population will drastically improve in the near future due to surveys including the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST; e.g., see [151, 43]). These surveys will probe both surviving subhalos, which are traced by satellite galaxies and may perturb stellar streams, and disrupted subhalos, which contribute to the buildup of the MW’s dark matter halo and may deposit their luminous content in the stellar halo.

A plethora of dark matter microphysics can impact the formation, internal structure, and distribution of subhalos within the MW (e.g., see Chapter 7 [332]). This situation is complicated by baryonic physics, which also impacts the survival and internal structure of subhalos orbiting the MW, largely due to the tidal effects of the Galactic disk (e.g., [144, 190]). It is therefore crucial to understand how different aspects of the MW subhalo population respond to both dark matter physics and baryons and to identify summary statistics that break degeneracies between these scenarios.

To explore these questions, we study the subhalo populations in simulations of MW-like systems that include realistic Large Magellanic Cloud (LMC) analogs. We present new SIDM re-simulations of these MW-like systems, and we compare the subhalo populations to those in a warm dark matter (WDM) model and in the presence of a Galactic disk using an analytic model calibrated to hydrodynamic simulations. We describe how the velocity function, radial distribution, and spatial distribution of MW subhalos can be used to disentangle SIDM, WDM, and disk suppression based on the distinct physical mechanisms that impact the formation or survival of subhalos in each scenario. Our results demonstrate that current observations of the MW satellite galaxy population, which probe the velocity function and radial distribution of low-mass subhalos in the MW, can already begin to distinguish dark matter models from both each other and baryonic effects.

Because of its recent, high-velocity infall into the MW, the LMC system provides a unique testing ground for dark matter physics. The LMC is accompanied by several ultra-faint dwarf

Model	Mass-dependent suppression	Radially-dependent suppression	Anisotropic suppression	$f_{\text{sub,LMC}}$	Physical driver
SIDM	✓	✓	✓	$48 \pm 13\%$	Ram-pressure + tidal stripping
WDM	✓	X*	X	$58 \pm 11\%$	Free-streaming
Disk	X	✓	X	$60 \pm 10\%$	Tidal disruption

Table 10.1: Impact of SIDM, WDM, and the Galactic disk on the MW subhalo population. The first column lists the MW subhalo population model. The second and third columns describe the suppression of the subhalo velocity function and radial distribution. The fourth column describes whether the suppression of the subhalo population is anisotropic. The fifth column lists the fraction of subhalos with $V_{\text{peak}} > 20 \text{ km s}^{-1}$ in the hemisphere that points from the MW halo center towards the LMC, averaged over our two MW-like simulations. The final column describes the main physical effect that impacts the MW subhalo population. The asterisk indicates that the shape of the WDM radial distribution may change for subhalos below our V_{peak} threshold.

galaxies [250, 356], , yielding an overabundance of satellites—and thus subhalos—in its vicinity, at a distance of only ~ 50 kpc [149, 333]. By using simulations tailored to match the LMC’s orbit and satellite population, our analysis accounts for effects specific to the MW–LMC system. We find that a significant fraction of the subhalos associated with the LMC are disrupted in our SIDM simulations due to self-interactions within the LMC halo at early times. Due to the LMC’s plunging, radial orbit, remaining LMC-associated subhalos are heavily stripped due to self-interactions with particles in the MW halo. This signature of SIDM physics erases the spatial anisotropy in the MW subhalo population caused by the LMC system in our CDM simulations, unlike WDM and disk disruption scenarios. This phenomenon is also relevant for interpreting rapidly advancing observations of satellite systems beyond the MW (e.g., [193, 316, 104]) in the context of SIDM.

This paper is organized as follows. In Section 10.3, we describe our MW-like CDM simulations, our new SIDM re-simulations of these systems, and our WDM and disk disruption models. We describe how each scenario impacts the MW subhalo population in Section 10.4. We discuss our results and conclude in Section 10.5. Throughout, we define virial quantities using the [93] critical overdensity $\Delta_{\text{vir}} \simeq 99.2$, set by the cosmological parameters in our simulations: $h = 0.7$, $\Omega_{\text{m}} = 0.286$, $\Omega_{\text{b}} = 0.047$, and $\Omega_{\Lambda} = 0.714$ [315].

10.3 Methods

10.3.1 MW-like CDM Simulations

Our analysis is based on two CDM dark matter–only zoom-in simulations of MW-like halos originally presented in [315]. The highest-resolution particles in these simulations have a mass of $3 \times 10^5 \text{ M}_{\odot} h^{-1}$, and the Plummer-equivalent softening length in the highest-resolution regions is 170 pc h^{-1} . Subhalos are well resolved down to a peak maximum circular velocity of $V_{\text{peak}} \approx$

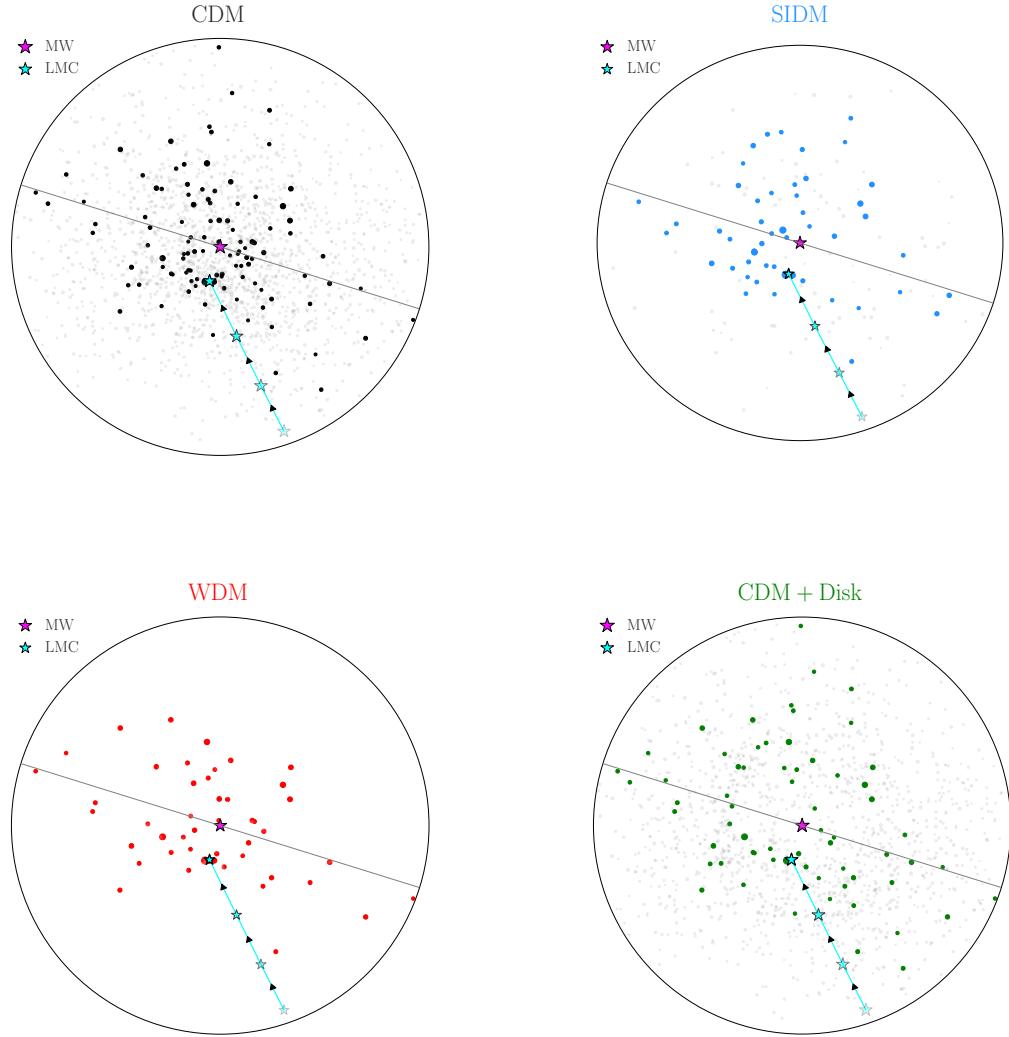


Figure 10.1: Projections of the subhalo population in one of our MW-like simulations in CDM (top left) and SIDM (top right), and for subhalos above a survival probability threshold corresponding to the suppression of the peak velocity function in our WDM (bottom left) and disk disruption (bottom right) models applied to the same CDM simulation. Bold markers show subhalos with $V_{\text{peak}} > 20 \text{ km s}^{-1}$, light gray markers show subhalos below this V_{peak} threshold, and marker size is proportional to V_{peak} . In each panel, the black circle shows the virial radius of the MW host halo ($\sim 300 \text{ kpc}$), the diagonal line shows the plane defined by the current position of the MW (magenta star) and LMC (cyan star), and transparent cyan stars show the LMC at previous snapshots as it falls into the MW.

10 km s^{-1} , where V_{peak} is the largest maximum circular velocity a halo attains over its entire history [315]. Halo catalogs and merger trees were generated using the ROCKSTAR halo finder and the CONSISTENT-TREES merger tree code [47, 48].

These two simulations include both *Gaia*-Enceladus-like mergers at early times and realistic LMC analog systems that fall into the MW within the last ~ 2 Gyr on orbits consistent with LMC proper motion measurements [250]. The subhalo populations in these simulations associated with observable satellite galaxies are consistent with current Dark Energy Survey and Pan-STARRS1 measurements of the MW satellite population down to a peak virial subhalo mass of $\sim 10^8 \text{ M}_\odot$. As visualized in the top-left panel of Figure 10.1, there is an overabundance of subhalos in the vicinity of the LMC, consistent with the number of LMC-associated satellites inferred from *Gaia* proper motion measurements [333].

10.3.2 SIDM Simulations

We re-run both simulations described above in the presence of dark matter self-interactions using the modified version of GADGET-2 presented in [39]. As emphasized in Chapter 9 [334], ROCKSTAR is not optimized for halo finding in SIDM simulations. To avoid difficulties introduced by halo finder incompleteness, we only study subhalos with $V_{\text{peak}} > 20 \text{ km s}^{-1}$, which are resolved with more than ~ 1000 particles at the time V_{peak} is achieved. In addition, this threshold roughly corresponds to the minimum V_{peak} associated with observed MW satellite galaxies based on abundance matching [333]. Thus, our results directly inform the interpretation of the subhalo population inferred from MW satellite observations.

Chapter 9 [334] reported significant levels of subhalo disruption in SIDM simulations of a MW-mass halo for SIDM models with cross sections of $\sim 1 \text{ cm}^2 \text{ g}^{-1}$. The efficiency of subhalo disruption is most sensitive to the self-interaction cross section at the orbital velocity scale, which is set by the potential of the MW host halo ($\sim 200 \text{ km s}^{-1}$). We therefore study an SIDM model with a total cross section of $0.1 \text{ cm}^2 \text{ g}^{-1}$ that is velocity independent on the scales of interest for our study, and we expect our main results to generalize to velocity-dependent models with larger cross sections at lower velocities. Because there are potentially significant quantitative differences between the subhalo populations predicted by various SIDM N -body codes (e.g., see the discussion in Chapter 9 [334]), we only focus on the qualitative trends imprinted by SIDM on the MW subhalo population. The top-right panel of Figure 10.1 shows the subhalo population in one of our SIDM simulations, which is clearly anisotropically suppressed relative to CDM, particularly at low subhalo masses.

10.3.3 WDM Model

We study the impact of thermal relic WDM on the subhalo populations in our MW-like simulations analytically. In particular, we assign a weight to each subhalo in our CDM simulations based on the probability it survives in a corresponding WDM simulation. We estimate this probability using a

fitting function that describes the suppression of the WDM subhalo mass function in simulations of MW-mass systems. Several fitting functions have been reported based on different WDM simulations [415, 26, 28, 299, 80]. Although these results vary, particularly below the cutoff scale [28], they are in reasonable agreement for the WDM particle masses and subhalo masses relevant for our study.

We express the suppression of the WDM subhalo mass function relative to that in CDM as

$$\frac{dN_{\text{WDM}}}{dM} \equiv f_{\text{WDM}}(M, M_{\text{hm}}) \frac{dN_{\text{CDM}}}{dM}, \quad (10.1)$$

where dN_{WDM}/dM (dN_{CDM}/dM) is the WDM (CDM) subhalo mass function. Here, f_{WDM} is a suppression factor that depends on peak subhalo mass virial mass M and WDM particle mass m_{WDM} via the half-mode mass [332]

$$M_{\text{hm}}(m_{\text{WDM}}) = 5 \times 10^8 \left(\frac{m_{\text{WDM}}}{3 \text{ keV}} \right)^{-10/3} \text{ M}_\odot. \quad (10.2)$$

We adopt the [299] suppression

$$f_{\text{WDM}}(M, m_{\text{WDM}}) = \left[1 + \left(\frac{\alpha M_{\text{hm}}(m_{\text{WDM}})}{M} \right)^\beta \right]^\gamma, \quad (10.3)$$

where $\alpha = 2.7$, $\beta = 1.0$, and $\gamma = -0.99$.

When evaluating subhalo population statistics below, we assign each subhalo a weight equal to f_{WDM} . We choose m_{WDM} such that the abundance of WDM subhalos predicted by Equation 10.3 matches that in our SIDM simulations for $V_{\text{peak}} > 20 \text{ km s}^{-1}$. This yields $m_{\text{WDM}} = 2.8 \text{ keV}$. This WDM particle mass is ruled out by a variety of small-scale structure probes (e.g., [459, 232, 195, 226, 341, 332]), but we emphasize that there are large theoretical uncertainties underlying the detailed relation between SIDM cross section and subhalo disruption [334]. We therefore leave a detailed investigation of SIDM constraints to future work.

The procedure outlined above assumes that the impact of WDM on the MW subhalo population can be expressed purely as a function of peak subhalo mass. Thus, our WDM model does not explicitly alter additional properties of the MW subhalo population, including its radial distribution (e.g., see [298]). Because subhalos of different masses above our V_{peak} threshold do not exhibit significantly different radial distributions in our CDM simulation, our predicted WDM and CDM subhalo populations therefore have similar radial profiles, consistent with the results of WDM simulations for the subhalo masses considered in this study (e.g., [26, 299, 80, 298]). The bottom-left panel of Figure 10.1 shows that WDM significantly suppresses the abundance of low-mass subhalos in our simulations relative to CDM.

10.3.4 Disk Disruption Model

In analogy to our WDM procedure, we model the impact of baryonic physics on the MW subhalo population by assigning a weight to each subhalo based on the probability it survives in a corresponding simulation with baryons. To estimate this probability, we adopt the fitting function from [404], which is derived using a suite of zoom-in simulations of MW-mass halos from the Feedback in Realistic Environments (FIRE) project.¹ The main physical driver of subhalo disruption in these hydrodynamic simulations is tidal disruption due to the central galaxy’s disk; this mechanism most strongly impacts subhalos with close pericentric passages (e.g., [190, 337, 383, 472]).

[404] find that subhalo disruption in hydrodynamic simulations can be described as a function of subhalo Galactocentric radius,

$$\frac{dN_{\text{disk}}}{dM} \equiv f_{\text{disk}}(r) \frac{dN_{\text{CDM}}}{dM}, \quad (10.4)$$

where

$$f_{\text{disk}}(r) = \begin{cases} 0, & 0 \leq r < r_0, \\ a \left[1 - \exp \left(-\frac{r-r_0}{r_1} \right) \right], & r \geq r_0. \end{cases} \quad (10.5)$$

We use the best-fit parameters from [404] for the differential suppression of subhalos with peak virial mass $M > 8 \times 10^8 M_{\odot}$, which are most relevant for our results. This corresponds to $a = 0.8$, $r_0 = 8$ kpc, and $r_1 = 78$ kpc in Equation ???. Following our WDM treatment, we assign each subhalo a weight of f_{disk} when evaluating subhalo population summary statistics.

Several other algorithms and fitting functions have been proposed to model the impact of baryons on subhalo populations in MW-mass systems. The prescription above predicts subhalo populations that agree reasonably well with those predicted by the Chapter 2 [337] model, with slightly more severe disruption at small Galactocentric radii. These predictions are also in broad agreement with other embedded disk potential (e.g., [190, 259, 391]) and hydrodynamic (e.g., [501, 383, 472]) simulations. We therefore adopt Equation 10.5 as our fiducial disk disruption model. The bottom-right panel of Figure 10.1 illustrates the suppression of subhalos in the inner regions of the MW caused by the Galactic disk.

10.4 Results

10.4.1 Subhalo Velocity Function

The left panel of Figure 10.2 shows the number of surviving subhalos as a function of V_{peak} in each model.² Disk disruption results in a roughly mass-independent suppression of subhalo abundance

¹<https://fire.northwestern.edu/>

²Our results are qualitatively similar when expressed in terms of subhalo mass functions (e.g., using peak virial mass) rather than velocity functions.

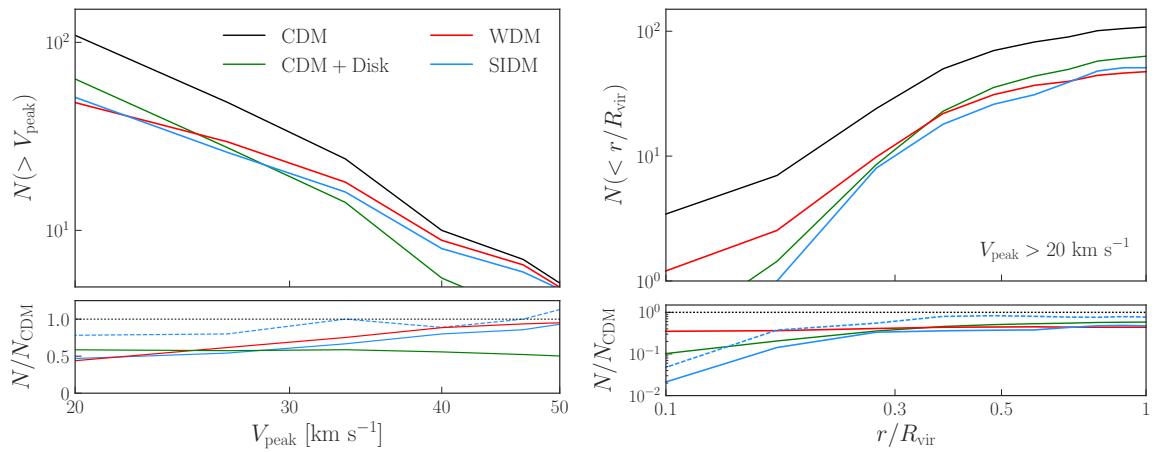


Figure 10.2: The effects of SIDM, WDM, and the Galactic disk on the velocity function and radial distribution of surviving subhalos in one of our MW-like simulations. Left panel: number of subhalos as a function of peak maximum circular velocity V_{peak} in our CDM simulations of MW-like systems (black dashed), SIDM simulations (blue), and predicted by applying our WDM (red) and disk disruption (green) models to our CDM simulations. Disruption due to the Galactic disk is approximately mass-independent, while SIDM and WDM preferentially disrupt less massive subhalos. Right panel: same as the left panel, but for the cumulative radial distribution of subhalos in units of the MW host halo virial radius. The suppression of subhalo abundance in WDM is not a strong function of Galactocentric radius, while SIDM and the disk preferentially disrupt subhalos at small radii. In the bottom panels, the dashed blue line shows the suppression of the subhalo population for the same SIDM model in a MW-mass system without a realistic LMC analog.

relative to CDM. This is due to the fact that subhalos in our disk disruption model are only suppressed based on their Galactocentric radius, which is not significantly correlated with subhalo mass in our CDM simulations. The subhalo V_{peak} function in our disk disruption model therefore retains an approximate power-law form with a lower normalization than in CDM.

Meanwhile, WDM and SIDM both preferentially reduce the abundance of low-mass subhalos relative to CDM. In WDM, this results from free-streaming effects that suppress the linear matter power spectrum on small scales and inhibit the formation of low-mass halos (e.g., [416]). In SIDM, this mass-dependent suppression is caused by enhanced tidal stripping in the presence of ram-pressure interactions between subhalos and the MW host halo, as described in Chapter 9 [334]. Thus, unlike for WDM, the reduction in subhalo abundance in our SIDM simulations does not necessarily imply a corresponding reduction in the abundance of satellite galaxies.

While we match the amplitude of WDM and SIDM subhalo abundances by construction, the detailed agreement between the corresponding V_{peak} functions shown in Figure 10.2 is coincidental. In particular, these V_{peak} functions deviate for $V_{\text{peak}} < 20 \text{ km s}^{-1}$ or for a different choice of the V_{peak} threshold at which subhalo abundances are matched. This is expected because low-mass subhalos are *dynamically* disrupted in SIDM, while their *formation* is suppressed in WDM due to the effects of free-streaming on the linear matter power spectrum. Nonetheless, our qualitative result—i.e., that the abundance of lower-mass subhalos is preferentially suppressed in WDM and SIDM, unlike in a disk disruption scenario—remains robust.

Interestingly, the suppression of the subhalo velocity function for the same SIDM model run on a MW-mass system without a realistic LMC analog (shown by the dashed blue line in the bottom panel of Figure 10.2) is less severe in terms of both amplitude and mass dependence. The difference between this alternate simulation and our fiducial results is consistent with the enhanced disruption of LMC subhalos in our MW-like simulations discussed in Section 10.4.3.

Thus, the Galactic disk suppresses the abundance of subhalos relative to CDM in a mass-independent manner, while the abundance of lower-mass subhalos is preferentially suppressed in WDM and SIDM.

10.4.2 Subhalo Radial Distribution

The right panel of Figure 10.2 shows the radial distribution of surviving subhalos in units of the MW host halo virial radius. We observe the following trends:

1. Zero subhalos with $V_{\text{peak}} > 20 \text{ km s}^{-1}$ survive within a Galactocentric distance of $\sim 30 \text{ kpc}$ in our fiducial disk disruption model;
2. SIDM subhalo abundances are suppressed relative to CDM at small Galactocentric radii;
3. The suppression of subhalo abundance in WDM is roughly independent of Galactocentric radius.

These results are consistent with simulations that include disk disruption (e.g., [190, 383, 472]) and SIDM physics (e.g., [391, 334, 435]). We emphasize that—in both the disk disruption and SIDM cases—the radius within which subhalos are most strongly disrupted is model and resolution-dependent. For example, [205] show that the abundance of surviving subhalos predicted by cosmological simulations within $\sim 0.1 R_{\text{vir}}$ is strongly resolution-dependent due to both stripping below the mass resolution limit and artificial disruption. These numerical uncertainties may be exacerbated in SIDM simulations due to imperfect halo finding [334].

Thus, disk disruption and SIDM preferentially suppress the abundance of subhalos in the inner regions of the MW, while WDM-like models that inhibit the formation of low-mass halos suppress the abundance of subhalos that are expected to host dwarf galaxies in a radially-independent manner.

10.4.3 Spatial Anisotropy in the MW Subhalo Population

We are uniquely poised to study the spatial distribution of MW subhalos because our simulations include realistic LMC analog systems. Averaged over our two CDM simulations, $60 \pm 10\%$ of the total MW subhalo population with $V_{\text{peak}} > 20 \text{ km s}^{-1}$ (including higher-order subhalos) is found in the hemisphere pointing towards the LMC from the host halo center, where uncertainties indicate 1σ Poisson scatter. This anisotropy results from the population of low-mass subhalos that fall into the MW with the LMC; we refer to subhalos that were within the LMC’s virial radius at the time of LMC infall as “LMC-associated.” In particular, $\sim 80\%$ of the LMC-associated subhalos above our V_{peak} threshold are found in the LMC hemisphere at $z = 0$, contributing $\sim 33\%$ to the total MW subhalo population in this direction. Recent discoveries and proper motion measurements of satellite galaxies in the Southern hemisphere indicate similar anisotropy and LMC association statistics in the data (e.g., [149, 356]).

The spatial anisotropy of the CDM subhalo distribution responds differently to various mechanisms that suppress the formation or survival of low-mass subhalos. In the presence of baryons, “pre-processing” in the LMC halo may reduce subhalo abundances by up to $\sim 30\%$, which is significantly lower than the $\sim 50\%$ disruption effect due to the MW disk [238]. Because the efficiency of subhalo disruption due to the LMC galaxy is highly uncertain, our disk disruption model only accounts for the impact of the MW disk on subhalo abundances. This model does not affect anisotropy statistics compared to our CDM simulations (although Poisson uncertainties increase due to the overall decrement in subhalo abundance) because the radial distribution of LMC-associated subhalos is approximately unbiased relative to the rest of the MW subhalo population.

The spatial anisotropy introduced by the LMC system is slightly reduced in WDM but remains consistent with our CDM result within 1σ Poisson uncertainty. This is due to the fact that the mass distributions for MW and LMC-associated subhalos are similar. Thus, WDM does not preferentially suppress the formation of LMC-associated subhalos relative to other MW subhalos with similar masses.

SIDM has the most noticeable impact on the spatial distribution of surviving subhalos, *erasing* the anisotropy in our CDM simulations. In particular, $48 \pm 13\%$ of subhalos with $V_{\text{peak}} > 20 \text{ km s}^{-1}$ are found in the LMC hemisphere, averaged over our two SIDM simulations. Two effects reduce the anisotropy of the MW subhalo population in SIDM:

1. Pre-processing in the LMC halo: self-interactions with LMC halo particles disrupt $\sim 50\%$ of LMC-associated subhalos above our V_{peak} threshold prior to LMC infall into the MW. These disrupted subhalos typically have orbital velocities above $\sim 200 \text{ km s}^{-1}$; thus, this mechanism mainly depends on the SIDM cross section at the same velocity scale that influences disruption within the MW halo.
2. Ram-pressure stripping and tidal disruption by the MW halo: our LMC analogs are on plunging radial orbits, which enhances ram-pressure stripping due to self-interactions and makes LMC-associated subhalos more susceptible to tidal disruption.³

Both of these effects depend on the specific history of the MW–LMC system; it is therefore crucial to study the impact of dark matter physics on subhalo populations in realistic settings constrained by the accretion history of the MW and the orbital properties of its most massive satellites. Although the anisotropy of the MW subhalo population is only affected at the $\sim 1\sigma$ level for subhalos with $V_{\text{peak}} > 20 \text{ km s}^{-1}$, Poisson uncertainties shrink quickly as this lower limit is decreased. Thus, probes of the MW subhalo population below the galaxy formation threshold—including the statistics and density profiles of stellar streams over the full sky—hold valuable information about SIDM physics.

Thus, the overabundance of subhalos in the direction of the LMC expected in CDM is largely unaffected by the Galactic disk and WDM, while SIDM can erase this anisotropy due to self-interactions within the LMC halo at early times and during its infall into the MW.

10.5 Discussion and Conclusion

We have demonstrated that the velocity function, radial distribution, and spatial anisotropy of the MW subhalo population respond in distinct ways to different kinds of dark matter and baryonic physics. These properties of the MW subhalo population can be inferred by forward modeling MW satellite galaxies and by studying the abundance, spatial distribution, and density perturbation profiles of stellar streams. Measurements of both probes are expected to rapidly progress in the coming years [151, 43].

Although we restricted our study to subhalos with $V_{\text{peak}} > 20 \text{ km s}^{-1}$, which MW satellite observations currently probe, many of our qualitative results hold for lower-mass subhalos probed by other tracers of surviving and disrupted substructure in the MW. For example, inferring the

³The same mechanism disrupts subhalos associated with the *Gaia*–Enceladus analogs in our simulations as they approach pericenter, exacerbating the effect discussed in [152].

population of low-mass subhalos that perturb stellar streams is a promising avenue for studying dark matter physics (e.g., [76, 41, 120]), and our results inform predictions for subhalo–stream encounters. For example, we predict that the relative rate of subhalo–stream encounters towards and away from the LMC is sensitive to SIDM physics. The velocity function, radial distribution, and spatial anisotropy trends we report also impact the interpretation of direct and indirect detection results that rely on estimates of the nearby subhalo population.

Additional properties of the MW subhalo population are altered in each scenario we consider. For example, surviving subhalos in our SIDM simulations can be “cored” by self-interactions and often have lower present-day maximum circular velocities than their counterparts in other models [334]. This effect is expected to be less prevalent in the WDM and disk disruption models we consider because disruption due to the Galactic disk is impulsive [190], while WDM free-streaming does not produce significant cores for the WDM particle masses and subhalos of interest (e.g., [312]).

We considered SIDM, WDM, and disk disruption separately because the efficiency of subhalo disruption due to the Galactic disk is not known in detail. For example, the [333] MW satellite population analysis finds that as few as $\sim 30\%$ (and up to $\sim 90\%$) of subhalos can survive disk disruption while remaining compatible with MW satellite observations. These constraints on disruption efficiency also depend on uncertainties in the MW halo mass (e.g., [341]). Combining SIDM or WDM physics with the effects of the disk may introduce degeneracies (e.g., between the SIDM and disk disruption radial distributions); thus, it is necessary to thoroughly understand the impact of each scenario separately.

Finally, we note that [391] performed a similar study, and we highlight several synergies and differences between this analysis and our work. [391] studied the subhalo V_{\max} function and radial distribution in CDM, SIDM, and disk disruption scenarios, finding that CDM and SIDM subhalo populations are similar in terms of their V_{\max} functions and radial distributions in the presence of an embedded disk potential. The radial distributions in our SIDM simulations are consistent with these findings, while our subhalo velocity functions are more suppressed relative to CDM at low masses. We emphasize that SIDM simulation codes may not be converged even with respect to the *abundance* of surviving subhalos (e.g., see the discussion in Chapter 9 [334]). Exploring these potential discrepancies in a dedicated code comparison study is essential before using the MW subhalo population to constrain the SIDM cross section.

Our results indicate that the impact of dark matter physics on the MW subhalo population can depend on the specific history of the MW–LMC system. Future work that expands the available suite of MW-like simulations ([94] in prep.) and explores modifications to initial conditions to engineer appropriate mass accretion histories (e.g., [396, 437]) is therefore compelling.

Chapter 11

Combining Milky Way Satellites and Strong Lensing

Abstract

Joint analyses of small-scale cosmological structure probes are relatively unexplored and promise to advance measurements of microphysical dark matter properties using heterogeneous data. Here, we present a multidimensional analysis of dark matter substructure using strong gravitational lenses and the Milky Way (MW) satellite galaxy population, accounting for degeneracies in model predictions and using covariances in the constraining power of these individual probes for the first time. We simultaneously infer the projected subhalo number density and the half-mode mass describing the suppression of the subhalo mass function in thermal relic warm dark matter (WDM), M_{hm} , using the semianalytic model `Galacticus` to connect the subhalo population inferred from MW satellite observations to the strong lensing host halo mass and redshift regime. Combining MW satellite and strong lensing posteriors in this parameter space yields $M_{\text{hm}} < 10^{7.0} M_{\odot}$ (WDM particle mass $m_{\text{WDM}} > 9.7$ keV) at 95% confidence and disfavors $M_{\text{hm}} = 10^{7.4} M_{\odot}$ ($m_{\text{WDM}} = 7.4$ keV) with a 20:1 marginal likelihood ratio, improving limits on m_{WDM} set by the two methods independently by $\sim 30\%$. These results are marginalized over the line-of-sight contribution to the strong lensing signal, the mass of the MW host halo, and the efficiency of subhalo disruption due to baryons and are robust to differences in the disruption efficiency between the MW and strong lensing regimes at the $\sim 10\%$ level. This work paves the way for unified analyses of next-generation small-scale structure measurements covering a wide range of scales and redshifts.

11.1 Paper Status and External Contributions

This chapter is in press in the *Astrophysical Journal* with the title, “Dark Matter Constraints from a Unified Analysis of Strong Gravitational Lenses and Milky Way Satellite Galaxies,” on which I am the corresponding author. It is the result of a collaboration with Simon Birrer, Daniel Gilman, Risa Wechsler, Xiaolong Du, Andrew Benson, Anna Nierenberg, and Tommaso Treu. Simon and Daniel led the strong lensing analysis that we combine with Milky Way satellite galaxies in this chapter; along with Risa, they worked with me to interface the analyses and interpret the joint constraints. Xiaolong and Andrew provided the `Galacticus` output used to facilitate this comparison, and all of the coauthors made editorial contributions to the text.

11.2 Introduction

The Λ CDM cosmological paradigm assumes a cold, collisionless dark matter (CDM) particle and therefore predicts a plethora of dark matter structure and substructure on extremely small cosmic scales (e.g., [204, 137, 466]). It is often argued that small-scale structure measurements represent an outstanding test to this prediction (e.g., see [97] for a review); yet, our understanding of the distribution of dark matter structure on nonlinear scales is rapidly progressing. Recent analyses of Milky Way (MW) satellite galaxies using data over nearly the full sky—including the population of ultrafaint dwarf galaxies discovered by deep photometric surveys over the last decade—have only recently been performed (e.g., [149, 333, 332]). Meanwhile, measurements of stellar streams from the *Gaia* mission are beginning to reach the requisite precision to infer the signatures of perturbations from nearby low-mass subhalos [41, 76]. On extragalactic scales, the number of compact-source strong gravitational lenses available for substructure analyses has drastically increased in recent years (e.g., [344, 343, 342]), and modeling efforts have advanced in step (e.g., [196, 195, 226]). Analyses of resolved distortion in extended strong lensing observations from adaptive optics and space-based imaging have also rapidly progressed (e.g., [217, 62, 455]).

All of the recent small-scale structure measurements outlined above are consistent with the CDM paradigm and have therefore been used to constrain microphysical properties of dark matter that would reduce its small-scale clustering [41, 195, 332]. Although analyses of different probes reach consistent dark matter constraints, to date they have been performed independently and with different modeling assumptions to address heterogeneous astrophysical systematics. Crucially, if evidence for a departure from the CDM paradigm arises, it must be confirmed across different redshifts and physical scales. It is therefore critical to jointly model and analyze small-scale structure probes. This effort will be particularly important to maximize small-scale structure measurements from next-generation surveys including the Rubin Observatory Legacy Survey of Space and Time (LSST), which will enable the discovery of vastly more strong gravitational lenses (e.g., [114]) and revolutionize the search for dwarf galaxies and measurements of stellar streams in the local universe

(e.g., [151]).

Jointly modeling the low-mass halo and subhalo populations relevant for various small-scale structure measurements requires precise theoretical predictions for the abundance and structure of these small systems—which probe highly nonlinear cosmological modes—as a function of redshift and environment. Even cosmological parameters play an important role given the precision of current data; for example, varying the running of the spectral index within Planck uncertainties significantly affects predictions for subhalo abundances [435], while other cosmological parameters including Ω_m and σ_8 have subleading effects that may become important to incorporate in models of next-generation small-scale structure data [146]. Moreover, a variety of other theoretical and numerical uncertainties must be marginalized over in joint likelihood analyses to robustly claim evidence for non-CDM physics. For example, specific systematics of interest for modeling the MW satellite galaxy population include the faint end of the galaxy–halo connection, the total mass of the MW halo, and the mass and accretion time of the Large Magellanic Cloud (LMC; [341, 332]). Meanwhile, the orbits of dark matter subhalos in the inner regions of the MW halo must be predicted precisely in a statistical sense while accurately modeling the effects of specific baryonic structures to infer dark matter properties from stellar stream measurements. For strong lensing, the mass–concentration relation in both CDM and alternative dark matter models is a key uncertainty that must be accounted for [217, 195, 197, 323], along with the host halo properties and selection functions of strong lenses, all while accurately modeling the differential signal contributed by substructure and small halos along the line of sight [133, 196].

Here, we perform a joint analysis of small-scale dark matter measurements by combining the results of recent strong gravitational lensing and MW satellite inferences in a multidimensional parameter space to break modeling degeneracies.¹ In particular, we combine these results in a parameter space that includes the mass scale describing the suppression of the subhalo mass function for thermal relic warm dark matter (WDM) and the amplitude of the projected subhalo mass function at the strong lensing host halo mass and redshift scale. In particular, we combine the constraints on these quantities derived from (i) the magnification and flux ratio data from quadruply imaged strong gravitational lenses presented in [344, 343, 342] and modeled in [195], and (ii) the abundance and properties of MW satellite galaxies over $\sim 75\%$ of the sky presented in [149] and modeled in Chapters 4–7 [333, 332]. We employ the semianalytic model `Galacticus` [56, 373] to translate the subhalo population inferred from MW satellite measurements to the strong lensing host halo mass and redshift scale by calibrating its predictions using cosmological zoom-in simulations of MW-like halos [315]. Thus, our work lays the foundation for joint semianalytic models of small-scale structure that are benchmarked by high-resolution simulations at each halo mass and redshift scale.

¹A joint small-scale structure analysis by [161] appeared during the preparation of this manuscript. We comment on the differences between the methodology and results of our study and this work in Section 11.8.2.

Our joint analysis breaks degeneracies among the amplitude of the projected subhalo mass functions inferred from MW satellite and strong lensing observations, thereby improving limits on deviations from the CDM paradigm that have been derived independently from these data. Specifically, we show that our combined analysis improves the lower limit on the WDM particle mass by $\sim 30\%$. The framework we develop for combining MW satellite and strong lensing data is particularly important because strong lensing is potentially sensitive to the presence of halos with masses below the threshold for galaxy formation, a mass scale that dwarf galaxy observations constrain. We therefore quantify the observational and theoretical advances necessary to robustly infer the presence of such dark halos, showing that this outcome is within the reach of next-generation small-scale structure measurements.

This paper is organized as follows. In Section 11.3, we describe the analytic model of dark matter substructure that underlies our joint analysis. We then describe the MW satellite data and model that enters our analysis in Section 11.4 and the strong lensing data and model in Section 11.5. We combine these analyses in Section 11.6, present our results in Section 11.7, discuss key systematics and compare them to previous work in Section 11.8, and conclude in Section 11.9. Throughout, we adopt the following cosmological parameters, following both [195] and [332]: $h = 0.7$, $\Omega_m = 0.286$, $\Omega_\Lambda = 0.714$, $\sigma_8 = 0.82$, and $n_s = 0.96$ [218].

11.3 Dark Matter Substructure Model

We begin by describing the analytic model of dark matter substructure used to connect the subhalo populations probed by MW satellite and strong lensing observations. In particular, we describe our model for the projected subhalo mass function (SHMF; Section 11.3.1), its dependence on host halo mass and redshift (Section 11.3.2), and the efficiency of subhalo disruption due to baryons (Section 11.3.3). We then describe our model for the impact of WDM physics on subhalo abundances (Section 11.3.4) and concentrations (Section 11.3.5).

11.3.1 Projected subhalo mass function

Strong lensing and MW satellites probe low-mass subhalos within host halos at different mass and redshift scales. Specifically, strong lensing probes both the projected dark matter substructure in the lens system and small-scale structure along the line of sight, while MW satellites probe the three-dimensional distribution of subhalos traced by luminous satellite galaxies within the MW. Because current strong lensing analyses are not highly sensitive to the radial distribution of subhalos within the host halo of the lens, we focus on the statistics of projected subhalo populations in this paper, although we will describe how observations of the radial distribution of MW satellites break model degeneracies.

To simultaneously predict the subhalo populations relevant for MW satellite and strong lensing

studies, we construct an analytic model for projected subhalo abundances that depends on the host halo mass, M_{host} , and redshift, z_{host} . In particular, we express the projected SHMF—i.e., the differential number of subhalos within a host halo, in projection—by generalizing the form in [195],

$$\frac{d^2N_{\text{sub}}}{dM dA} \equiv \frac{\Sigma_{\text{sub}}(M_{\text{host}}, z_{\text{host}})}{M_0} \left(\frac{M}{M_0} \right)^\alpha \mathcal{F}_{\text{CDM}}(M_{\text{host}}, z_{\text{host}}), \quad (11.1)$$

where M denotes subhalo mass, A denotes the unit area, $M_0 = 10^8 M_\odot$ and α is the power-law slope of the SHMF. In Equation 11.1, $\Sigma_{\text{sub}}(M_{\text{host}}, z_{\text{host}})$ is the projected number density of subhalos within the virial radius of a host halo of mass M_{host} at redshift z_{host} , including the effects of baryonic physics, and $\mathcal{F}_{\text{CDM}}(M_{\text{host}}, z_{\text{host}})$ captures the dependence of the projected SHMF on M_{host} and z_{host} in CDM only (i.e., modulo the effects of baryonic physics) as described in Section 11.3.2. We discuss the impact of halo boundary definitions in Section 11.6.

Our model of the MW satellite population is based on the Chapter 7 [332] analysis, which defines subhalo mass using the peak Bryan–Norman virial mass M_{peak} (see Appendix F.1 for details). Meanwhile, our strong lensing constraints are based on the [195] analysis, which uses M_{200} values relative to the critical density at $z = 0$, with subhalo masses evaluated at infall to compute the WDM SHMF and mass–concentration relation. Here, we simply interpret the peak virial mass values from our MW satellite analysis as M_{200} values at infall. The peak virial masses of subhalos relevant for our work are on average a factor of ~ 2 larger than their M_{200} values at infall in the cosmological zoom-in simulations our MW satellite analysis is based on, largely due to pre-infall tidal stripping (e.g., [46, 480]). Thus, converting M_{peak} to M_{200} would further strengthen our joint WDM constraints; however, because changing this mass definition would nontrivially affect the abundance-matching model used in our satellite analysis, we leave a detailed investigation of this point to future work that combines satellite and lensing inferences at the likelihood level.

11.3.2 CDM host mass and redshift dependence

We model the dependence of the projected SHMF on host halo mass and redshift with the functions $\Sigma_{\text{sub}}(M_{\text{host}}, z_{\text{host}})$ and $\mathcal{F}_{\text{CDM}}(M_{\text{host}}, z_{\text{host}})$. Both of these terms play a key role in our joint analysis because they allow us to relate the subhalo populations corresponding to low-redshift, group-mass strong lens host halos ($M_{\text{lens}} \sim 10^{13} M_\odot$, $z_{\text{lens}} \sim 0.5$) to the regime of the MW halo today ($M_{\text{MW}} \sim 10^{12} M_\odot$, $z_{\text{MW}} = 0$).

$\mathcal{F}_{\text{CDM}}(M_{\text{host}}, z_{\text{host}})$ captures the dependence of the SHMF on host halo mass and redshift in CDM only (i.e., without baryons), including the effects of tidal disruption by the dark matter host halo. This scaling depends on both the statistics of subhalo populations at infall, which can be predicted reasonably precisely using extensions of the Press–Schechter formalism [371], and on the dynamical evolution of subhalos after infall into a host. Detailed semianalytic models calibrated to N -body simulations are necessary to model this evolution; we therefore follow [195] in using the

`Galacticus` model [56, 373, 490], which predicts

$$\log \mathcal{F}_{\text{CDM}}(M_{\text{host}}, z_{\text{host}}) = k_1 \log \left(\frac{M_{\text{host}}}{10^{13} M_{\odot}} \right) + k_2 \log(z_{\text{host}} + 0.5), \quad (11.2)$$

where $k_1 = 0.88$ and $k_2 = 1.7$.

Section 11.3.3 describes our model for subhalo disruption due to baryons, which captures the leading-order corrections to $\mathcal{F}_{\text{CDM}}(M_{\text{host}}, z_{\text{host}})$. We do not model the impact of additional host halo, central galaxy, and environmental variables on the projected SHMF, noting that this is an important area for future work that ongoing observational efforts like the Satellites Around Galactic Analogs (SAGA) survey are informing at the MW-mass scale [193, 316]. However, we emphasize that the Chapter 7 [332] MW satellite analysis our work is based on self-consistently uses simulations that are consistent with key secondary MW halo properties, including concentration, the existence of a realistic LMC analog system, and a formation history constrained by Gaia observations. Meanwhile, the [195] lensing analysis is not sensitive to host-to-host variation in Σ_{sub} beyond that modeled by $\mathcal{F}_{\text{CDM}}(M_{\text{host}}, z_{\text{host}})$ given the current number of strong lenses studied and the information available per lens.

11.3.3 Subhalo disruption efficiency Due to baryons

We model $\Sigma_{\text{sub}}(M_{\text{host}}, z_{\text{host}})$ with explicit host halo mass and redshift dependence to capture the impact of baryonic physics on the projected SHMF. This extra dependence relative to the CDM scaling is not captured by the $\mathcal{F}_{\text{CDM}}(M_{\text{host}}, z_{\text{host}})$ term predicted by `Galacticus`, although baryonic effects can be modeled in future `Galacticus` implementations. Although Σ_{sub} is not modeled with explicit host mass and redshift dependence in [195], we include this dependence here because the subhalo populations probed by strong lensing and MW satellites are subject to baryonic effects that potentially impact the two regimes differently. Of these effects, the most important is tidal disruption due to the central galaxy, which suppresses the SHMF at the $\sim 50\%$ level (e.g., [132, 190, 203, 259, 383, 472]).² Tidal disruption due to the central galaxy most strongly suppresses the abundance of subhalos in the inner regions of the host halo or (more precisely) subhalos that accrete early and have close pericentric passages [190, 337].

The projected SHMF is largely driven by the plethora of subhalos in the outer regions of the host halo, which mitigates the impact of uncertainties in the strength and radial dependence of these baryonic effects on our probe combination. Nonetheless, our joint analysis is sensitive to both the amplitude of and differences in the efficiency of subhalo disruption due to baryonic physics as a function of host halo mass and redshift. We measure $\Sigma_{\text{sub}}(M_{\text{host}}, z_{\text{host}})$ in units of its value for

²Supernova feedback within sufficiently massive subhalos can also reduce their inner densities (e.g., [201, 369, 378]) and accelerate disruption, but hydrodynamic simulations suggest that this process has a subleading effect on the SHMF compared to disruption by the central galaxy.

strong lens host halos, and we define the differential subhalo disruption efficiency due to baryons as

$$q \equiv \frac{\Sigma_{\text{sub}}(M_{\text{MW}}, z_{\text{MW}})}{\Sigma_{\text{sub}}(M_{\text{lens}}, z_{\text{lens}})} \equiv \frac{\Sigma_{\text{sub,MW}}}{\Sigma_{\text{sub}}}, \quad (11.3)$$

where Σ_{sub} hereafter denotes the projected subhalo number density for strong lens host halos, following [195], and $\Sigma_{\text{sub,MW}}$ denotes the same quantity for the present-day MW system. In Equation 11.3, q represents the efficiency of subhalo disruption due to baryonic physics in the MW at $z = 0$ in units of the efficiency of subhalo disruption due to baryonic physics in the group-mass halos and at the redshifts probed by strong lensing. Note that larger (smaller) values of q represent less efficient (more efficient) subhalo disruption in the MW relative to strong lenses and that differences in the radial dependence of subhalo disruption at these scales (which we do not model) do not affect our joint analysis of projected SHMFs.

Motivated by the results of hydrodynamic simulations, we assume that subhalo disruption due to baryonic physics results in a mass-independent rescaling of the MW and strong lens projected SHMFs. This allows us to respectively express the projected SHMFs probed by strong lensing and MW satellite observations as

$$\left(\frac{d^2 N_{\text{CDM}}}{dM dA} \right)_{\text{lensing}} = \frac{\Sigma_{\text{sub}}}{M_0} \left(\frac{M}{M_0} \right)^\alpha \mathcal{F}_{\text{CDM}}(M_{\text{lens}}, z_{\text{lens}}), \quad (11.4)$$

$$\begin{aligned} \left(\frac{d^2 N_{\text{CDM}}}{dM dA} \right)_{\text{MW}} &= \frac{\Sigma_{\text{sub,MW}}}{M_0} \left(\frac{M}{M_0} \right)^\alpha \mathcal{F}_{\text{CDM}}(M_{\text{MW}}, z_{\text{MW}}) \\ &= \frac{q \Sigma_{\text{sub}}}{M_0} \left(\frac{M}{M_0} \right)^\alpha \mathcal{F}_{\text{CDM}}(M_{\text{MW}}, z_{\text{MW}}). \end{aligned} \quad (11.5)$$

As noted above, strong lenses typically have halo masses of $M_{\text{lens}} \approx 10^{13} M_\odot$ and redshifts of $z_{\text{lens}} \approx 0.5$, and host massive elliptical galaxies [191, 36, 195]. In contrast, the MW has a halo mass of $M_{\text{MW}} \sim 10^{12} M_\odot$ (e.g., [102, 106]) at $z_{\text{MW}} = 0$, and is largely typical for a spiral galaxy of its stellar mass, although it has a relatively quiescent formation history (e.g., [65, 168]). Subhalo disruption due to the central galaxy in hydrodynamic simulations of MW-mass systems reduces the amplitude of the SHMF by $\sim 50\%$ relative to corresponding dark-matter-only simulations. This effect is roughly mass independent and is not a strong function of redshift at late times. Although hydrodynamic simulations of group-mass systems yield similar levels of SHMF suppression [172, 203, 384], this regime is less well studied. We therefore adopt $q = 1$ in our fiducial analysis—i.e., equally efficient subhalo disruption due to baryons in the MW and strong lens host halos—and we also test values within a range of $q \in [0.5, 2]$ when translating the projected SHMF amplitude inferred from MW satellites to the strong lens host halo regime. We emphasize that, in detail, subhalo disruption is expected to depend on the mass and formation history of the central galaxy along with host

halo mass and redshift, and its efficiency will therefore differ among strong lenses. Although our phenomenological model for differences in subhalo disruption due to baryonic physics is very simple, we will demonstrate that the corresponding uncertainties do not significantly impact our joint dark matter constraints.

11.3.4 Warm dark matter subhalo mass function

The half-mode mass, M_{hm} , represents a characteristic mass scale describing the suppression of the linear matter power spectrum due to non-CDM physics; in particular, it corresponds (in linear theory) to the wavenumber at which the ratio of the linear matter power spectrum drops to 25% of that in CDM (e.g., [335]). In the case of thermal relic WDM, free streaming suppresses the power spectrum on small scales, leading to a turnover in the halo and subhalo mass functions below M_{hm} , which in turn depends on the WDM particle mass, m_{WDM} (e.g., [415]). MW satellites constrain this suppression by tracing the abundance of low-mass halos, while the subhalos surrounding the main deflector in strong lenses affect image flux ratios.

The WDM SHMF can be expressed as

$$\frac{dN_{\text{WDM}}}{dM} \equiv f_{\text{WDM}}(M, M_{\text{hm}}) \frac{dN_{\text{CDM}}}{dM}, \quad (11.6)$$

where dN_{WDM}/dM (dN_{CDM}/dM) is the WDM (CDM) SHMF, and f_{WDM} is a multiplicative suppression factor that depends on subhalo mass M and the WDM particle mass m_{WDM} via M_{hm} . We follow both [195] and [332] by assuming that this SHMF suppression does not alter the (normalized) radial distribution of subhalos, consistent with the findings of WDM simulations (e.g., [299, 77]). Thus, the same multiplicative factor $f_{\text{WDM}}(M, M_{\text{hm}})$ dictates the suppression of the projected SHMFs in our model, i.e.,

$$\left(\frac{d^2N_{\text{WDM}}}{dM dA} \right)_{\text{lensing}} = f_{\text{WDM}}(M, M_{\text{hm}}) \left(\frac{d^2N_{\text{CDM}}}{dM dA} \right)_{\text{lensing}}, \quad (11.7)$$

$$\left(\frac{d^2N_{\text{WDM}}}{dM dA} \right)_{\text{MW}} = f_{\text{WDM}}(M, M_{\text{hm}}) \left(\frac{d^2N_{\text{CDM}}}{dM dA} \right)_{\text{MW}}. \quad (11.8)$$

For concreteness and to allow for an apples-to-apples comparison between lensing and MW satellite analyses, we focus on the case of thermal relic WDM, for which the SHMF can be expressed as [297]

$$f_{\text{WDM}}(M, M_{\text{hm}}) = \left[1 + \left(\frac{\alpha M_{\text{hm}}(m_{\text{WDM}})}{M} \right)^\beta \right]^\gamma, \quad (11.9)$$

where M_{hm} is related to m_{WDM} in our fiducial cosmology via [332]

$$M_{\text{hm}}(m_{\text{WDM}}) = 5 \times 10^8 \left(\frac{m_{\text{WDM}}}{3 \text{ keV}} \right)^{-10/3} M_\odot. \quad (11.10)$$

In Equation 11.9, α , β , and γ are free parameters fit to simulation results. The analysis in Chapter 7 [332] uses the SHMF from [299], which corresponds to $\alpha = 2.7$, $\beta = 1.0$, and $\gamma = -0.99$, while [195] adopt an alternative fit to the SHMF from [299], corresponding to $\alpha = 1$, $\beta = 1$, and $\gamma = -1.3$. As described in Section 11.4, we rerun the MW satellite analysis with the [195] choice of WDM SHMF suppression in order to self-consistently combine the posterior distributions from these analyses according to the procedure in Section 11.6.3.

11.3.5 Warm dark matter mass–concentration relation

The delay in the collapse of small-scale density perturbations in WDM suppresses the central densities of halos with masses near M_{hm} , altering the mass–concentration relation for both field and subhalos. Because flux ratios in strong lenses are highly sensitive to the central densities of subhalos, the altered mass–concentration relation provides crucial information relevant for forward-modeling strong lensing signals [197]. We implement the WDM mass–concentration relation in a similar manner to the suppression of the SHMF [195],

$$c_{\text{WDM}}(M) \equiv f'_{\text{WDM}}(M, M_{\text{hm}}) c_{\text{CDM}}(M), \quad (11.11)$$

where $c_{\text{WDM}}(M)$ ($c_{\text{CDM}}(M)$) is the WDM (CDM) mass–concentration relation, and f'_{WDM} is a concentration suppression factor analogous to f_{WDM} . In particular, we follow [195] by using $c_{\text{CDM}}(M)$ from [138] with 0.1 dex scatter [155] and $c_{\text{WDM}}(M)$ from [77],

$$f'_{\text{WDM}}(M, M_{\text{hm}}) = (1+z)^{\beta(z)} \left(1 + 60 \frac{M_{\text{hm}}}{M}\right)^{-0.17}, \quad (11.12)$$

where $\beta(z) = 0.026z - 0.04$.

Halo concentrations are affected over an order of magnitude in mass above the turnover in the mass function set by M_{hm} . Thus, the mass–concentration relation must be accounted for to self-consistently constrain WDM-like models using strong lensing data. Meanwhile, MW satellite abundances are relatively insensitive to the mass–concentration relation because subhalo disruption is mainly determined by subhalos’ orbital properties. Moving beyond abundances, the internal dynamics of relatively bright MW satellite galaxies are often subject to baryonic effects that make it difficult to robustly infer halo concentration (e.g., [378]). Meanwhile, it is difficult to obtain precise dynamical measurements given the limited number of spectroscopically confirmed stars associated with the faintest MW satellites; however, future spectroscopic measurements of these galaxies may reach the precision necessary to provide complementary constraints [423].

11.4 Milky Way Substructure Modeling

We now review the key components of the Chapters 4–7 [333, 332] MW satellite analyses our study uses. These analyses, which respectively constrain the galaxy–halo connection in CDM and non-CDM scenarios, are based on a forward model of the MW satellite population that combines high-resolution simulations of halos selected to resemble the MW combined with an empirical model for the galaxy–halo connection. These studies account for observational selection functions to fit the MW satellite population in a statistical framework and infer the underlying SHMF, which in turn constrains dark matter physics. For brevity, we mainly describe the Chapter 7 [332] WDM analysis, and we refer the reader to specific sections of Chapter 4 [333] for further methodological details throughout the following subsections.

11.4.1 Milky Way Satellite Data

Chapter 7 [332] analyzes the kinematically confirmed and candidate MW satellites from [149], which were identified using Dark Energy Survey (DES) and Pan-STARRS1 (PS1) data. In particular, [332] analyze 34 satellite galaxies with stellar masses from $\sim 10^2 M_\odot$ to $10^7 M_\odot$. Together, the DES and PS1 datasets cover more than $\sim 75\%$ of the high-Galactic-latitude sky and provide exquisite sensitivity near the LMC due to deep DES photometry in that region. Thus, [332] incorporate both inhomogeneity and incompleteness in the observed MW satellite population through by using the observational selection functions from [149]. Unlike other semiempirical models of the MW satellite population (e.g., [241, 264, 340, 341]), [332] account for the effect of the LMC system on the observed MW satellite population, which is essential to fit the full dataset.

11.4.2 Milky Way Satellite Model

Milky Way Zoom-in Simulations

The MW satellite model used in Chapter 7 [332] is based on high-resolution dark-matter-only zoom-in simulations selected from the suite of 45 MW-mass hosts presented in [315]; technical details on these simulations, which resolve subhalos with virial masses as small as $\sim 10^7 M_\odot$ at $z = 0$, are provided in Appendix F.1. In particular, [332] use the two most “MW-like” host halos in this simulation suite to model the MW satellite population. These hosts have mass and concentration values consistent with recent inferences based on Gaia data [102, 106]. In addition, they have early major mergers that resemble the Gaia-Enceladus event as well as nearby, recently accreted LMC analogs that match the satellite population and kinematics of the real LMC system (see [333] Section 7.2).

Galaxy–Halo Connection Model

To infer the present-day abundance of subhalos in the MW, [332] combine the simulations described above with an empirical model of the galaxy–halo connection (introduced in Chapters 3–4 [336, 333]), which populates subhalos with satellite galaxies in a parametric fashion. By combining these predictions with observational selection functions derived from satellite searches in DES and PS1 data [149], the model is compared to observations assuming that satellites in each survey footprint populate the parameter space of surface brightness and heliocentric distance according to a Poisson process (see [333] Section 6). By marginalizing over the underlying Poisson rate in the calculation of the likelihood for each surface brightness bin, the galaxy–halo connection and dark matter model parameters are fit to data in a Markov Chain Monte Carlo (MCMC) framework.

The majority of the parameters in the Chapter 7 [332] WDM analysis govern the relationship between satellite galaxies and the subhalos they inhabit. For example, these include the slope and scatter of the abundance-matching relation between galaxy luminosity and peak halo maximum circular velocity; the amplitude, scatter, and power-law slope of the relation between galaxy size and halo size; and parameters governing the fraction of low-mass dark matter halos that host observable galaxies. These parameters are not directly relevant for our strong lensing joint analysis because lensing is sensitive to the integrated amount of matter in the lens galaxy and along the line of sight, which is dominated by dark matter. However, they are crucial for robustly modeling the MW satellite population and are marginalized over in our probe combination.

Here, we highlight the aspects of the [332] model that are most relevant for our joint analysis:

(i) *Minimum halo mass (\mathcal{M}_{\min})*: The [332] satellite analysis is consistent with CDM predictions down to a characteristic halo mass scale referred to as the minimum halo mass (\mathcal{M}_{\min}). The minimum halo mass is defined as the peak virial mass of the smallest surviving subhalo inferred to host observed MW satellite galaxies and therefore represents the lowest mass down to which the SHMF is directly constrained by current MW satellite observations. \mathcal{M}_{\min} is jointly inferred along with the fraction of halos that host observable galaxies, which is consistent with 100% down to \mathcal{M}_{\min} . The upper limit on \mathcal{M}_{\min} is calculated by marginalizing over the full posterior distribution, which yields $\mathcal{M}_{\min} < 3.2 \times 10^8 M_{\odot}$ at 95% confidence in the [333] CDM fit (see [333] Sections 4.4, 7.4, and 7.5).

(ii) *Baryonic disruption efficiency (\mathcal{B})*: The efficiency of subhalo disruption due to the Galactic disk. Disruption probabilities due to baryonic physics for the subhalos in the dark-matter-only simulations described above are predicted using the Chapter 2 [337] subhalo disruption model, which is calibrated to hydrodynamic simulations from the Feedback in Realistic Environments (FIRE) project [190]. We note that several subsequent DM-only plus disk [259] and hydrodynamic simulations [383, 404] of MW-mass halos report comparable amounts of subhalo disruption relative to the FIRE simulations used to calibrate the [337] disruption model.

To account for uncertainties resulting from the limited statistics of these training simulations, [332] parameterize the efficiency of subhalo disruption by assigning the following disruption probability to each subhalo in the MW-like zoom-in simulations:

$$p_{\text{disrupt}} \equiv (p_{\text{disrupt},0})^{1/\mathcal{B}}, \quad (11.13)$$

where $p_{\text{disrupt},0}$ is the fiducial disruption probability predicted by the Chapter 2 [337] model, which is a function of the orbital properties (including pericentric distance and accretion time) and internal properties (including mass and maximum circular velocity at accretion) of subhalos. Adopting a lognormal prior on \mathcal{B} centered on the hydrodynamic training simulations (i.e., $\mathcal{B} = 1$), the WDM MW satellite analysis in Chapter 7 [332] yields disruption efficiencies that are consistent with hydrodynamic simulations and rule out very efficient ($\mathcal{B} > 1.9$) and very inefficient ($\mathcal{B} < 0.2$) subhalo disruption at 95% confidence (see [333] Sections 4.3 and 7.4).

(iii) *WDM half-mode mass (M_{hm})*: The characteristic mass scale describing the suppression of the WDM SHMF as defined in Equations 11.6 and 11.10. Due to the population statistics of faint satellites corresponding to low-mass halos, MW satellite analyses have recently achieved upper limits on the half-mode mass that fall below the minimum halo mass associated with observed systems. In particular, for thermal relic WDM, [332] infer $M_{\text{hm}} < 10^{7.5} M_{\odot}$ ($m_{\text{WDM}} > 7.0$ keV) at 95% confidence. [332] scale this constraint by a factor of the maximum possible ratio of the MW halo mass relative to the average host halo mass of the two realistic MW zoom-in simulations used in the inference, which increases M_{hm} constraints by $\sim 25\%$ and yields a fiducial constraint of $M_{\text{hm}} < 10^{7.6} M_{\odot}$ ($m_{\text{WDM}} > 6.5$ keV) at 95% confidence. We discuss the role of the MW halo mass in detail in Section 11.8.1.

11.4.3 Constraints from Milky Way Satellite Observations

Here, we rerun the Chapter 7 [332] WDM MW satellite analysis using priors and a WDM SHMF parameterization chosen to match the [195] lensing analysis, which allows us to self-consistently perform our multidimensional satellite–lensing probe combination. In particular, we rerun the MW satellite analysis adopting a uniform prior of $\mathcal{B} \sim \mathcal{U}(0, 3)$, which ensures that we match the shape of the Σ_{sub} prior used in [195] based on the linear relation between \mathcal{B} and Σ_{sub} we derive in Section 11.6.2. The use of a uniform (rather than lognormal) prior on \mathcal{B} weakens the upper and lower limits of the marginalized posterior from [332] (i.e., $0.2 < \mathcal{B} < 1.9$) from 95% to 68% confidence constraints. In Appendix F.4, we show that the choice of this prior does not significantly impact our joint WDM constraints.

We also use the WDM SHMF and the $\log M_{\text{hm}} \sim \mathcal{U}(5, 10)$ prior assumed in [195]. The resulting marginalized posterior distribution yields $M_{\text{hm}} < 10^{7.4} M_{\odot}$ ($m_{\text{WDM}} > 7.4$ keV) at 95% confidence after MW host halo mass scaling, which is more constraining than the Chapter 7 [332] result despite

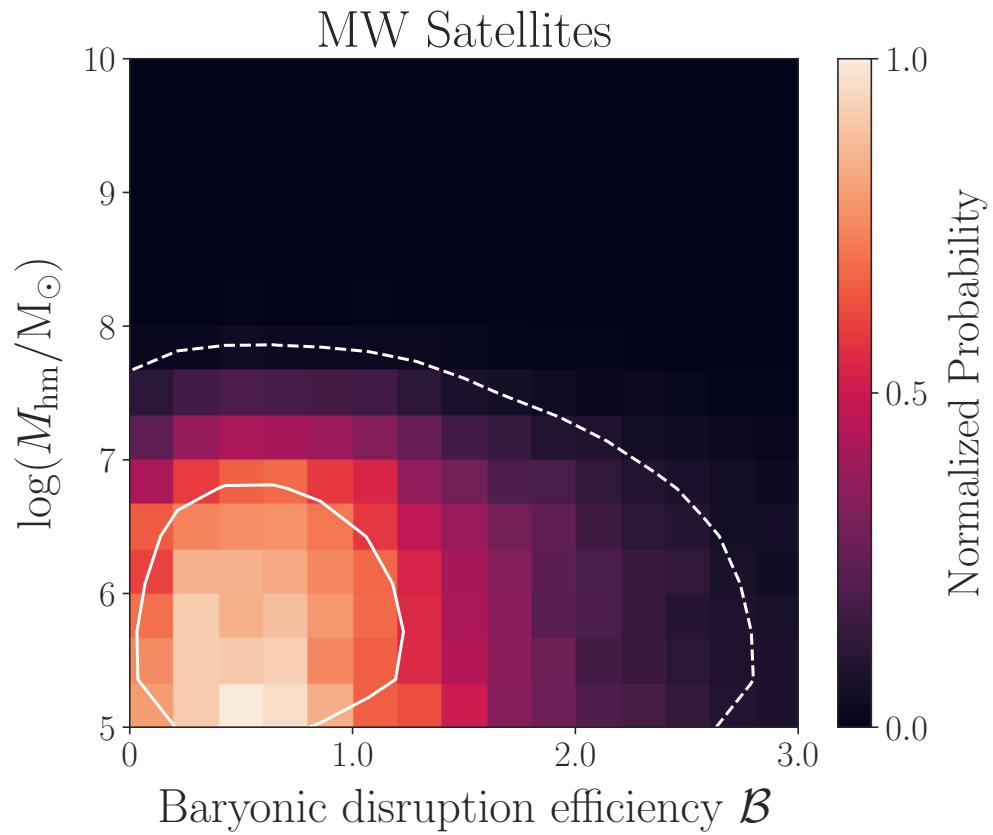


Figure 11.1: Posterior distribution of WDM half-mode mass versus baryonic disruption efficiency from our MW satellite analysis. $\mathcal{B} = 0$ corresponds to zero additional subhalo disruption relative to CDM, and larger values of \mathcal{B} correspond to more efficient subhalo disruption due to baryons. The color map shows the probability density normalized to its maximum value in this parameter space. Solid (dashed) white lines indicate 1σ (2σ) contours for a two-dimensional Gaussian distribution.

using a slightly less suppressed WDM SHMF. This is caused by the change in the lower limit of our $\log M_{\text{hm}}$ prior, which is two orders of magnitude lower than that adopted in [332]. As described in Section 11.7.1 the lower limit of the M_{hm} prior is arbitrary unless we assume that WDM physics manifests at a particular halo mass scale. Thus, we also quote likelihood ratios for both our independent and combined constraints. We find that $M_{\text{hm}} = 10^{7.9} M_{\odot}$ ($m_{\text{WDM}} = 5.2 \text{ keV}$) is disfavored relative to the peak of the marginalized posterior at $10^5 M_{\odot}$ with a 20:1 ratio, consistent with the [332] result.

Figure 11.1 shows the posterior from our updated WDM fit to MW satellite data in the two-dimensional parameter space of M_{hm} versus \mathcal{B} , marginalized over seven other galaxy–halo connection parameters. In Figure 11.1 and subsequent plots, we do not scale parameters to account for MW host halo mass uncertainty unless explicitly noted. We reiterate that our MW satellite analysis only probes systems down to a peak halo mass threshold of $\sim 3 \times 10^8 M_{\odot}$ at 95% confidence and that M_{hm} constraints below this mass scale are driven by the population statistics of halos near the minimum observable halo mass. This is demonstrated in Chapter 7 [332] (see Figure 7.1), which shows that the WDM model ruled out by MW satellites at 95% confidence yields $\sim 25\%$ suppression in subhalo abundances relative to CDM at the minimum halo mass, which is about one order of magnitude larger than M_{hm} .

There is not a strong degeneracy between \mathcal{B} and M_{hm} in Figure 11.1 because \mathcal{B} models the disruptive effects of the MW disk, which suppresses the inner radial distribution of MW satellites in an approximately mass-independent fashion, while M_{hm} models the mass-dependent suppression of the projected SHMF caused by WDM free streaming. Figures 11.2 and F.1 illustrate the effects of \mathcal{B} on the projected SHMF and radial distribution of our MW-like simulations.

11.5 Strong Lens Substructure Modeling

Next, we describe the data and constraints from the [195] quadruply lensed quasar flux ratio analysis our study is based on. Briefly, this analysis combines recent observations of the flux ratios and image positions from eight quadruply imaged quasars with a forward model for the dark matter substructure and line-of-sight halo populations to statistically infer the abundance and concentration of low-mass halos, which in turn constrains the WDM particle mass. Again, we refer the reader to specific sections of [195] for modeling details throughout the following subsections.

11.5.1 Strong Lensing Data

[195] analyze the narrow-line emission from six background quasars presented in [342] and from two additional quasars presented in [344, 343]. These sources have a range of redshifts from $0.8 \lesssim z_s \lesssim 3.7$, while the deflectors span redshifts of $0.2 \lesssim z_d \lesssim 1$ and consist of massive elliptical galaxies. Priors on the masses of the deflector halos are estimated using the stellar mass–velocity dispersion

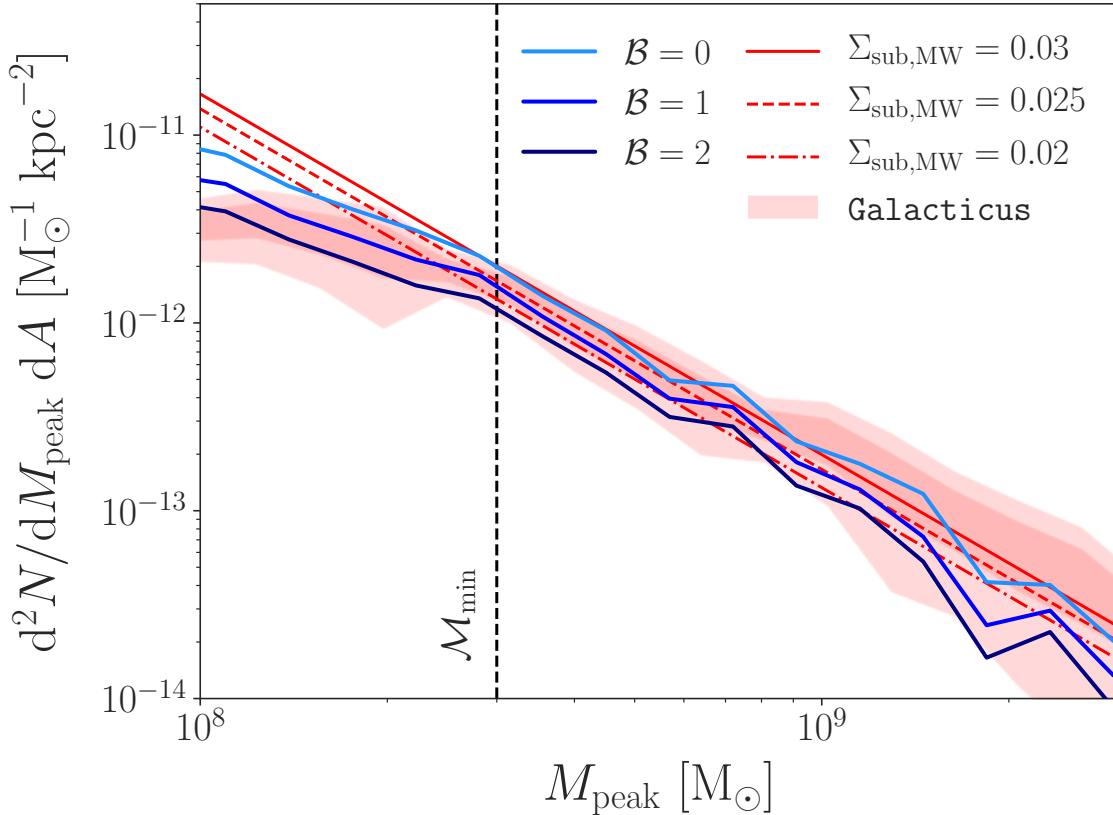


Figure 11.2: Projected SHMFs for MW-like host halos. Blue lines show the average results from the zoom-in simulations used in our MW satellite inference as a function of baryonic disruption efficiency \mathcal{B} ($\mathcal{B} = 0$ corresponds to CDM only and larger values of \mathcal{B} correspond to more efficient subhalo disruption due to baryons). Red lines show our analytic SHMF (Equation 11.5) using the host halo mass and redshift scaling predicted by **Galacticus**, evaluated at the average halo mass of our MW-like simulations with a slope of $\alpha = -1.92$. $\Sigma_{\text{sub},\text{MW}}$ is chosen such that the SHMF amplitude matches our simulations at the subhalo mass corresponding to the faintest observed MW satellites, \mathcal{M}_{min} (dashed vertical line). Dark (light) red contours show 68% (95%) confidence intervals from **Galacticus** for host halos with characteristics matched to our MW-like simulations. We impose the resolution cuts described in Appendix F.1 on the simulation and **Galacticus** results.

relation derived for strong lens galaxies by [36], and typically peak at $\sim 10^{13.3} M_\odot$. We note that [195] excluded quads with main lensing galaxies that contain stellar disks from their analysis.

11.5.2 Strong Lensing Model

The presence of small-scale structure in the lens and along the line of sight can perturb the magnified fluxes of unresolved quasar emission regions. The occurrence rate of the distribution of perturbed flux ratios between the multiple images is therefore sensitive to the underlying population of (potentially dark) subhalos within the host halo of the lens. Importantly, the strong lensing image position and flux ratio data described above are also sensitive to dark matter structure along the entire line of sight from the observer to the lensed quasars.

To perform the inference on the underlying subhalo and line-of-sight mass function population parameters, [195] forward-model the quasar flux ratio with a large set of realizations of the small-scale lensing structure through a multiplane ray-tracing scheme, which accounts for the finite emitting source size and satisfies the astrometric constraints on the positions of the images. The likelihood for the individual lenses' population parameters was constructed with Approximate Bayesian Computation (ABC), and the joint posterior inference was performed by multiplying the individual likelihoods (see [195] Section 2).

Unlike in the case of subhalos, the line-of-sight dark matter structure is unaffected by tidal stripping and disruption. Thus, [195] modeled the line-of-sight dark matter distribution using a Sheth–Tormen [421] mass function with the same WDM SHMF and concentration suppression factors described above for subhalos, along with a contribution from the two-halo term near the main deflector's host halo and an overall scaling factor that allows for uncertainty in the halo mass function amplitude of 20% (see [195] Section 5.3).

11.5.3 Constraints from Strong Lensing Observations

The [195] strong lensing analysis is consistent with CDM predictions for the slope and amplitude of the halo and subhalo mass functions. In particular, [195] derive constraints on the SHMF slope that are consistent with N -body simulations and find that the line-of-sight contribution is consistent with Sheth–Tormen mass function predictions. Moreover, [197] demonstrate that these data are compatible with standard predictions for the CDM mass–concentration relation while self-consistently modeling the effects of tidal stripping on subhalos.

Here, we highlight the constraints from [195] that enter our multidimensional MW satellite–lensing probe combination:

(i) *Projected subhalo number density* (Σ_{sub}): The amplitude of the projected SHMF defined in Equation 11.4. [195] place a lower limit of $\Sigma_{\text{sub}} > 8 \times 10^{-3} \text{ kpc}^{-2}$ at 95% confidence, which (given the lens sample) implies the presence of halos in the $10^6 M_\odot$ – $10^9 M_\odot$ range. Lower values of Σ_{sub}

do not yield sufficient perturbations to reproduce the observed flux ratios, and the [195] analysis does not place an upper limit on Σ_{sub} within the prior range of $\Sigma_{\text{sub}} \sim \mathcal{U}(0, 0.1) \text{ kpc}^{-2}$ (see [195] Sections 3.2 and 6.2).

(ii) *WDM half-mode mass* (M_{hm}): The characteristic mass scale describing the suppression of the WDM SHMF defined in Equations 11.6 and 11.10. For thermal relic WDM, [195] infer $M_{\text{hm}} < 10^{7.8} M_{\odot}$ ($m_{\text{WDM}} > 5.6 \text{ keV}$) at 95% confidence. This constraint results from the fact that warmer models suppress the abundance and concentrations of low-mass halos that contribute to the lensing signal (see [195] Sections 3.4 and 6.2).

Here, we reanalyze the [195] marginalized M_{hm} posterior using a slightly higher lower limit of the $\log M_{\text{hm}}$ prior. We find $M_{\text{hm}} < 10^8 M_{\odot}$ ($m_{\text{WDM}} > 4.9 \text{ keV}$) at 95% confidence, which is slightly less constraining than the [195] result. Again, we also calculate likelihood ratios due to the ambiguity of the M_{hm} prior and find that $M_{\text{hm}} = 10^{8.7} M_{\odot}$ ($m_{\text{WDM}} = 3.0 \text{ keV}$) is disfavored relative to the peak of the marginalized posterior at $10^{6.4} M_{\odot}$ with a 20:1 ratio, consistent with [195].³

The right panel of Figure 11.3 shows the posterior distribution from the fit to strong lensing data in [195] in the two-dimensional parameter space of M_{hm} versus Σ_{sub} , marginalized over the SHMF slope and line-of-sight mass function amplitude. There is a moderate degeneracy between Σ_{sub} and M_{hm} , particularly at high values of Σ_{sub} ; in this regime, it is difficult to distinguish the coincident suppression of the WDM SHMF and mass-concentration relation relative to CDM from changes to the normalization of the CDM SHMF.

11.6 Joint Analysis Methodology

Having described the data, models, and constraints that enter our joint analysis, we now describe our procedure for combining MW satellite and strong lensing constraints in a shared, multidimensional parameter space. In particular, we qualitatively outline our probe combination procedure (Section 11.6.1) and present our method for translating the subhalo disruption efficiency inferred from our MW satellite analysis to projected subhalo number density at the strong lensing scale (Section 11.6.2). We then describe the statistics of our probe combination (Section 11.6.3).

11.6.1 Probe Combination Procedure

Our probe combination qualitatively proceeds as follows; these steps are described in detail in the following subsections:

1. We compare **Galacticus** predictions for MW-mass halos to the projected SHMF inferred from the MW satellite population (Figure 11.2) to construct a relation between the amplitude of

³We have resolved minor errors in the $M_{\text{hm}} - m_{\text{WDM}}$ conversion and likelihood ratios quoted in the original version of [195].

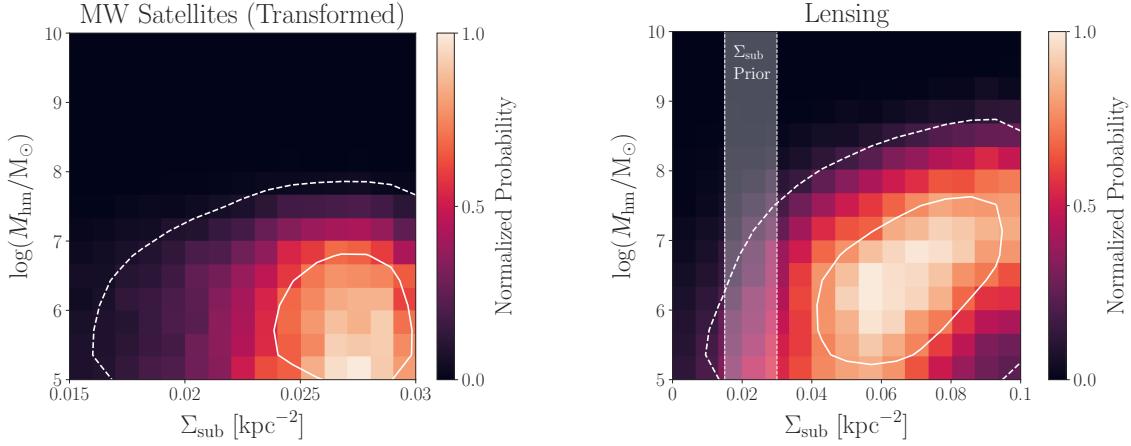


Figure 11.3: Left panel: posterior distribution of WDM half-mode mass versus projected subhalo number density at the strong lensing scale inferred from the MW satellite posterior, transformed according to the procedure in Section 11.6.2, with $q = 1$ (i.e., for equally efficient subhalo disruption due to baryons at the MW and strong lensing host halo mass and redshift scales). Right panel: posterior distribution in the same parameter space from the [195] strong lensing analysis. The vertical band labeled “ Σ_{sub} Prior” shows the range of Σ_{sub} inferred from the MW satellite posterior in our fiducial joint analysis (i.e., $0.015 \text{ kpc}^{-2} \leq \Sigma_{\text{sub}} \leq 0.03 \text{ kpc}^{-2}$). In both panels, color maps show the probability density normalized to its maximum value in each parameter space, and solid (dashed) white lines indicate 1σ (2σ) contours for a two-dimensional Gaussian distribution.

the projected SHMF ($\Sigma_{\text{sub,MW}}$) and the efficiency of subhalo disruption due to baryons in the MW (\mathcal{B});

2. We use this relation to translate the \mathcal{B} – M_{hm} posterior from our MW satellite analysis (Figure 11.1) into a $\Sigma_{\text{sub,MW}}\text{–}M_{\text{hm}}$ posterior distribution;
3. For a given value of the differential subhalo disruption efficiency q , we use Equation 11.3 to translate $\Sigma_{\text{sub,MW}}$ to the strong lensing host halo mass and redshift regime, which yields a $\Sigma_{\text{sub}}\text{–}M_{\text{hm}}$ posterior from MW satellites that can be combined with the strong lensing posterior (Figure 11.3);
4. We construct a joint $\Sigma_{\text{sub}}\text{–}M_{\text{hm}}$ likelihood by multiplying the MW satellite and strong lensing distributions according to the procedure in Section 11.6.3 (Figure 11.4).

This method relies on several simplifying assumptions that could yield additional information if they are self-consistently addressed in a joint likelihood analysis using a model that simultaneously predicts the halo and subhalo distributions relevant for MW satellite and strong lensing analyses. We describe these areas for future work in Section 11.8.1.

11.6.2 Inferring Σ_{sub} from Milky Way Satellites

To connect the host halo mass and redshift regimes probed by strong lensing and MW satellites, we first construct a relation between the subhalo disruption efficiency \mathcal{B} inferred from our MW satellite analysis and the projected subhalo number density $\Sigma_{\text{sub},\text{MW}}$ predicted by evaluating **Galacticus** at the MW halo mass scale. We then translate $\Sigma_{\text{sub},\text{MW}}$ to Σ_{sub} at the strong lensing host halo mass and redshift scale using the dark matter substructure model described in Section 11.3.

Figure 11.2 shows the average $z = 0$ projected SHMF from the two realistic zoom-in simulations used in our MW satellite analysis for our fiducial disruption model calibrated to hydrodynamic simulations (i.e., $\mathcal{B} = 1$) and for bracketing values of the subhalo disruption efficiency (i.e., $\mathcal{B} = 0$ and $\mathcal{B} = 2$) that are ruled out at greater than 68% confidence by MW satellite data as discussed in Section 11.4. To construct the projected SHMF predicted by our analytic substructure model, we use Equation 11.5 with a slope of $\alpha = -1.92$ and with $\mathcal{F}_{\text{CDM}}(M_{\text{MW}}, z_{\text{MW}})$ evaluated at the mean virial mass of our simulated host halos, $M_{\text{MW}} = 1.4 \times 10^{12} M_{\odot}$, and $z_{\text{MW}} = 0$ to account for the CDM dependence on host halo mass and redshift. The zoom-in simulation predictions shown in these panels include both conservative resolution thresholds based on subhalos' peak and $z = 0$ maximum circular velocity values as well as orphan subhalos (i.e., disrupted subhalos that are re-inserted into the simulation and analytically evolved until $z = 0$) using the Chapter 3 [336] model; we provide additional details in Appendix F.1. These choices allow for a more direct comparison to the semianalytically evolved subhalo populations predicted by **Galacticus**, which are less prone to artificial disruption [454, 453, 166].

We construct a relation between $\Sigma_{\text{sub},\text{MW}}$ and \mathcal{B} by matching our analytic prediction from Equation 11.5 to the average projected SHMF predicted by our MW-like simulations as a function of \mathcal{B} , as illustrated in Figure 11.2. In particular, we match the amplitudes of the projected SHMFs inferred from our MW satellite analysis and predicted by Equation 11.5 at the minimum observable halo mass of $3.2 \times 10^8 M_{\odot}$ and within a fixed radius of 300 kpc (roughly corresponding to the virial radius of the MW host halo), chosen to match the Chapter 7 [332] analysis. Our choice to match these SHMFs at the minimum halo mass is conservative because our MW satellite analysis is not sensitive to subhalos below this mass scale at 95% confidence. We then translate $\Sigma_{\text{sub},\text{MW}}$ to the strong lensing regime using Equation 11.3, which yields

$$\frac{\Sigma_{\text{sub}}}{\text{kpc}^{-2}} = \frac{0.03 - 0.005\mathcal{B}}{q}. \quad (11.14)$$

This relation allows us to transform the $\mathcal{B}-M_{\text{hm}}$ MW satellite posterior from Figure 11.1 into the $\Sigma_{\text{sub}}-M_{\text{hm}}$ parameter space for a given value of q . Note that Σ_{sub} scales linearly with MW halo mass because it measures the projected SHMF amplitude within a fixed physical radius in our model.

The $\mathcal{B}-\Sigma_{\text{sub},\text{MW}}$ relation constructed above is only based on the amplitude of the projected SHMF from our two realistic MW-like simulations, measured at the minimum halo mass scale of $M_{\text{peak}} =$

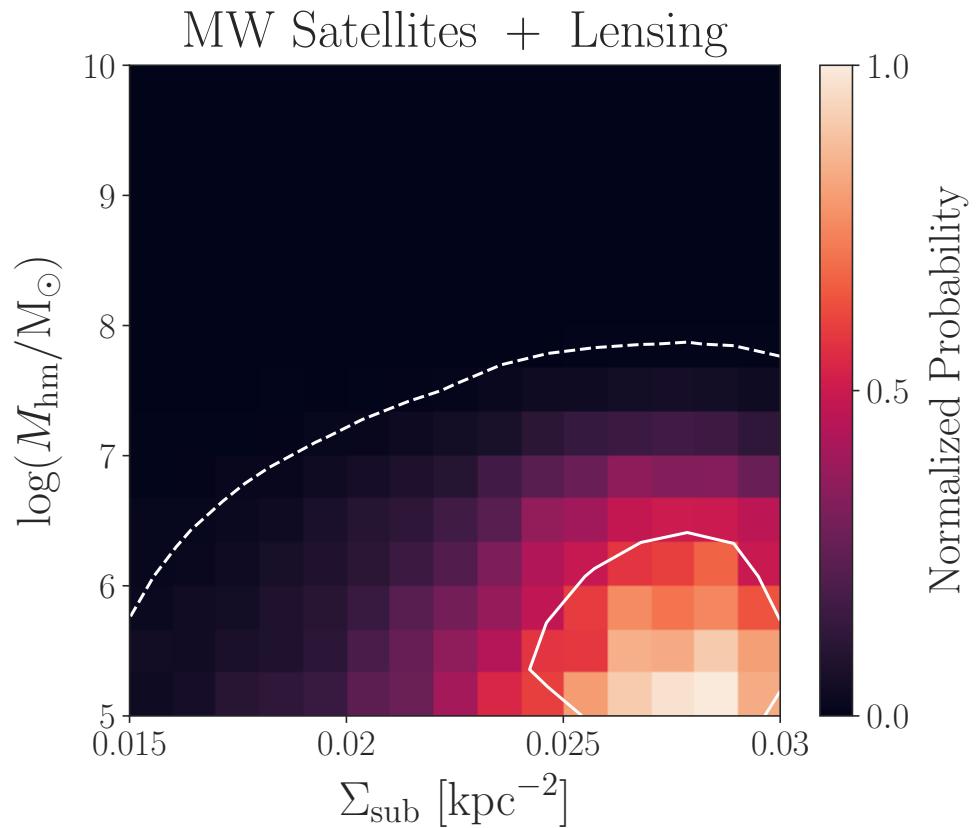


Figure 11.4: Joint marginal likelihood of WDM half-mode mass versus projected subhalo number density at the strong lensing scales from our combined MW satellite–strong lensing posterior, transformed according to the procedure in Section 11.6.2, with $q = 1$. The colormap shows the probability density normalized to its maximum value in this parameter space. Solid (dashed) white lines indicate 1σ (2σ) contours for a 2-dimensional Gaussian distribution.

$3 \times 10^8 M_\odot$. In general, the procedure to infer $\Sigma_{\text{sub},\text{MW}}$ from an estimate of the MW SHMF should account for both host-to-host variation in the SHMF within the range of MW host halo properties allowed by observations and Poisson scatter in the SHMF given the range of subhalo masses probed by MW satellite measurements. Our SHMF matching procedure is intentionally simplistic because—as demonstrated in Appendix F.3—host-to-host and Poisson scatter in the projected SHMF near the minimum halo mass scale are both subleading uncertainties compared to the range of differential subhalo disruption efficiencies due to baryons that we explore. We therefore leave a statistically rigorous construction of the \mathcal{B} – Σ_{sub} relation to future work that propagates such uncertainties into the joint analysis at the likelihood level.

The result of the transformation in Equation 11.14 is shown in the left panel of Figure 11.3 for our fiducial model of $q = 1$ (i.e., for equally efficient subhalo disruption in the MW and strong lensing host halos). The typical Σ_{sub} values favored by the MW satellite posterior for this choice of q are significantly smaller than the largest values allowed by the [195] lensing analysis; we return to this point below.

The lack of degeneracy observed between \mathcal{B} and M_{hm} in Figure 11.1, which results from the joint constraining power of the MW satellite radial distribution and luminosity function for subhalos near the minimum observable halo mass, persists in the $\Sigma_{\text{sub}}\text{--}M_{\text{hm}}$ parameter space. On the other hand, strong lensing flux ratio statistics probe an integrated combination of subhalo masses and concentrations. The lensing analysis is currently less sensitive to subhalos in specific mass ranges than MW satellite population statistics and therefore exhibits a stronger $\Sigma_{\text{sub}}\text{--}M_{\text{hm}}$ degeneracy in Figure 11.3. However, because lensing measurements do not depend on the connection between subhalos and luminous matter, they can probe subhalos below the minimum observable halo mass.

In Figure 11.2 and Appendix F.2, we also compare our simulation results to the SHMF in terms of both peak and present-day subhalo mass and to the radial subhalo distribution predicted by **Galacticus** for 14 halos selected to match the characteristics of our MW-like simulations. In all cases, we apply the same cuts on peak and present-day subhalo maximum circular velocity when comparing **Galacticus** to our zoom-in simulations; the details of these resolution cuts and our **Galacticus** runs are described in Appendix F.2. Note that these **Galacticus** predictions should be compared to our $\mathcal{B} = 0$ simulation results because the current **Galacticus** implementation does not model subhalo disruption due to central galaxies. The **Galacticus**-predicted SHMFs agree well with our simulations in terms of both peak and present-day subhalo mass, lending confidence to the choice of $\mathcal{F}_{\text{CDM}}(M_{\text{host}}, z_{\text{host}})$ that enters our calibration of the projected subhalo number density at the MW scale.

Because we construct a $\Sigma_{\text{sub}}\text{--}\mathcal{B}$ relation based on our specific MW-like simulations, which have been shown to match the MW satellite population, we imposed several host halo selection criteria on the **Galacticus** runs used for the comparison above. These conditions include the existence of a realistic LMC analog system and a Gaia-Enceladus-like merger event. We emphasize that validating

seianalytic halo and subhalo population predictions with self-consistent simulation suites of both MW-like and strong-lens-like halos is an important avenue for future work. As discussed in Section 11.8.1, increasingly precise near-field observations and complementary data for strong lens systems will allow us to mitigate the impact of additional host halo properties including concentration (which is known to influence subhalo populations at fixed halo mass, e.g., [496, 499, 236, 315, 173]) and the characteristics of the local dark matter environment in future joint analyses.

11.6.3 Probe Combination Statistics

Having placed the lensing and MW satellite posteriors on the same footing, we now proceed to combine them to construct a joint $\Sigma_{\text{sub}}-M_{\text{hm}}$ likelihood as follows. Formally, we write our joint MW satellite and strong lensing analysis as a combined Bayesian inference problem,

$$\begin{aligned} P(\boldsymbol{\theta}|\mathbf{D}) &\propto P(\mathbf{D}|\boldsymbol{\theta}) \times P(\boldsymbol{\theta}) \\ &= P(\mathbf{D}_{\text{MW}}|\boldsymbol{\theta}_{\text{MW}}) \times P(\mathbf{D}_{\text{lensing}}|\boldsymbol{\theta}_{\text{lensing}}) \times P(\boldsymbol{\theta}), \end{aligned} \quad (11.15)$$

where $\boldsymbol{\theta}$ is the vector of parameters in both the lensing and satellite analyses including q (where the shared parameters Σ_{sub} and M_{hm} only appear once), $\boldsymbol{\theta}_{\text{MW}}$ ($\boldsymbol{\theta}_{\text{lensing}}$) are the parameters entering the MW satellite (strong lensing) inference, $\mathbf{D} = [\mathbf{D}_{\text{MW}}, \mathbf{D}_{\text{lensing}}]$ is the joint datavector, and $P(\boldsymbol{\theta})$ is the prior distribution over all model parameters.

Next, for a given value of $q \in [0.5, 2]$, we marginalize over the independent parameters (i.e., the seven galaxy–halo connection parameters in the satellite analysis described in Section 11.4 and the SHMF slope and line-of-sight contribution in the lensing analysis described in Section 11.5) to arrive at a combined $\Sigma_{\text{sub}}-M_{\text{hm}}$ posterior distribution,

$$\begin{aligned} P(\Sigma_{\text{sub}}, M_{\text{hm}}|\mathbf{D}) &= P(\Sigma_{\text{sub}}, M_{\text{hm}}|\mathbf{D}_{\text{MW}}) \times P(\Sigma_{\text{sub}}, M_{\text{hm}}|\mathbf{D}_{\text{lensing}}) \\ &\propto P(\mathbf{D}|\Sigma_{\text{sub}}, M_{\text{hm}}) \times P(\Sigma_{\text{sub}}, M_{\text{hm}}). \end{aligned} \quad (11.16)$$

We assume independent priors for Σ_{sub} and M_{hm} ; in particular, use the prior distributions from [195],

$$P(\Sigma_{\text{sub}}/\text{kpc}^{-2}) = \mathcal{U}(0, 0.1) \quad (11.17)$$

$$P(\log(M_{\text{hm}}/M_{\odot})) = \mathcal{U}(5, 10). \quad (11.18)$$

As described in Sections 11.4–11.5, we choose a lower limit of $\log(M_{\text{hm}}/M_{\odot}) = 5$ because models with even lower M_{hm} values are indistinguishable from CDM in both our MW satellite and strong lensing analyses. For simplicity, we repeat our analysis at several fixed values of q rather than marginalizing over this parameter. We note that our WDM limits marginalized over q are nearly

identical to our fiducial ($q = 1$) result in the absence of a well-motivated, nonuniform prior for q based on hydrodynamic simulations (see Section 11.7.3).

Based on Equation 11.14, our MW satellite inference only samples $\Sigma_{\text{sub}}/\text{kpc}^{-2} \in [0.015q^{-1}, 0.03q^{-1}] \in [0.0075, 0.06]$ given our prior of $\mathcal{B} \sim \mathcal{U}(0, 3)$ and our assumed range of $q \in [0.5, 2]$ (note that $\mathcal{B} \geq 0$ by definition). Thus, our fiducial (i.e., $q = 1$) analysis is restricted to the range $0.015 \leq \Sigma_{\text{sub}}/\text{kpc}^{-2} \leq 0.03$, labeled “ Σ_{sub} Prior” in Figures 11.3 and 11.5. This range is set by combining the MW satellite posterior and zoom-in simulations with our analytic SHMF prediction and is narrower than the Σ_{sub} range considered in [195], which did not enforce priors based on cosmological simulations. This difference limits the range of Σ_{sub} values from the lensing analysis relevant for our probe combination, but it does not formally affect our calculation of the marginal likelihood because the effective prior on Σ_{sub} from the MW satellite analysis is nevertheless uniform.

Thus, exploiting the fact that our priors are uniform in both Σ_{sub} and $\log M_{\text{hm}}$, we arrive at a joint marginal likelihood for these quantities in terms of the marginalized MW satellite and strong lensing posteriors,

$$\begin{aligned} P(\mathbf{D}_{\text{MW}}, \mathbf{D}_{\text{lensing}} | \Sigma_{\text{sub}}, \log M_{\text{hm}}) \propto & \ P(\Sigma_{\text{sub}}, \log M_{\text{hm}} | \mathbf{D}_{\text{MW}}) \\ & \times P(\Sigma_{\text{sub}}, \log M_{\text{hm}} | \mathbf{D}_{\text{lensing}}). \end{aligned} \quad (11.19)$$

This joint marginal likelihood is illustrated in Figure 11.4 and analyzed in Section 11.7. Because the independently derived M_{hm} distributions are consistent (Section 11.7.2) and the Σ_{sub} distributions are only in mild tension (Section 11.7.4), we do not formally test for statistical consistency between the MW satellite and strong lensing analyses before constructing the joint likelihood.

Because we fix the slope of the projected SHMF in our analytic substructure model at $\alpha = -1.92$, our joint analysis is effectively performed at a thin slice in α of the [195] posterior (we remind the reader that α is defined in a CDM context well above the cutoff scale). We do not expect this choice to significantly affect our results because α is not highly degenerate with the parameters of interest (i.e., M_{hm} and Σ_{sub}) in the [195] analysis. However, in Section 11.8.1 we emphasize the importance of jointly inferring the SHMF slope in future work, and we discuss the role of the remaining line-of-sight dimension of the [195] posterior, modeled by the amplitude of the Sheth–Tormen halo mass function, that is marginalized over in our analysis.

11.7 Results

We now present our probe combination results, which are summarized in Figures 11.4–11.5 and Table 11.1. We first describe our conventions for calculating WDM constraints in Section 11.7.1. We then describe our fiducial joint analysis results in Section 11.7.2, and we explore the impact of varying the differential subhalo disruption efficiency due to baryons in Section 11.7.3 and our

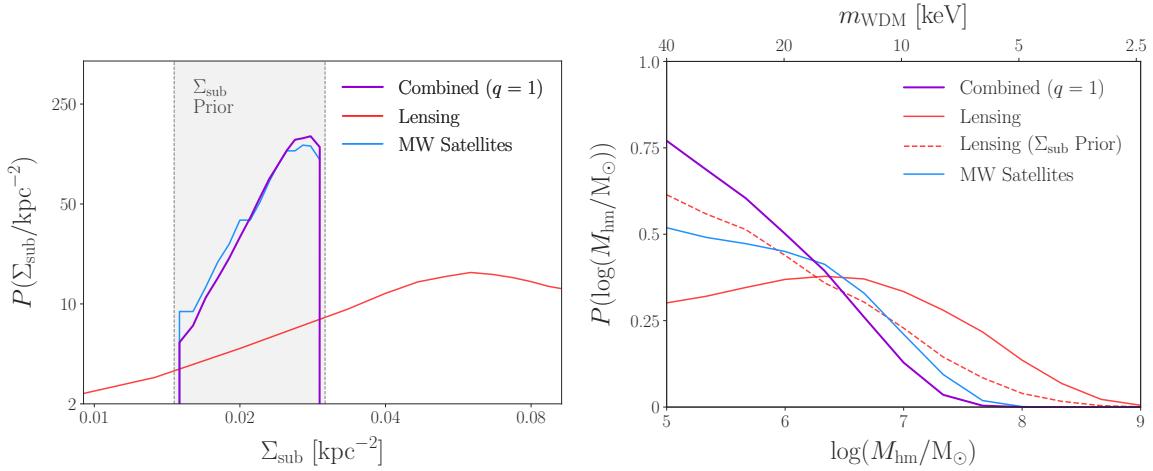


Figure 11.5: Marginal distributions from our joint MW satellite–strong lensing likelihood (Figure 11.4) for projected subhalo number density at the strong lensing scale (left panel) and WDM half-mode mass (right panel), assuming equally efficient subhalo disruption due to baryons in the MW and strong lens systems ($q = 1$). The marginalized MW satellite posterior is shown in blue, the marginalized strong lensing posterior is shown in red, and results obtained from our probe combination and marginalized over the remaining dimension are shown in purple. In the left panel, the vertical band labeled “ Σ_{sub} Prior” shows the range of Σ_{sub} inferred from the MW satellite posterior in our fiducial joint analysis (i.e., $0.015 \text{ kpc}^{-2} \leq \Sigma_{\text{sub}} \leq 0.03 \text{ kpc}^{-2}$, slightly offset from the posteriors for visual clarity), and the dashed red line on the right panel shows the lensing half-mode mass posterior restricted to this range of Σ_{sub} values.

constraints on the projected SHMF amplitude in Section 11.7.4.

11.7.1 Conventions for WDM Constraints

To quantify the WDM constraints corresponding to the joint likelihood derived in Section 11.6.3, we marginalize over the Σ_{sub} dimension and construct the following summary statistics from the marginal M_{hm} likelihood:

- (i) *Confidence intervals*: defined as the range of parameter values enclosing a particular fraction of the integrated marginal likelihood. Following common practice in the WDM literature, we quote upper limits on M_{hm} and lower limits on m_{WDM} at 95% confidence.
- (ii) *Marginal likelihood ratios*: defined as the parameter value at which the marginal likelihood probability density falls to a particular fraction of its peak value. Following [195], we quote the M_{hm} and m_{WDM} values disfavored with 20:1 marginal likelihood ratios.

Although confidence intervals capture more information about the shape of the probability density and are commonly quoted in the WDM literature (e.g., [459, 232]), they depend on the arbitrary choice of a lower limit on the M_{hm} prior (or equivalently, an upper limit on the m_{WDM} prior) as noted

	One-dimensional	$q = 0.5$	$q = 1$	$q = 2$
95% confidence level M_{hm} [M_{\odot}]	$10^{7.2}$	$10^{7.1}$	$10^{7.0}$	$10^{6.9}$
95% confidence level m_{WDM} [keV]	8.4	9.1	9.7	10.4
20:1 likelihood ratio M_{hm} [M_{\odot}]	$10^{7.7}$	$10^{7.6}$	$10^{7.4}$	$10^{7.3}$
20:1 likelihood ratio m_{WDM} [keV]	6.0	6.4	7.4	7.9

Table 11.1: 95% confidence and 20:1 likelihood ratio upper limits on M_{hm} and corresponding lower limits on m_{WDM} for our multidimensional probe combination for various differential subhalo disruption efficiency values q , and for an analysis that combines the fully marginalized one-dimensional M_{hm} distributions. $q = 0.5$ corresponds to twice as efficient subhalo disruption due to baryons in the MW relative to strong lenses, $q = 1$ (our fiducial model) corresponds to equally efficient subhalo disruption due to baryons, and $q = 2$ corresponds to twice as efficient subhalo disruption due to baryons in strong lenses.

above. In particular, small-scale structure data are currently consistent with CDM and therefore yield one-sided limits on M_{hm} or m_{WDM} ; without assuming a preferred scale for a small-scale structure cutoff due to WDM (or other non-CDM) physics, this makes the lower limit of the M_{hm} prior arbitrary. This situation motivated several authors (e.g., [161, 195]) to quote alternative summary statistics including marginal likelihood ratios that are less dependent on the choice of M_{hm} prior, and we follow this practice here. Similarly, we follow both [195] and [332] by adopting a logarithmic prior on M_{hm} because any other choice would not be invariant to rescaling m_{WDM} (e.g., see the discussion in [241]).

11.7.2 Fiducial WDM Constraints

We now present the results of our joint analysis for our fiducial subhalo disruption efficiency model of $q = 1$, which assumes equally efficient subhalo disruption due to baryons in the MW and in strong lens host halos, which is broadly compatible with the results of hydrodynamic simulations (see Section 11.3.3). The combined $\Sigma_{\text{sub}}-M_{\text{hm}}$ marginal likelihood is shown in Figure 11.4 and the corresponding one-dimensional marginalized likelihoods for Σ_{sub} and M_{hm} are shown in Figure 11.5. The joint marginal likelihood retains the shape of the $\Sigma_{\text{sub}}-M_{\text{hm}}$ distribution from the transformed MW satellite posterior and from the lensing analysis limited to the range of Σ_{sub} inferred from our MW satellite analysis according to the procedure in Section 11.6.2. Moreover, the joint marginal likelihood visibly prefers lower values of M_{hm} than either posterior alone, demonstrating the unique constraining power accessible when combining independent small-scale structure probes in a multidimensional parameter space.

Consistent with these qualitative aspects of the joint $\Sigma_{\text{sub}}-M_{\text{hm}}$ likelihood, the upper limit of the marginal M_{hm} likelihood shown in the right panel of Figure 11.5 is noticeably lower than either of the individual constraints from MW satellites or strong lensing. Quantitatively, the upper limit on M_{hm} from our joint analysis improves upon those set by the MW satellite and strong lensing

analyses individually by $\sim 60\%$, leading to a $\sim 30\%$ increase in the strength of the lower limit on m_{WDM} . Specifically, the 95% confidence limit of $M_{\text{hm}} < 10^{7.4} M_{\odot}$ ($m_{\text{WDM}} > 7.4 \text{ keV}$) from our MW satellite analysis improves to $M_{\text{hm}} < 10^{7.0} M_{\odot}$ ($m_{\text{WDM}} > 9.7 \text{ keV}$). We find a similar level improvement in terms of likelihood ratios, with $M_{\text{hm}} = 10^{7.4} M_{\odot}$ ($m_{\text{WDM}} = 7.4 \text{ keV}$) ruled out at 20:1 relative to the peak of the marginal likelihood at the lower limit of the prior at $10^5 M_{\odot}$.

To derive these limits, we conservatively increased the M_{hm} values returned by our joint analysis by a factor of $\sim 25\%$ to account for the maximum mass of the MW halo relative to the average host halo masses of our zoom-in simulations, following [332]. As demonstrated in the following subsection, propagating the MW halo mass uncertainty into the Σ_{sub} dimension would have a negligible impact on the results compared to uncertainties in the efficiency of subhalo disruption due to baryons, so we do not perform this scaling for simplicity.

Our fiducial constraint of $m_{\text{WDM}} > 9.7 \text{ keV}$ at 95% confidence is one of the most stringent limits on the WDM particle mass set by small-scale structure observations to date. Moreover, it is set using only existing strong lensing and MW satellite measurements, underscoring the importance of unified, multidimensional small-scale structure analyses as the corresponding measurements continue to improve. Joint model-building efforts that further incorporate Ly α forest [459, 232] and stellar stream [41] constraints while retaining the unique information provided by each probe will therefore be particularly fruitful.

11.7.3 Impact of the Differential Subhalo Disruption Efficiency due to Baryons

We now explore the impact of the differential efficiency of subhalo disruption due to baryons on our WDM constraints. Table 11.1 lists the M_{hm} and m_{WDM} 95% confidence level and 20:1 likelihood ratio limits for $q = 0.5, 1$, and 2 , and the right panel of Figure 11.6 shows the corresponding joint marginal likelihoods. In Table 11.1 and Figure 11.6, we also show the result of combining the fully marginalized one-dimensional M_{hm} posteriors from our MW satellite and strong lensing analyses.

As demonstrated in the right panel of Figure 11.6, the joint marginal likelihoods for M_{hm} become increasingly constraining as q increases. This is due to the fact that the transformed $\Sigma_{\text{sub}}-M_{\text{hm}}$ posterior distribution from MW satellites (Figure 11.3 left panel) breaks the degeneracy between these parameters present in the strong lensing posterior. In particular, larger values of q correspond to more efficient subhalo disruption in strong lens host halos relative to the MW and yield lower inferred values of Σ_{sub} at the strong lensing scale according to Equation 11.14. This shifts the region of two-dimensional parameter space in which we multiply the MW satellite and strong lensing posteriors toward lower values of Σ_{sub} . Thus, because the low- Σ_{sub} region of the lensing posterior does not allow for large values of M_{hm} , larger values of q yield more stringent joint M_{hm} constraints (and vice versa for smaller values of q). Indeed, as shown in the right panel of Figure 11.5, restricting the strong lensing posterior to the range of Σ_{sub} inferred from our MW satellite analysis for $q = 1$

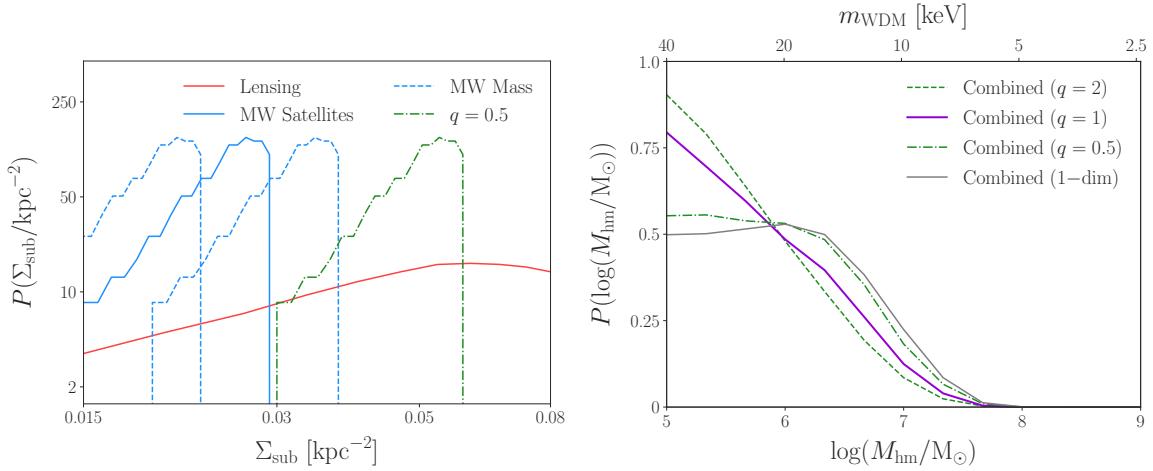


Figure 11.6: Left panel: the impact of systematics on the marginalized one-dimensional posterior distributions of projected subhalo number density at the strong lensing scale. The marginalized posterior distribution from our MW satellite analysis is shown in blue, the marginalized strong lensing posterior is shown in red, the dashed blue distributions indicate additional uncertainty in our MW satellite inference due to the mass of the MW halo, and the dotted-dashed green distribution illustrates the effects of systematic uncertainty in the differential efficiency of subhalo disruption due to baryons at the MW and strong lensing host halo scales. Right panel: joint marginal likelihood of WDM half-mode mass for our MW satellites plus strong lensing probe combination. Joint likelihoods are shown for equally efficient subhalo disruption in the MW and strong lens host halo mass and redshift regimes ($q = 1$, purple), twice as efficient disruption due to baryons in the MW relative to strong lens halos ($q = 0.5$, dotted-dashed green), and twice as efficient disruption in strong lens halos relative to the MW ($q = 2$, dashed green). The gray distribution shows the result of combining the fully marginalized one-dimensional M_{hm} posteriors derived from strong lensing and MW satellites.

significantly strengthens the M_{hm} constraint set by lensing alone.

Despite the qualitative effects of varying the differential subhalo disruption efficiency described above, varying q within a reasonably broad range only impacts the results of our probe combination at the $\sim 10\%$ level in terms of m_{WDM} . As discussed in Section 11.8.1, the differential efficiency of subhalo disruption due to baryons is one of several systematics that impact our probe combination at this level, all of which must be controlled in a joint modeling framework to claim a detection of non-CDM physics at the corresponding level of precision. Figure 11.6 and Table 11.1 demonstrate that combining the MW satellite and strong lensing posteriors with any value of q —i.e., performing the combination in multiple dimensions—is more constraining than combining the fully marginalized M_{hm} posteriors, as expected.

11.7.4 Projected Subhalo Number Density Constraints

As demonstrated in the left panel of Figure 11.5, the marginalized posterior for Σ_{sub} from strong lensing accommodates significantly larger values than we infer from our MW satellite analysis. In particular, our fiducial joint analysis yields a marginalized posterior distribution from MW satellites that peaks at $\Sigma_{\text{sub}} \approx 0.025 \text{ kpc}^{-2}$; moreover, $\Sigma_{\text{sub}} > 0.03 \text{ kpc}^{-2}$ is not sampled because these projected SHMF amplitudes are larger than the average of our MW-like simulations.⁴ Meanwhile, Σ_{sub} values in this range are disfavored in the lensing posterior relative to its mild peak at $\Sigma_{\text{sub}} \approx 0.067 \text{ kpc}^{-2}$ by a ratio of $\sim 2:1$. Although this is not a significant tension, it is worth exploring in future work that places Σ_{sub} constraints at various host mass and redshift scales in the context of expectations from cosmological simulations. For example, [287] identify potentially significant contributions from backsplash halo populations near strong lenses beyond those captured by the two-halo term used in [195], which (if modeled) may lower the inferred range of Σ_{sub} and strengthen the corresponding WDM constraints. Furthermore, there are potential differences between the surviving subhalo populations inferred from our MW satellite and strong lensing analyses caused by tidal stripping, although heavily stripped halos do not dominate the signal in either case. Thus, although it is unlikely because the subhalos that contribute to strong lensing flux ratio statistics are usually tidally truncated well outside of their NFW scale radius [195, 323], a careful analysis of whether these systems can be stripped severely enough such that their luminous content is affected warrants detailed investigation in future work.

In the left panel of Figure 11.6, we show how the Σ_{sub} posterior from our MW satellite analysis shifts as a function of both the differential subhalo disruption efficiency due to baryons, q , and the MW halo mass, where we use the MW host halo mass uncertainties discussed in Section 11.4 and assume that $\Sigma_{\text{sub}} \propto M_{\text{MW}}$. The Σ_{sub} distribution from MW satellites is clearly sensitive to both of these systematic uncertainties, which we discuss further in Section 11.8. Because varying q changes the inferred Σ_{sub} distribution in the strong lens halo mass and redshift regime, this quantity can potentially be constrained as the precision of Σ_{sub} constraints from strong lensing increases. Although we do not attempt to constrain q here, we note that our results disfavor simultaneously high MW halo mass and low subhalo disruption efficiency due to baryons in the MW relative to strong lens host halos, which is physically reasonable.

11.8 Discussion

We now place the WDM and SHMF constraints from our MW satellite–strong lensing probe combination in context by discussing key systematics (Section 11.8.1) and comparing our study to other recent analyses (Section 11.8.2).

⁴The upper limit of this prior increases to $\Sigma_{\text{sub}} \sim 0.04 \text{ kpc}^{-2}$ when accounting for uncertainties in the mass of the MW host halo, which is still much lower than the largest Σ_{sub} inferred in [195].

11.8.1 Systematics

The analysis presented above casts MW satellite and strong lensing constraints in a shared, multi-dimensional parameter space for the first time. We emphasize that our WDM constraints (Section 11.7.2) are conservative due to our broad priors on key systematics and are robust to the modeling uncertainties directly addressed in the joint analysis at the $\sim 10\%$ level (Section 11.7.3). Our work therefore provides important foundations for more detailed modeling frameworks that simultaneously constrain MW satellite and lensing observables at the likelihood level.

Nevertheless, our analysis makes several simplifying assumptions that circumvent a joint likelihood analysis. We regard these as crucial areas for future model-building work in preparation for next-generation facilities and surveys, both for the MW satellite–strong lensing probe combination presented here and to further combine these probes with analyses of stellar stream perturbations, the Ly α forest, and any other novel probes of small-scale structure. In general, joint small-scale structure constraints may be sensitive to additional “nuisance parameters” distinct from those governing non-CDM physics, which must be simultaneously measured to robustly claim evidence for a deviation from CDM. This underscores the importance of our multidimensional approach and of the following systematics, which we plan to build a joint model to simultaneously infer in future work.

SHMF slope. We assume a particular value of the SHMF slope α when constructing the \mathcal{B} – Σ_{sub} relation in Section 11.6, thereby taking a thin slice through this dimension of the posterior from [195]. Although current MW satellite analyses do not strongly constrain the SHMF slope, future constraints from the MW satellite population probed by LSST may be sensitive to this quantity due to excellent observational sensitivity at the faint end of the satellite luminosity function throughout the MW virial radius (e.g., [237, 214, 151]). Meanwhile, the [195] strong lensing analysis already mildly constrains the SHMF slope, and this sensitivity will drastically increase with larger lens samples. Exploiting all of these data will require self-consistent suites of high-resolution simulations of both MW-like systems (including realistic LMC analogs) and strong-lens-like systems, which we are currently developing. Few such high-resolution zoom-in simulations at the group-mass scale have been performed, and these are particularly valuable to validate the predictions of semianalytic models like `Galacticus` used to inform strong lens substructure models. These studies must be coupled with detailed models for the impact of baryonic physics on small-scale dark matter structure because it is expected to significantly affect both the amplitude and slope of SHMF at low halo masses [57].

Line-of-sight halo mass function. We marginalized over the amplitude of the line-of-sight halo mass function in our probe combination, noting that the [195] lensing analysis our work is based on does not constrain this quantity within a broad prior range of $\pm 20\%$ relative to the mean Sheth–Tormen prediction. However, detailed zoom-in simulations of strong lens analogs coupled with realizations from cosmological simulations of the line-of-sight halo populations may provide more informative theoretical priors that—combined with upcoming strong lens discoveries and follow-up imaging and spectroscopy—will yield more decisive differential measurements of the line-of-sight

and substructure contributions to the lensing signal (see [287] for a recent discussion). This will ultimately allow Σ_{sub} to be measured more precisely, breaking degeneracies with WDM physics and facilitating a more direct combination with MW satellite data.

Subhalo disruption efficiency due to baryons. We combined MW satellite and strong lensing constraints at fixed values of the differential subhalo disruption efficiency due to baryons, q . Although q does not significantly affect the joint WDM limits presented here (Section 11.7), this quantity represents a key systematic that must be addressed in dedicated modeling work. In particular, it will be fruitful to analyze samples of hydrodynamic simulations at the MW and group-mass scales to refine subhalo disruption models that can be applied to larger simulation suites efficiently (e.g., [337]). Constructing a physically motivated model for the differential efficiency of subhalo disruption due to baryons in strong lens systems and the MW will again allow for more informative theoretical priors in joint analyses, enabling robust constraints on deviations from CDM predictions.

Milky Way and strong lens host halo properties. The mass of the MW halo remains a key systematic for the interpretation of MW satellite measurements in terms of the underlying SHMF which then propagates into joint small-scale structure constraints. The MW halo mass is a particularly important nuisance parameter for setting non-CDM constraints because the (lack of a) turnover in low-mass subhalo abundances is inferred from the SHMF corresponding to MW satellite observations, while the SHMF amplitude scales linearly with host halo mass. In our analysis, uncertainty in the MW halo mass significantly affects our M_{hm} and m_{WDM} constraints, and we currently take a conservative approach to marginalize over this dependence. Forthcoming Gaia data releases will increase the precision of MW halo mass measurements, and combining detailed simulation suites of MW-like halos spanning the inferred mass range with next-generation observations of the MW satellite population will allow us to derive joint constraints on the MW halo mass and SHMF (e.g., see [341]).

Meanwhile, strong lensing measurements are less sensitive to host halo mass uncertainty because they probe both the SHMF, small-scale structure along the line of sight, and the concentrations of low-mass halos and subhalos. Nonetheless, the details of strong lens host halo selection functions are relatively unexplored (e.g., see [430]), and will be better quantified using a variety of data including weak lensing and satellite velocity dispersion measurements. These efforts will lead to more precise constraints on the masses, secondary properties, and environments of strong lens host halos, further mitigating key theoretical uncertainties in forward models of strong lensing data.

11.8.2 Comparison to Recent Studies

[161] recently presented a joint analysis of small-scale structure probes including MW satellite galaxies and gravitational imaging, with several distinct assumptions underlying the individual and joint modeling of these probes relative to our work. Here, we discuss the most important aspects of our individual models for MW satellites and strong lensing flux ratio statistics as well as our probe

combination procedure relative to the [161] study.

For MW satellites, the Chapter 7 [332] study upon which we base our analysis explicitly includes realistic LMC analog systems in the simulations used to perform the inference. This allows us to use the entire population of observed MW satellite galaxies—and particularly those within and near the DES footprint—without down-weighting systems based on the probability they are associated with the LMC, strengthening our dark matter constraints relative to the [340, 341] MW satellite analyses that the [161] joint constraints are based on. In addition, unlike [340, 341], we follow [332] by using the newest and most precise versions of DES and PS1 observational selection functions from [149]. Importantly, these selection functions depend on satellite galaxy size, which is a crucial driver of satellite detectability that directly informs the translation from MW satellite observations to the underlying SHMF. This highlights the importance of including a model for the relationship between subhalo and satellite galaxy size like the one used in our analysis. As discussed in Section 11.4, we also marginalize over MW halo mass and the efficiency of subhalo disruption due to baryonic physics, which are both key systematics in the MW satellite inference. Our MW host halo mass marginalization procedure is analytic, unlike the simulation-based method employed in [341], due to the limited statistics of MW-like simulations that include realistic LMC analogs. The significant improvements in sensitivity to non-CDM physics afforded by modeling the LMC satellite system further reinforce the importance of simulation suites of MW-like systems including realistic LMC analog systems.

On the strong lensing side, the [195] study upon which we base our analysis uses flux ratio statistics that are significantly more constraining than the gravitational imaging data underlying the [161] joint analysis. This additional constraining power results from the fact that current gravitational imaging data probes $\sim 10^9 M_\odot$ subhalos while flux ratio anomalies are sensitive to the presence of lower-mass subhalos. In terms of modeling, [195] explicitly account for the host mass and redshift dependence of the SHMF using **Galacticus**—these are leading-order effects in predicting the SHMF for a given lens and its lens-to-lens variation—while the [455] and [387] analyses that the joint constraints in [161] are based on do not. In addition, [195] self-consistently account for the reduction in halo concentration in WDM, which significantly increases the sensitivity of lensing observations to WDM effects and also models the effects of tidal stripping on subhalos after infall, which is again crucial to accurately forward-model flux ratio observations.

Finally, we emphasize the following key aspects of our probe combination relative to the procedure in [161], which combines fully marginalized one-dimensional M_{hm} distributions from various small-scale structure probes including MW satellites and gravitational imaging to derive joint WDM constraints:

1. We cast the subhalo populations inferred from MW satellites and from the group-mass, $z \sim 0.5$ host halos probed by strong lensing into a common, multidimensional parameter space of projected subhalo number density Σ_{sub} versus WDM half-mode mass M_{hm} ;

2. We combine these $\Sigma_{\text{sub}}-M_{\text{hm}}$ distributions to construct a joint marginal likelihood that is strictly more constraining and informative than the joint M_{hm} distribution resulting from fully marginalizing over all additional parameters (see the right panel of Figure 11.6), improving the precision of our joint analysis; and
3. We model the differential efficiency of subhalo disruption due to the central galaxies in the different host halo mass and redshift regimes probed by MW satellites and strong lensing, finding that our results are robust to uncertainties in these effects at the $\sim 10\%$ level, which lends confidence to the robustness of our results.

The differences in the underlying data used in our inference—and particularly the inclusion of LMC-associated satellites in the MW satellite analysis and the use of strong lensing flux ratio statistics that probe lower-mass subhalos than current gravitational imaging data—therefore result in more precise joint constraints than those obtained in [161] and allow us to significantly improve upon their WDM limit. Moreover, the joint analysis choices described above lend to the robustness and accuracy of our results.

11.9 Conclusions

In this paper, we performed a multidimensional joint analysis of the distribution of small-scale dark matter structure inferred from MW satellite galaxies and strong gravitational lensing. In particular, we combined state-of-the-art dark matter substructure measurements derived from (i) the MW satellite galaxy population over $\sim 75\%$ of the sky and (ii) the flux ratio statistics and image positions from eight quadruply imaged quasars. By combining constraints on the projected subhalo number density and the half-mode mass describing the suppression of the subhalo mass function in thermal relic WDM, we improved lower limits on the WDM particle mass derived independently by breaking degeneracies among the inferred subhalo distributions at each scale for the first time. Our m_{WDM} constraint is more stringent than any limit set by independent analyses of small-scale structure probes to date.

Our key results are summarized below:

1. Our multidimensional joint analysis extracts information that was not accessed by MW satellite or strong lensing analyses independently, improving WDM constraints by $\sim 30\%$, with $M_{\text{hm}} < 10^{7.0} M_{\odot}$ ($m_{\text{WDM}} > 9.7 \text{ keV}$) at 95% confidence, or $M_{\text{hm}} = 10^{7.4} M_{\odot}$ ($m_{\text{WDM}} = 7.4 \text{ keV}$) disfavored with a 20:1 marginal likelihood ratio. (Figures 11.4–11.5);
2. Our joint WDM constraint is robust to uncertainties in the differential efficiency of subhalo disruption between the MW and strong lens host halo mass and redshift regimes at the $\sim 10\%$ level;

3. Projected subhalo number density constraints from MW satellites and strong lensing flux ratio statistics are in mild tension but are sensitive to uncertainties in the efficiency of subhalo disruption in the corresponding host halo mass and redshift regimes;
4. We discuss key systematics that are conservatively marginalized over in the current analysis but which must be mitigated in future work to claim a detection of non-CDM physics from small-scale structure measurements. These systematics include the line-of-sight contribution to the strong lensing signal, the differential efficiency of subhalo disruption due to baryons at the MW and lensing host halo mass and redshift scales, and the properties of the MW and strong lens host halos (Figure 11.6);
5. Inferences of the small-scale dark matter structure from MW satellites and strong lensing are consistent despite the completely different nature of these probes and differences in their corresponding host halo mass and redshift regimes.

Recent studies have identified a variety of microphysical dark matter properties that suppress small-scale structure in a manner quantitatively similar to WDM, including the strength of velocity-independent interactions between dark matter and protons (see Chapter 6 [335]), the production mechanism of nonthermal dark matter in early matter-dominated cosmologies [322], and the dark matter formation redshift in models of “late-forming” dark matter [123]. Our jointly derived WDM constraints directly inform all of these properties. Dark matter models that feature qualitatively different suppression of small-scale structure compared to WDM can also be constrained by constructing a conservative mapping; for example, such mappings have been applied to constrain fuzzy dark matter [418], models with velocity-dependent self- and Standard Model dark matter interactions [448, 308], and models within the ETHOS framework [75]. Such dark matter physics may manifest differently in small-scale structure probes like MW satellites and strong lensing that are sensitive to halo abundances and concentrations in unique ways, and we regard this as a particularly compelling avenue for future work.

We expect the relative improvement offered by our probe combination to continue to increase as both techniques progress due to both additional data from existing instruments and next-generation observational facilities. Excitingly, the sample sizes of both nearby ultrafaint dwarf galaxies and quadruply lensed quasars are expected to drastically increase with LSST [237], Euclid Space Telescope [285], and Nancy Grace Roman Space Telescope [431] observations. Forthcoming facilities including the Maunakea Spectroscopic Explorer [438] will also help to confirm the nature of candidate MW satellites and faint dwarf galaxies throughout the Local Volume, while wide-aperture and extremely large telescopes (ELTs) will provide detailed information about the dynamical masses of these systems, which is key to refine galaxy–halo connection and WDM constraints [423]. Meanwhile, the unprecedented sample of strong lenses expected to be discovered within 5–10 yr will yield precise measurements of the differential line-of-sight and substructure contributions to lensing signal and

will allow the selection functions of strong lenses to be better quantified. Observations of extended source emission will also help constrain lens macromodels.

With sufficiently stringent limits on the minimum luminous halo mass from nearby dwarf galaxies and on the mass scale of a cutoff in the subhalo mass function, our procedure for combining satellite galaxy and strong lensing posteriors can potentially provide evidence for the existence of dark subhalos—i.e., subhalos devoid of observable baryonic components—which are a key, unverified prediction of many viable dark matter models. [333] estimate that the lowest-mass halo expected to host a dwarf galaxy is more massive than $\sim 10^7 M_\odot$. Thus, the future observations discussed above, which are expected to constrain the subhalo mass function at and below these mass scales, will either yield evidence for a cutoff in galaxy (or halo) formation or evidence for halos devoid of observable baryonic matter. We plan to pursue these measurements by developing a multipronged theoretical framework to jointly infer the distribution of small-scale structure using heterogeneous data.

Chapter 12

Forecasts for Future Surveys

Abstract

Forthcoming cosmic surveys promise to sharpen the galaxy–halo connection and dark matter constraints from nearby dwarf galaxies presented in the preceding chapters. The Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST) will be particularly sensitive to ultra-faint dwarf galaxies within the virial radius of the Milky Way halo and potentially throughout the Local Volume. This chapter forecasts the galaxy formation and dark matter constraints that will be enabled by LSST, summarized respectively by the minimum halo mass corresponding to an observable dwarf galaxy and the warm dark matter half-mode mass that describes the suppression of the halo and subhalo mass functions. We predict that the LSST-discovered satellite population can probe galaxy occupation in halos with peak virial masses down to $\sim 5 \times 10^7 M_\odot$ and half-mode mass scales down to $\sim 10^7 M_\odot$, where the extra sensitivity to the half-mode scale is driven by the population statistics of the faintest Milky Way satellites.

12.1 Paper Status and External Contributions

Section 12.2 appears in slightly modified form in the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST) dark matter white paper [151] with the title, “Probing the Fundamental Nature of Dark Matter with the Large Synoptic Survey Telescope,” on which I am a co-author. It is the result of an effort that I led within the LSST Dark Matter Working Group with contributions from Keith Bechtol, Alex Drlica-Wagner, Mitch McNanna, Andrew Pace, Yao-Yuan Mao, Erik Tollerud, Risa Wechsler, Francis-Yan Cyr-Racine, Mei-Yu Wang, Kimberly Boddy, and Arka Banerjee. Keith and Alex provided helpful input on the LSST dwarf galaxy detection sensitivity forecasts, and all of the coauthors made editorial contributions to the text. Section 12.3 is unpublished and authored by me.

12.2 Forecasts for the Minimum Halo Mass Inferred from LSST Milky Way Satellite Observations

The least luminous galaxies currently known contain only a few hundred stars and have been found exclusively in the inner regions of the Milky Way due to observational selection effects. Although the census of Milky Way dwarf galaxies has grown from ~ 25 to more than 50 in recent years (e.g., with the Dark Energy Survey; [42, 272, 150], our current census is certainly incomplete. For example, the HSC-SSP Collaboration has detected two ultra-faint galaxy candidates in the first 300 deg^2 of the survey [222, 220]; these galaxies are faint and distant enough to have been undetectable in previous optical imaging surveys. HSC is representative of the depth that will be achieved by the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST) over half the sky—an area 60 times larger than the current HSC-SSP footprint. Thus, based on the results of SDSS, HSC, DES, etc., several groups have predicted that LSST could detect tens to hundreds of new low-luminosity Milky Way satellites, mainly at larger distances and fainter luminosities than those accessible with current-generation surveys [271, 440, 214, 340, 241, 336, 264]. In addition, novel techniques, such as the use of the correlated phase space motions of stars [29, 112] or clustering of variable stars [37] could further expand the sample of ultra-faint galaxies. LSST observations of Milky Way satellites therefore offer an exciting testing ground for dark matter models; for example, the measured abundance, luminosity function, and radial distribution of Milky Way satellites *already* place competitive constraints on warm dark matter particle mass at the level of $\sim 2\text{--}5 \text{ keV}$ (e.g., [241, 264]).

To relate these questions to LSST observations, we have analyzed simulated ultra-faint galaxies as they would appear in LSST WFD coadd object catalogs to quantify LSST’s ability to detect nearby satellite galaxies. We detect ultra-faint galaxies as arcminute-scale statistical overdensities of individually resolved stars; in ground-based optical imaging surveys, it is often challenging to classify low signal-to-noise catalog objects near the detection threshold as either foreground stars or unresolved background galaxies. LSST will reach depths at which the galaxy counts far outnumber stellar counts, so the search sensitivity for ultra-faint galaxies will largely be determined by our ability to accurately perform star-galaxy separation at magnitudes $24 < r < 27.5$; importantly, our sensitivity analyses include these effects. We find that Milky Way satellites within 300 kpc are well-detected with a surface brightness detection threshold of $\mu = 32 \text{ mag arcsec}^{-2}$ [237] and an absolute magnitude cutoff of $M_V = 0 \text{ mag}$.

Figure 12.1 shows the minimum subhalo mass that LSST can probe via observations of Milky Way satellites, obtained by folding our search sensitivity estimates through a cosmological model of the Milky Way satellite population that predicts satellite luminosity functions, radial distributions, and size distributions that agree well with current observations. In particular, we generate many mock Milky Way satellite populations using the model presented in Chapter 3 [336] given a “true”

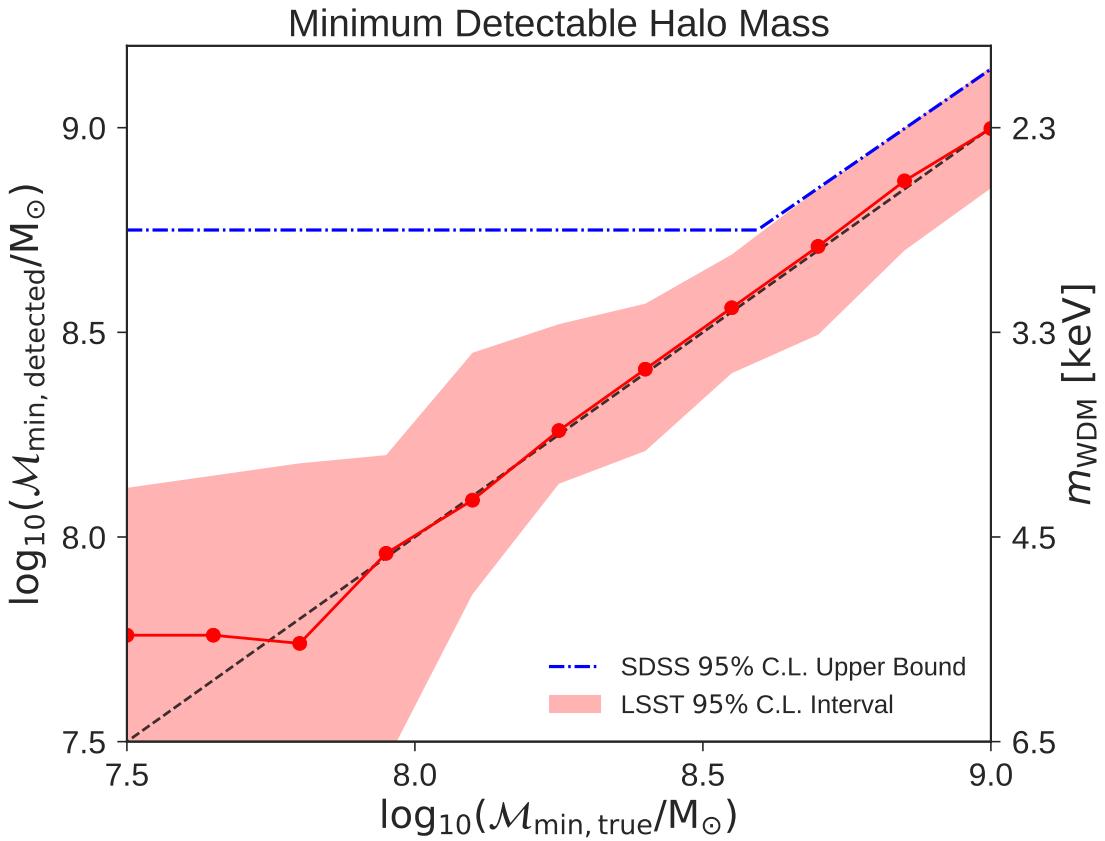


Figure 12.1: Forecast for the minimum dark matter subhalo mass probed by LSST via observations of Milky Way satellites. The red band shows the 95% confidence interval from our MCMC fits to mock satellite populations as a function of the true peak subhalo mass necessary for galaxy formation. Note that we marginalize over the relevant nuisance parameters associated with the galaxy–halo connection—including the effects of baryons using the Chapter 2 model calibrated on subhalo disruption in hydrodynamic simulations [336]—in our sampling. We indicate the corresponding constraints on the warm dark matter mass assuming $M_{\text{hm}} = M_{\text{min}}$.

value of the minimum peak subhalo virial mass necessary for galaxy formation, $\mathcal{M}_{\min, \text{true}}$, marginalizing over the remaining galaxy–halo connection parameters. We then perform mock observations of these generated satellite populations using the LSST selection function, and we compare these to the true satellite populations by MCMC sampling \mathcal{M}_{\min} and the remaining galaxy–halo connection parameters assuming that satellite number counts are Poisson distributed in bins of absolute magnitude (see Chapter 3 [336] for details on the fitting procedure). For each value of $\mathcal{M}_{\min, \text{true}}$, this procedure yields a posterior distribution over the minimum halo mass inferred by LSST observations. The red band in Figure 12.1 illustrates the recovered 95% confidence interval as a function of $\mathcal{M}_{\min, \text{true}}$, and the blue dot-dashed line indicates the minimum halo mass inferred from known classical and SDSS-detected Milky Way satellites. For small $\mathcal{M}_{\min, \text{true}}$, the 95% confidence level upper bound on the lowest detectable subhalo mass improves by a factor of ~ 5 with LSST, from $\sim 5 \times 10^8 \text{ M}_\odot$ to $\sim 10^8 \text{ M}_\odot$; this translates to a conservative lower bound of $\sim 5 \text{ keV}$ on the WDM particle mass.

Although we have presented a “population-based” forecast for dark matter constraints from LSST-detected ultra-faint satellites, we note that kinematic data obtained by follow-up spectroscopy of newly discovered satellites also offers a powerful probe of dark matter microphysics. We estimate the number of LSST-detected Milky Way satellites that can be spectroscopically confirmed in Section 4.1 of [151], and we forecast the constraints offered by these stellar velocity dispersion measurements for WDM and SIDM in Section 3.1.4 of [151].

Further extending the sensitivity of LSST to a power spectrum cut-off on scales smaller than the mass threshold for galaxy formation requires techniques that are independent of satellite luminosity and that can detect subhalos purely through their gravitational signatures.

12.3 Half-mode Mass Forecasts Including the Large Magellanic Cloud Satellite Population

The warm dark matter sensitivity forecasts in Section 12.2 are extremely conservative because they assume that the half-mode mass scale describing the suppression of the halo and subhalo mass functions is equal to the minimum observable halo mass inferred from dwarf galaxy observations. However, as demonstrated in Chapter 7 [332], the half-mode mass constraints derived from Dark Energy Survey and Pan-STARRS1 Milky Way satellite population data exceeds this expectation, with $M_{\text{hm}} \lesssim 4 \times 10^7 \text{ M}_\odot$ despite the upper limit of $\mathcal{M}_{\min} \lesssim 3 \times 10^8 \text{ M}_\odot$ derived from the same data. This improvement is driven by the increased number of observed dwarf galaxies inferred to inhabit subhalos with peak virial masses near \mathcal{M}_{\min} relative to the classical and Sloan Digital Sky Survey-discovered satellite population. In turn, these population statistics raise the confidence that models resulting in even moderate subhalo mass function suppression near the minimum observable halo mass are ruled out.

To include the impact of these population statistics and the Large Magellanic Cloud satellite population in our forecasts, we re-run the analysis described in Section 12.2 using the satellite model from Chapters 4–7 [333, 332]. Thus, these updated forecasts include a smoothly varying galaxy occupation fraction (parameterized by the peak virial mass at which 50% of subhalos host satellites, M_{50}) and thermal relic warm dark matter subhalo mass function suppression (parameterized by M_{hm}), along with the comprehensive models for satellite luminosity, size, and disruption described in Chapter 4.

Figure 12.2 summarizes the results of these sensitivity analyses, which again assume a limiting surface brightness detection threshold of $\mu = 32$ mag arcsec $^{-2}$ and absolute magnitude cutoff of $M_V = 0$ mag. In particular, Figure 12.2 shows how the 95% confidence upper limits inferred on M_{50} and M_{hm} scale as our mock satellite observations approach LSST sensitivity. Given the upper limits on M_{50} and M_{hm} derived from current Milky Way satellite data, there is only a small amount of parameter space in which LSST Milky Way satellite observations *alone* can distinguish a cutoff in galaxy formation from a cutoff in halo abundances due to dark matter physics (Region I. of Figure 12.2). In the remaining parameter space (Region II. of Figure 12.2), LSST satellite observations alone cannot uniquely determine the astrophysical or dark matter origin of a cutoff in dwarf galaxy abundances. Nonetheless, these observations are expected to probe the abundance of halos with peak virial masses down to $\sim 5 \times 10^7 M_\odot$, improving the sensitivity of current minimum halo mass measurements by nearly an order of magnitude. These promising forecasts motivate further advances in galaxy–halo connection modeling and simulations of Milky Way-like systems in preparation for LSST data.

The lower limit of the horizontal axis on Figure 12.2 corresponds to the minimum halo mass that can *in principle* host an observable dwarf galaxy as derived in Chapter 4. Thus, the detection of halo abundances consistent with CDM predictions below this mass scale would provide decisive evidence for the existence of completely dark halos and must be enabled by probes that do not rely on the luminous content of these systems (e.g., see Chapter 11).

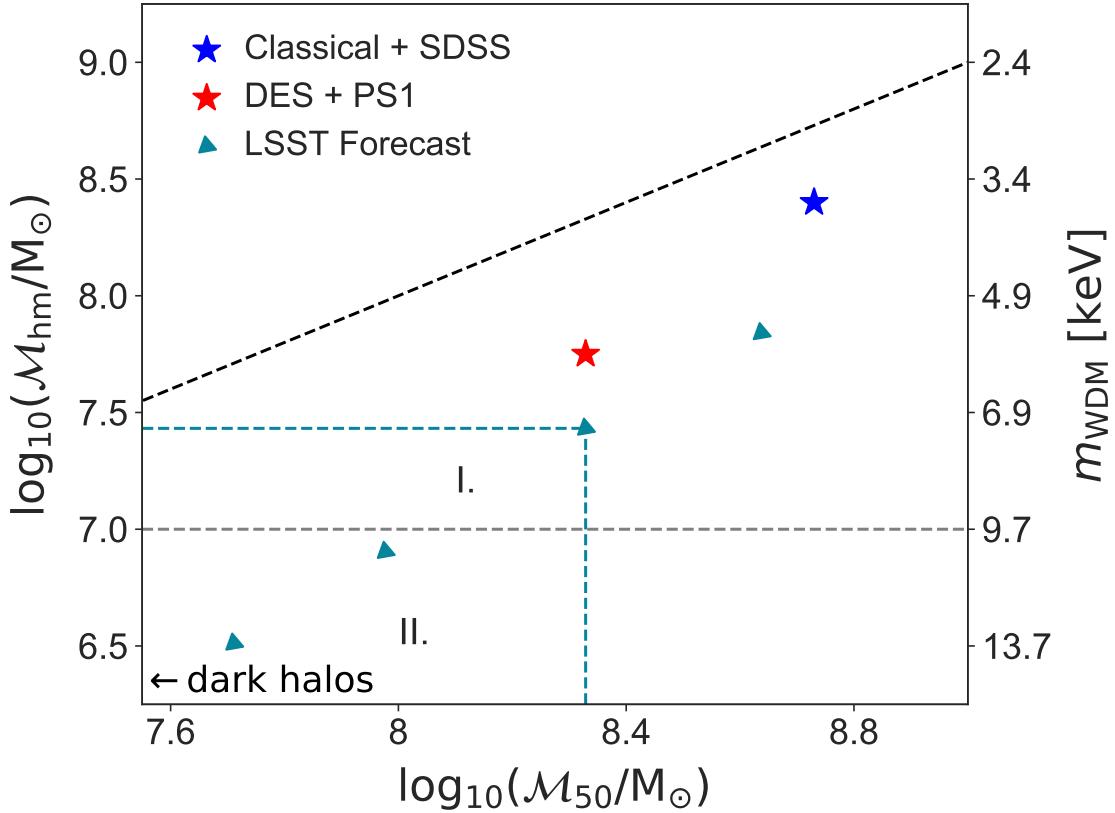


Figure 12.2: Forecasts for the galaxy formation and dark matter sensitivity of LSST Milky Way satellite observations using the satellite model from Chapters 4 and 7, which includes realistic Large Magellanic Cloud analog systems. The x -axis shows the 95% confidence upper limit on M_{50} and the y -axis shows the 95% confidence upper limit on M_{hm} for warm dark matter from classical and SDSS-discovered satellites (blue star; Chapter 3 [336]), DES and Pan-STARRS1 satellites (red star; 4 [333]), and mock Milky Way satellite observations as they approach LSST sensitivity (blue-green triangles). The black dashed line shows the one-to-one relation and the gray dashed line shows the current lower limit on the warm dark matter mass from our joint Milky Way satellite and strong lensing analysis in Chapter 11. A cutoff in satellite abundances in Region I. must be astrophysical in origin, a cutoff in Region II. can be caused by either galaxy formation or dark matter physics, and the detection of halos below the lower limit of the M_{50} axis at $\sim 5 \times 10^7 M_{\odot}$ would provide evidence for the existence of completely dark halos.

Chapter 13

Conclusions

In this thesis, I addressed small-scale dark matter and galaxy formation physics by modeling the connection between faint dwarf galaxies and the low-mass dark matter halos they inhabit. Understanding these nonlinear systems entails cosmology, galaxy formation, high-energy particle physics, and statistical methods to be tied together by complex numerical simulations. Despite the computational nature of this work, it is motivated by the principles of effective field theory, which guarantee that the microphysical nature of dark matter is reflected in its behavior at progressively higher energy scales and shorter length scales. The small-scale modes that seed the formation of dwarf galaxy halos enter the cosmological horizon within months of the Big Bang—well before the epoch of the cosmic microwave background—meaning that dwarf galaxies trace both early Universe physics (through their connection to small halos) and galaxy formation over cosmic time.

In Chapters 2–5, I presented an empirical model for the connection between faint galaxies and small dark matter halos applicable to both Milky Way-mass systems and to the Milky Way specifically. In particular, Chapter 2 developed a machine-learning algorithm trained on hydrodynamic simulations to capture the disruptive effects of Milky Way-like central galaxies on subhalo populations. Chapter 3 integrated this subhalo disruption algorithm with a new, flexible galaxy–halo connection model based on abundance matching to predict the relationship between satellite galaxy luminosities and sizes and their subhalo properties. Crucially, this model includes a rigorous statistical framework to facilitate comparisons between its predictions and observed satellite populations. As demonstrated in Chapter 3 classical and Sloan Digital Sky Survey-discovered Milky Way satellite galaxies, which have been known for more than ten years, place interesting constraints on the faint-end galaxy–halo connection and provide strong evidence for the existence of halos with peak virial masses below $\sim 5 \times 10^8 M_\odot$ in abundances consistent with cold dark matter predictions.

Chapter 4 combines the model described above with (i) tailored cosmological simulations of Milky Way-like systems including realistic Large Magellanic Cloud analogs, (ii) recent satellite galaxy observations from the Dark Energy Survey and Pan-STARRS1, and (iii) state-of-the-art observational

selection functions for these galaxies constrained directly by photometric data from each survey. These observations of the Milky Way satellite population over nearly the full sky tighten constraints on the galaxy–halo connection and minimum halo mass while providing compelling statistical evidence for the existence of Large Magellanic Cloud satellites. In the dark matter context, these satellites-of-satellites occupy second-order subhalos in a beautiful signature of hierarchical structure formation. This work also places the first robust limits on the fraction of low-mass dark matter halos that host dwarf galaxies, which is consistent with 100% down to the minimum observable halo mass of $\sim 3 \times 10^8 M_\odot$.

To close the first part of this thesis, Chapter 5 takes the first steps towards placing the galaxy–halo connection measurements derived from Milky Way satellites in a cosmological context. The Milky Way satellite luminosity function and the galaxy–halo connection model constrained only by Milky Way satellite data are consistent with measurements of satellite populations around 36 Milky Way analogs in the low-redshift Universe from the Satellites Around Galactic Analogs survey. Forthcoming work that derives joint constraints using Milky Way satellites, Local Volume satellite data, and the full sample of 100 complete SAGA satellite systems promises to sharpen this comparison.

The second part of this thesis (Chapters 6–11) uses the framework described above to explore and constrain dark matter physics that suppresses the present-day abundance of small halos. Chapter 6 derives an analytic relation between the cross section of dark matter–Standard Model interactions and the minimum halo mass. Further, it numerically demonstrates that the suppression of the linear matter power spectrum for velocity-independent dark matter–proton scattering is similar to that for thermal relic warm dark matter, and it exploits this correspondence in the galaxy–halo connection model from Chapter 3 to place the first constraints on these interactions using dwarf galaxies. These constraints complement terrestrial direct detection experiments and improve upon limits from other cosmological probes by several orders of magnitude.

Chapters 7 and 8 constrain the warmth, Standard Model interactions, minimum particle mass, and formation epoch of dark matter using the Milky Way satellite model from Chapter 4, which includes realistic Large Magellanic Cloud analogs. The warm, interacting, and late-forming dark matter constraints all follow from a 6.5 keV lower limit on the mass of thermal relic warm dark matter derived from Dark Energy Survey and Pan-STARRS1 satellite abundances. This limit implies that modes entering the cosmological horizon more than one week after the Big Bang cannot be significantly suppressed relative to cold dark matter predictions by *any* kind of dark matter micro-physics. Chapter 7 also translates these constraints into parameter spaces relevant for specific dark matter particle models such as sterile neutrinos and ultra-light axions.

Next, Chapter 9 studies the signatures of dark matter self-interactions on subhalos in Milky Way-mass systems, and Chapter 10 highlights the signatures that are unique to Milky Way-like systems with Large Magellanic Cloud analogs. Although the self-interacting dark matter literature largely focuses on its impact on dwarf galaxy halo density profiles, Chapter 9 reveals that subhalo

disruption is sensitive to self-interactions. In particular, the momentum transfer cross section at the orbital velocity scale set by the dark matter host halo dictates the strength of ram-pressure stripping and the resulting efficiency of subhalo disruption; in turn, the impact of this process on subhalo abundances is tied to the velocity-dependence of the self-interaction cross section. Chapter 10 shows that the recent, high-velocity infall of the Large Magellanic Cloud system enhances late-time subhalo disruption due to dark matter self-interactions, opening up new avenues to disentangle different kinds of dark matter and baryonic physics that impact subhalo populations.

Finally, Chapter 11 describes and implements a unified framework that combines dark matter constraints from Milky Way satellite galaxy abundances and strong gravitational lensing flux ratio statistics. This probe combination significantly improves dark matter constraints set independently, resulting in a lower limit on the thermal relic warm dark matter mass of ~ 10 keV, and breaks degeneracies among the subhalo populations inferred by each technique. Based on the plethora of dwarf galaxy and strong lens discoveries expected from next-generation surveys including the Vera C. Rubin Observatory Legacy Survey of Space and Time (LSST), this work paves the way for next-generation surveys of small-scale cosmic structure to deliver unprecedented insights into dark matter physics. Specifically, Chapter 12 shows that upcoming LSST observations will likely either provide evidence for the existence of halos without any luminous content or detect a deviation from the cold dark matter paradigm at the corresponding mass scales.

Observational evidence for the existence and properties of dark matter has always run through its connection to the visible Universe. This is clearly illustrated in its historical underpinnings, including in early studies by Zwicky [504] and Rubin [397], which respectively demonstrate that the dynamics of cluster and spiral galaxies require the existence of invisible mass. Dwarf galaxies, as the most dark matter-dominated systems currently known, therefore provide compelling laboratories for precision tests of dark matter physics. The promise and the peril of using these faint galaxies to test dark matter and galaxy formation has played out in response to claims of “missing satellites” over the last two decades. Fortunately, we are now reaching the necessary observational, theoretical, and computational precision to deliver on the promise of these remarkable systems and to robustly push the boundaries of dark matter and galaxy formation theory.

Appendix A

Random Forest Scatter, Feature Selection, & Resolution

We perform several tests to check the robustness of our results in Chapter 2. First, we examine the scatter in the random forest predictions for the `m12i` and `m12f` V_{\max} functions and whether the details of the training data affect our results. Figure A.1 shows the V_{\max} functions from 200 realizations of our fiducial classifier, along with the most probable realization of classifiers trained only on subhalos from `m12i` or `m12f` with $V_{\text{peak}} > 10 \text{ km s}^{-1}$. The scatter about the most probable prediction for our fiducial classifier is small; in particular, the intrinsic scatter of the random forest prediction is comparable to or smaller than the Poisson noise over the entire velocity function for each host. Thus, even though the prediction for the total number of surviving subhalos is different for classifiers trained on `m12i` or `m12f` separately, this uncertainty does not propagate to our fiducial classifier.

Next, we explore the choice of training features. In particular, we test how adding subhalo features affects our results for the V_{\max} functions and radial distributions of the surviving subhalo populations predicted from DMO simulations of `m12i` and `m12f`. In Figure A.2, we plot the most probable V_{\max} functions and radial distributions predicted by five classifiers that each use an additional training feature, corresponding to the rows of Table 2.1 and the columns of Figure 2.4. As we add subhalo features, the predicted distributions approach the FIRE results. Interestingly, d_{peri} (or a_{peri} or a_{acc}) alone provides most of the information needed to match the total number of surviving subhalos with $V_{\text{peak}} > 10 \text{ km s}^{-1}$ and $r < 300 \text{ kpc}$, but adding additional features improves the predictions at large V_{\max} and small radii.

Finally, we study how our results depend on the resolution limits used for the training data. In Figure A.3, we show the `m12i` and `m12f` velocity functions and radial distributions predicted by a classifier trained on subhalos from both hosts with $V_{\text{peak}} > 5 \text{ km s}^{-1}$, which is less restrictive

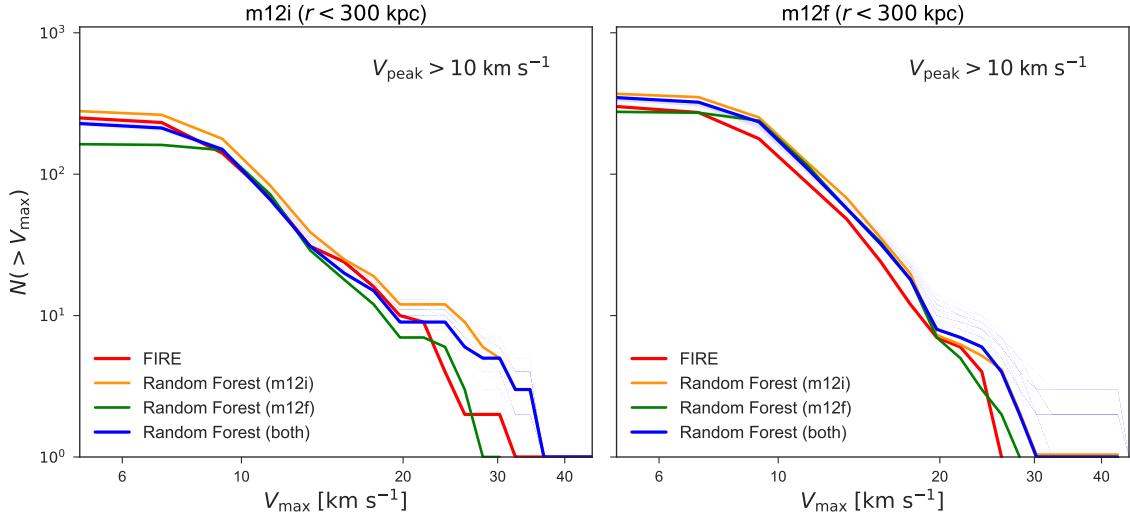


Figure A.1: Velocity functions for m12i (left) and m12f (right) predicted by the most probable realization of our random forest classifier when trained only on subhalos from m12i (orange) or m12f (green) with $V_{\text{peak}} > 10 \text{ km s}^{-1}$. Blue lines show 200 realizations of the prediction for our fiducial classifier, which is trained on subhalos from both hosts, and red lines show the FIRE results. While there is a difference between the total number of surviving subhalos predicted by classifiers trained only on m12i or m12f , the scatter about the most probable prediction for our fiducial classifier is small.

than the $V_{\text{peak}} > 10 \text{ km s}^{-1}$ cut used in our primary analysis. We plot the results for subhalos with $V_{\text{max}} > 5 \text{ km s}^{-1}$, where V_{max} is the maximum circular velocity evaluated at $z = 0$, which allows for a direct comparison with the results in GK17. Our conclusions are unaffected by changing the minimum circular velocity. In fact, our predictions match the hydrodynamic results even more closely than before in the low- V_{max} regime, since this less restrictive cut significantly increases the number of subhalos at the low- V_{max} end of the training set. Thus, our classifier can be applied to simulations with a range of resolution thresholds if appropriate cuts are applied to the training data.

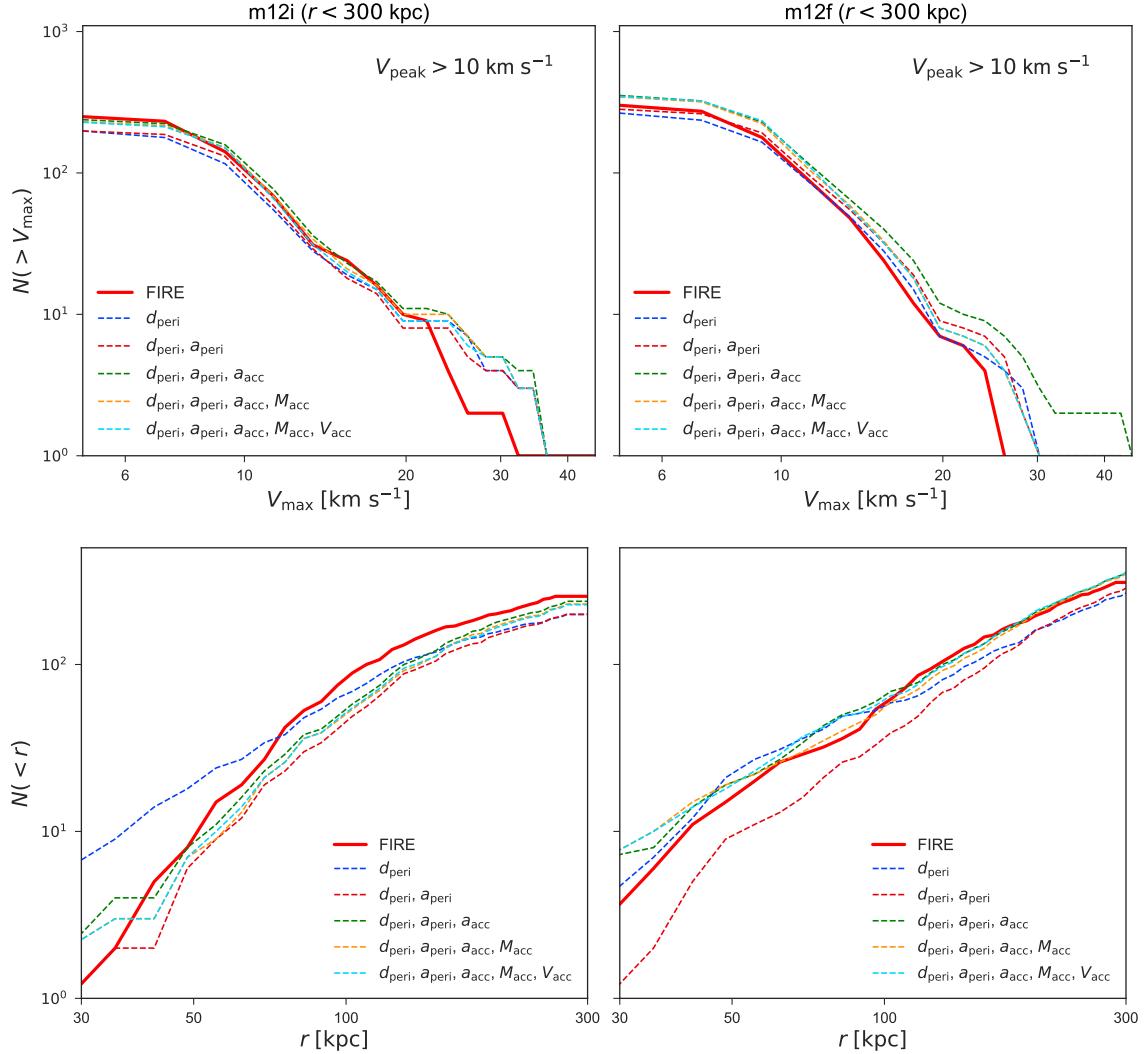


Figure A.2: Velocity functions and radial distributions of subhalos in m12i (left) and m12f (right) predicted by the most probable realization of random forest classifiers trained on subhalos from both hydrodynamic simulations with $V_{\text{peak}} > 10 \text{ km s}^{-1}$. The classifiers use the features d_{peri} (blue); d_{peri} and a_{peri} (red); d_{peri} , a_{peri} , and a_{acc} (green); d_{peri} , a_{peri} , a_{acc} , and M_{acc} (orange); and d_{peri} , a_{peri} , a_{acc} , M_{acc} , and V_{acc} (cyan), corresponding to the rows of Table 2.1 and the columns of Figure 2.4. The solid red lines show the FIRE results.

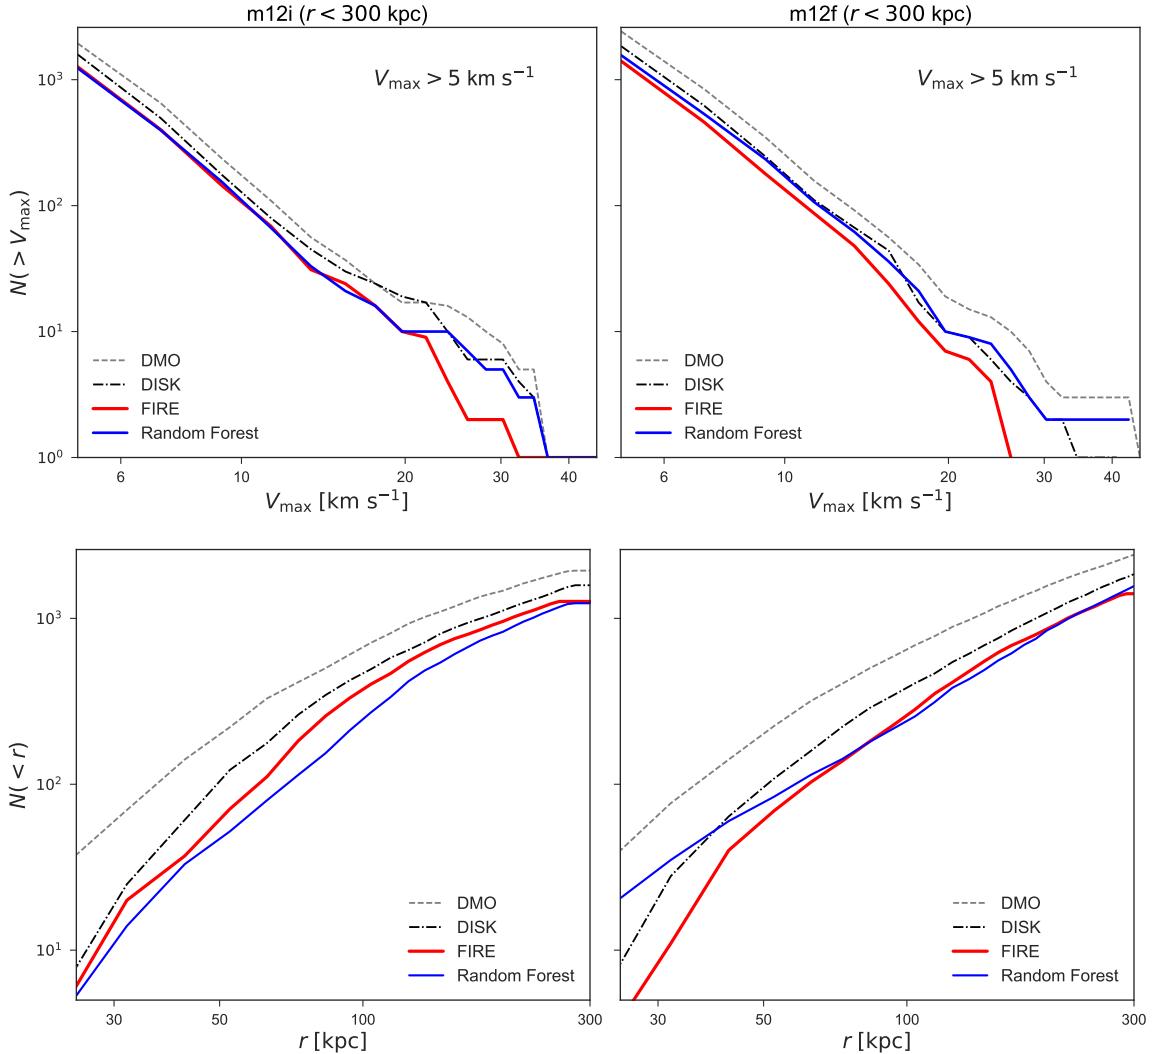


Figure A.3: Velocity functions (top) and radial distributions (bottom) for subhalos hosted by m12i (left) and m12f (right), predicted by the most probable realization of our random forest classifier trained on subhalos from m12i and m12f with $V_{\text{peak}} > 5 \text{ km s}^{-1}$ (blue). The FIRE (red), DISK (dot-dashed), and DMO (dashed) results are shown for comparison. We restrict these plots to subhalos within 300 kpc of their respective hosts at $z = 0$ and with $V_{\text{max}} > 5 \text{ km s}^{-1}$, where V_{max} is the maximum circular velocity evaluated at $z = 0$, to allow for a direct comparison with the results in [190].

Appendix B

Satellite Model Statistical Methods

This Appendix supplements the statistical methods underlying the galaxy–halo connection model presented in Chapter 3.

B.1 Convergence of Improved Likelihood to the Poisson Distribution

The main assumption underlying the derivation of our likelihood (Equation 3.13) is that the number of observed satellites and mock satellites in a given magnitude bin are drawn from the same Poisson distribution. Thus, in the limit of many mock observations, any likelihood that marginalizes over an unknown rate parameter should converge to the underlying Poisson distribution. This holds for the likelihood used in our analysis: in particular, if $\hat{n}_{i,1}, \dots, \hat{n}_{i,\hat{N}} \sim \text{Poisson}(\lambda_i)$, then

$$\lim_{\hat{N} \rightarrow \infty} P(n_i | \hat{n}_{i,1}, \dots, \hat{n}_{i,\hat{N}}) = \frac{e^{-\lambda_i} \lambda_i^{n_i}}{n_i!} = \text{Poisson}(\lambda_i), \quad (\text{B.1})$$

where the scaling with n_i follows from Equation 3.13 using $\hat{n}_{i,1} + \dots + \hat{n}_{i,\hat{N}} = \hat{N} \lambda_i$ for large \hat{N} , and the prefactor follows from the fact that $P(n_i | \hat{n}_{i,1}, \dots, \hat{n}_{i,\hat{N}})$ is a normalized distribution or by taking the $\hat{N} \rightarrow \infty$ limit of Equation 3.13.

We illustrate the fact that our likelihood satisfies this property in Figure B.1. We compare our likelihood to the Poisson distribution that the $\hat{n}_{i,j}$ are sampled from for various values of λ_i , n_i , and \hat{N} . We also plot the likelihood used in [241], which is equivalent to the mean of $P(n_i | \hat{n}_{i,j})$ over the realizations $j = 1, \dots, \hat{N}$, and we note that it does *not* converge to the underlying Poisson distribution in the limit of large \hat{N} . We verify that this is the expected behavior by computing central moments of $P(n_i | \hat{n}_{i,j})$ for $\hat{n}_{i,j} \sim \text{Poisson}(\lambda_i)$ and using these quantities to derive central moments of $P(n_i | \hat{n}_{i,j})$ averaged over many mock observations $j = 1, \dots, \hat{N}$. For example, our

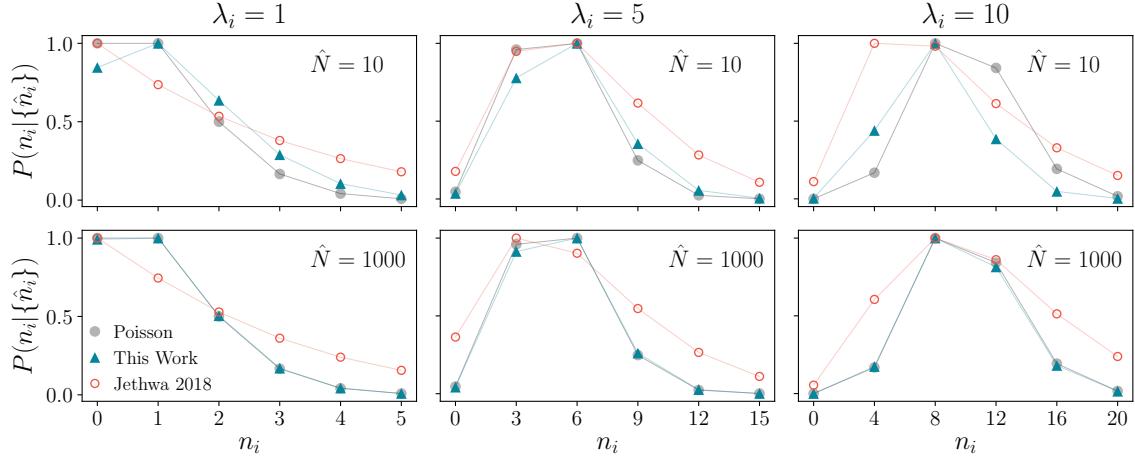


Figure B.1: Likelihood of observing n_i counts given \hat{N} mock observations $\hat{n}_{i,1}, \dots, \hat{n}_{i,\hat{N}}$ assuming that n_i and all \hat{n}_i are drawn from a Poisson distribution with rate parameter λ_i . We show results for $\hat{N} = 10$ (top row) and $\hat{N} = 1000$ (bottom row) mock observations given true rate parameters $\lambda_i = 1$ (left column), $\lambda_i = 5$ (middle column), and $\lambda_i = 10$ (right column), computed using (i) the Poisson likelihood given the true rate parameter (gray points); (ii) the likelihood used in our analysis (Equation 3.13), which marginalizes over λ_i given all \hat{n}_i simultaneously (blue triangles); and (iii) the likelihood used in [241], which averages the likelihoods obtained from multiple mock observations (red circles). Our likelihood converges to the underlying Poisson distribution in the limit of many mock observations, while the averaged version does not. Note that we rescale each version of $P(n_i | \{\hat{n}_i\})$ by its maximum value in every panel.

analytic calculations yield $\lambda_i + 1$ and $3\lambda_i + 2$ for the mean and variance of the averaged likelihood (versus λ_i for both quantities in the Poisson distribution), which agree with our numerical tests to high precision. Heuristically, this occurs because $P(n_i | \hat{n}_{i,j})$ is broader than (and biased with respect to) $\text{Poisson}(\lambda_i)$, so an averaged version of this likelihood cannot converge to the underlying Poisson distribution.

B.2 Bayesian Evidence for Radial Scaling, Size Reduction, and Orphan Satellites

Here we explore whether our radial scaling, satellite size reduction, and orphan models — which were held fixed in the preceding analysis — are favored by the classical-plus-SDSS luminosity distribution by computing Bayes factors for fits with and without these effects. In particular, we compare the Bayesian evidence for fits with $\chi = 0.8$ (satellite radii scaled inward by a factor of 0.8 relative to subhalo radii) versus $\chi = 1$ (satellite radii set equal to subhalo radii), $\beta = 0$ (satellite sizes set at accretion) versus $\beta = 1$ (satellite sizes reduced based on tidal stripping), and $\mathcal{O} = 0$ (no orphan

satellites) versus $\mathcal{O} = 1$ (fiducial orphan model). In all three cases, we vary α , σ_M , \mathcal{M}_{\min} , and \mathcal{B} , and we fix all remaining parameters, except for either χ , β , or \mathcal{O} , according to Table 3.1.

We calculate the evidence for each model using the bounded harmonic mean estimate, which effectively averages the inverse product of the likelihood and the prior over Markov Chain samples drawn from an ellipsoid in a high-density region of the posterior distribution. In particular, we select samples of the free parameters $\boldsymbol{\theta}$ within a fixed Mahalanobis distance [313] of a point $\boldsymbol{\theta}_0$ in a high posterior density (HPD) region according to

$$(\boldsymbol{\theta} - \boldsymbol{\theta}_0) \boldsymbol{\Sigma}^{-1} (\boldsymbol{\theta} - \boldsymbol{\theta}_0)^T \leq \delta, \quad (\text{B.2})$$

where $\boldsymbol{\Sigma}$ is the covariance matrix from our MCMC fit and δ is chosen to isolate a HPD region. To estimate the (inverse of) the evidence, we average the inverse of the posterior probability for points sampled from the ellipsoid defined by Equation B.2, with δ chosen to yield a particular HPD region. We then divide this quantity by the volume of the HPD region, which we compute by finding the convex hull of the sampled points. We repeat this procedure several times for values of δ chosen to yield HPD regions containing 10%–25% of our Markov Chain samples, which correspond to $\delta \approx 1$ for our chains; we then average the resulting values of the evidence. We have verified that our evidence estimates are converged by varying the range of HPD regions over which this mean is computed. We refer the reader to [388] and references therein for further details on this procedure.

Finally, we calculate Bayes factors K by taking the ratio of the evidence for each set of fits. We find $K(\chi = 0.8/\chi = 1) = 2.5$, $K(\beta = 0/\beta = 1) = 0.97$, and $K(\mathcal{O} = 0/\mathcal{O} = 1) = 0.98$; these results represent weak evidence (or, in terms of the Jeffreys scale [239], evidence that is “barely worth mentioning”) against our $\chi = 1$, $\beta = 0$, and $\mathcal{O} = 0$ models.¹ Thus, neither radial scaling, size reduction due to tidal stripping, nor orphan satellites significantly impact the fit presented herein, and our evidence calculations justify fixing χ , β , and \mathcal{O} for our fit to the classical-plus-SDSS luminosity distribution. The fact that we only find weak evidence in favor of $\beta = 1$ suggests that the classical-plus-SDSS luminosity distribution is fairly consistent with a subhalo–satellite size relation set at accretion. Similarly, the fact that orphans do not appreciably impact our fit hints that it is not necessary to invoke a significant amount of spurious subhalo disruption to fit the luminosity distribution of classical-plus-SDSS satellites, which correspond to subhalos with $V_{\text{peak}} \gtrsim 20 \text{ km s}^{-1}$ in our best-fit model. As discussed previously, smaller values of χ and larger values of \mathcal{O} might be favored by a joint fit to the observed luminosity distribution and radial distribution of MW satellites, in which case the Bayesian evidence for models with $\chi < 1$ and $\mathcal{O} > 0$ would increase accordingly.

¹We have checked that these Bayes factors are monotonic functions of χ , β , and \mathcal{O} in order to rule out the possibility that intermediate values of these parameters (e.g., $0 < \mathcal{O} < 1$) are preferred.

Appendix C

Satellite Model Data Comparison Details

C.1 Galaxy–Halo Connection Model Details

In this appendix, we examine several components of the galaxy–halo connection model from Chapter 4 to determine whether any of our assumptions significantly impact the results presented above.

C.1.1 MW Host Halo Mass

The analysis in this paper is restricted to two fiducial MW-like hosts with virial masses of 1.57 and $1.26 \times 10^{12} M_\odot$. However, we expect the uncertainty in the MW host halo mass, M_{MW} , to impact our constraints. Consider a toy model in which N_{sat} satellites brighter than a limiting magnitude $M_{V,\text{lim}}$ must be predicted in order to match the observed luminosity function. In this toy model, the predicted number of satellites is given by

$$\begin{aligned} N_{\text{sat}}(< M_{V,\text{lim}}) &= \int_{M_{V,\text{lim}}}^{-\infty} \frac{dN_{\text{sat}}}{dM_V} dM_V \\ &= \int_{\mathcal{M}_{\text{min}}(M_{V,\text{lim}})}^{\infty} f\left(\frac{dN_{\text{sub}}}{d\mathcal{M}_{\text{peak}}}, \boldsymbol{\theta}\right) d\mathcal{M}_{\text{peak}}, \end{aligned} \quad (\text{C.1})$$

where $dN_{\text{sub}}/d\mathcal{M}_{\text{peak}}$ is the subhalo mass function, \mathcal{M}_{min} is the lowest halo mass populated by an observed satellite, and f encapsulates the observational selection, subhalo disruption, and galaxy occupation effects that determine whether each halo hosts an observable satellite, all of which depend on galaxy–halo connection model parameters $\boldsymbol{\theta}$. Neglecting the dependence of the latter effects on host mass (which we expect to be subdominant compared to the overall rescaling of subhalo abundances), using the standard linear relationship between subhalo abundance and host mass, and

assuming a standard subhalo mass function slope of $dN_{\text{sub}}/d\mathcal{M}_{\text{peak}} \propto \mathcal{M}_{\text{peak}}^{-2}$ (e.g., [315]), we have

$$N_{\text{sat}} \propto M_{\text{MW}} \int_{\mathcal{M}_{\text{min}}}^{\infty} \mathcal{M}_{\text{peak}}^{-2} d\mathcal{M}_{\text{peak}} \propto \frac{M_{\text{MW}}}{\mathcal{M}_{\text{min}}}. \quad (\text{C.2})$$

Thus, for a fixed observed satellite count N_{sat} , we expect our 95% confidence level upper limit on \mathcal{M}_{min} to scale linearly with host mass. In addition, because the error on MW mass is independent of the error on \mathcal{M}_{min} , we expect these uncertainties to add in quadrature.

Given our fiducial minimum halo mass of $2.5 \times 10^8 M_{\odot}$ derived for an average host mass of $1.4 \times 10^{12} M_{\odot}$, we therefore expect $\mathcal{M}_{\text{min}} < 3.2 \times 10^8 M_{\odot}$ ($\mathcal{M}_{\text{min}} < 2 \times 10^8 M_{\odot}$) for a maximally high-mass (maximally low-mass) host halo given the current 2σ observational uncertainty on the MW virial mass of $1.0 \times 10^{12} < M_{\text{MW}}/M_{\odot} < 1.8 \times 10^{12}$ [102, 106, 292, 291]. We expect the remaining galaxy–halo connection model parameters and associated errors to remain largely unchanged, although rerunning our analysis using additional simulations is required to confirm this hypothesis. We expect the inferred total satellite count to scale linearly with MW mass; thus, given our fiducial prediction of 220 total MW satellites with $M_V < 0$ mag and $r_{1/2} > 10$ pc, we expect 280 (170) such satellites for a maximally high-mass (maximally low-mass) host halo.

Above, we implicitly assumed that our galaxy–halo connection model is capable of adjusting the number of faint satellites, which make the largest contribution to N_{sat} , while simultaneously matching the bright end of the observed satellite luminosity functions. This assumption holds because we have fixed the abundance-matching prescription to the relation derived from GAMA data for $M_V < -13$ mag while allowing the faint-end slope to vary. We note that the results of [340] suggest that the inferred number of satellites within a fixed physical radius is independent of M_{MW} . We find that the total number of satellites inferred within the virial radius scales almost exactly linearly with M_{MW} , as expected from the linear scaling of subhalo abundance with host halo mass, and further confirming the consistency of the results among our two fiducial simulations. In addition, we find that \mathcal{M}_{min} is roughly independent of M_{MW} for our fiducial simulations.

C.1.2 Mass-dependent Scatter

Here we test a model where the abundance-matching scatter in luminosity at fixed V_{peak} , σ_M , depends on peak halo mass. Motivated by the model in [187], we set

$$\sigma_M \equiv \sigma_{M,0} - \gamma_M (\log \mathcal{M}_{\text{peak}} - \log \mathcal{M}_1), \quad (\text{C.3})$$

where $\sigma_{M,0}$ is a free parameter that captures the amplitude of the luminosity scatter, γ_M is a free parameter that captures its mass dependence, and $\mathcal{M}_1 = 10^{11} M_{\odot}$ is fixed. By rerunning our fit with γ_M as an additional ninth free parameter, we find that large values of γ_M are ruled out by the DES and PS1 satellite populations at high statistical significance, with $\gamma_M < 0.07$ at 95% confidence.

Large values of γ_M are disfavored because abundant, low-mass halos host satellites that upscatter to observable luminosities too often to match the observed DES and PS1 luminosity functions; however, the same caveats noted in §4.8.4 for our constraint on σ_M apply to γ_M , so this upper limit should be interpreted with caution. Introducing mass-dependent scatter does not significantly affect our inferred upper bound on \mathcal{M}_{50} , implying that our fiducial minimum halo mass constraint does not depend on the details of our luminosity scatter model.

C.1.3 Radial Scaling

To account for potential biases in our radial subhalo distributions due to artificial disruption and halo finder incompleteness, we define the parameter χ by

$$r_{\text{sat}} \equiv \chi r_{\text{sub}}, \quad (\text{C.4})$$

where r_{sat} is a satellite’s distance from the center of its host halo, which we equate to its galactocentric distance, and r_{sub} is the galactocentric distance of the corresponding subhalo.

In our main analysis, we take subhalo positions directly from the simulation data and therefore assume $\chi = 1$. However, as noted above, our fiducial model slightly underpredicts the observed radial distribution of satellites close to the center of the MW in the PS1 footprint. We plot the predicted DES and PS1 radial distributions for our fiducial model in Figure C.1; to illustrate the effect of varying χ , we also show the 68% confidence interval for our fiducial posterior evaluated with $\chi = 0.5$.

To test the impact of radial scaling, we refit the DES and PS1 satellite populations with χ as an additional ninth free parameter. As expected, decreasing χ reduces the tension between the predicted and observed inner radial distribution of PS1 satellites; however, doing so does not significantly affect the goodness of fit for the observed luminosity functions and size distributions. Moreover, our key constraints, including the upper limit on \mathcal{M}_{50} , and our conclusions regarding the impact of the LMC system are not affected. In particular, the Bayes factors in favor of our fiducial LMC model relative to the alternative LMC scenarios defined in §4.8.2 are unchanged. We note that, since we have only fit to observed absolute magnitudes and surface brightnesses, the discrepancy with the observed radial distribution for $\chi = 1$ might not persist for a fit that includes galactocentric distance; we comment on the technical difficulties associated with such a fit in Appendix C.2.4.

[78] suggested that the Gaia-Enceladus accretion event, in which an LMC-mass galaxy merged with the MW 8–11 Gyr ago, might lead to a relative overabundance of ultrafaint satellites in the MW. Because of dynamical friction, this overabundance would be particularly evident in the innermost regions of the MW and might affect the observed radial satellite distribution. Interestingly, both host halos used in this work experience a Gaia-Enceladus-like accretion event, following the definition in [78] of a massive ($\sim 10^{11} M_{\odot}$) halo merging with the MW halo between $z = 1$ and 2. Given that we

still predict an underabundance of observed satellites in the inner regions of both fiducial host halos, the Gaia-Enceladus-like events they experience do not seem to be sufficient to ease the tension with the observed radial distribution. Nonetheless, exploring the relationship between the mass accretion history of the MW and the present-day radial distribution of observed ultrafaint satellites in detail is an interesting avenue for future work.

C.1.4 Tidal Stripping

Following [336], we test a model for the evolution of satellite sizes by changing the mean sizes predicted by Equation 4.1 to

$$r'_{1/2} \equiv r_{1/2} \left(\frac{V_{\max}}{V_{\text{acc}}} \right)^\beta, \quad (\text{C.5})$$

where $r'_{1/2}$ denotes the satellite half-light radius at $z = 0$, $r_{1/2}$ is the half-light radius at accretion predicted by Equation 4.1, V_{\max} (V_{acc}) is the maximum circular velocity of a subhalo today (at accretion), and $\beta > 0$ is a parameter that controls the strength of size reduction due to tidal stripping. We set $\beta = 0$ in our fiducial analysis, meaning that satellite sizes are fixed based on halo sizes at accretion. However, tidal stripping after infall can shrink satellite sizes; for example, [359] found that $1 < \beta < 2$ describes the results of high-resolution simulations well.

In Figure C.2, we illustrate predicted size distributions for our fiducial posterior evaluated with $\beta = 3$; a large value of β was chosen to test an extreme dependence of satellite sizes on tidal stripping. We find that even this extreme model does not impact the observed satellite size distributions, indicating that our results are robust to assumptions about tidal stripping. Our simulations lack the spatial resolution to test whether the [359] prescription holds in detail and alters observed satellite size distributions, but this—along with an exploration of size enlargement due to tidal heating—is an interesting avenue for future work.

C.1.5 Concentration-dependent Satellite Sizes

[244] found that galaxy sizes in two hydrodynamic simulations follow a size relation similar to that in [275], with an additional dependence on halo concentration. In particular, the size relation

$$r_{1/2} \equiv \mathcal{A} \left(\frac{c}{10} \right)^\gamma \left(\frac{R_{\text{vir}}}{R_0} \right)^n, \quad (\text{C.6})$$

with $\mathcal{A} = 0.02$, $n = 1$, $\gamma = -0.7$, $R_0 = 1$ kpc, and halo concentration c measured as a function of redshift, fits the hydrodynamic simulation results in [244] with a residual scatter of ~ 0.15 dex. This relation implies that more concentrated halos host less extended stellar systems at a fixed virial radius in these simulations.

To test whether a concentration-dependent size model is favored by the DES and PS1 data, we

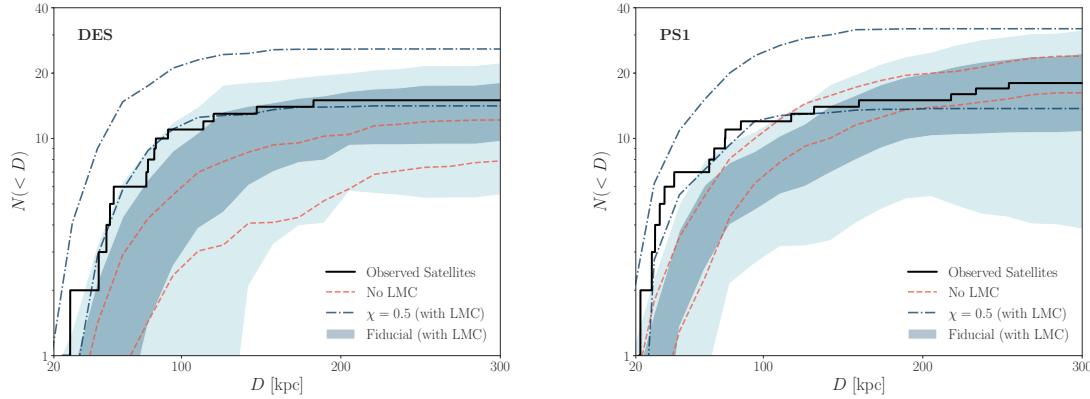


Figure C.1: Radial distributions derived from our fit to the DES and PS1 satellite populations. Our fiducial eight-parameter galaxy occupation fraction model is shown in blue. Dark (light) blue bands correspond to 68% (95%) confidence intervals, dashed red lines show the 68% confidence interval for a model using host halos without LMC analogs (“No LMC”), and black lines show the observed radial distributions. Dotted-dashed blue lines show the 68% confidence interval for a model with a radial scaling parameter of $\chi = 0.5$.

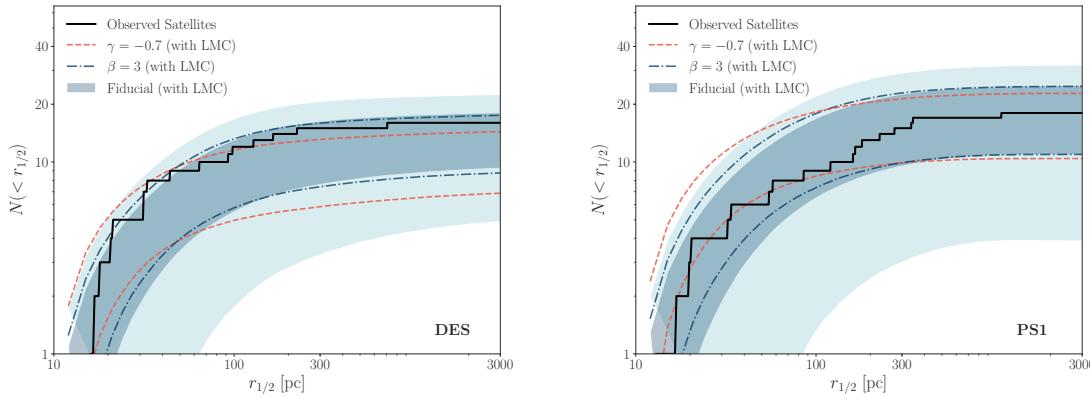


Figure C.2: Size distributions derived by fitting to the DES and PS1 satellite populations. Our fiducial eight-parameter galaxy occupation fraction model is shown in blue. Dark (light) blue bands correspond to 68% (95%) confidence intervals, dashed red lines show the 68% confidence interval for a model with a concentration-dependent galaxy–halo size relation, and dotted-dashed blue lines show the 68% confidence interval for a model with an extreme dependence of satellite size on tidal stripping.

Parameter	Prior	Motivation
α	$\arctan \alpha \sim \text{unif}(-1.1, -0.9)$	Jeffreys prior for $-2 < \alpha < -1.2$
σ_M	$\sigma_M \sim \text{unif}(0, 2)$ dex	Conservative upper limit [187, 288]
\mathcal{M}_{50}	$\log(\mathcal{M}_{50}/M_\odot) \sim \text{unif}(7.5, 11)$	Simulation resolution limit [315]
\mathcal{B}	$\ln(\mathcal{B}) \sim \mathcal{N}(\mu = 1, \sigma = 0.5)$	Hydrodynamic simulations [337, 336]
σ_{gal}	Scatter in galaxy occupation fraction	Hydrodynamic simulations [175, 202]
\mathcal{A}	$\mathcal{A} \sim \text{unif}(0, 0.5)$ kpc	Empirical galaxy–halo size relation [275]
$\sigma_{\log R}$	$\sigma_{\log R} \sim \text{unif}(0, 2)$ dex	Empirical galaxy–halo size relation [275]
n	$n \sim \mathcal{N}(\mu = 1, \sigma = 0.5)$	Empirical galaxy–halo size relation [275]

Table C.1: Prior distributions for the parameters varied in our fiducial eight-parameter fit to the DES and PS1 satellite populations. Here $\mathcal{N}(\mu, \sigma)$ denotes a normal distribution with mean μ and standard deviation σ .

refit these satellite populations with γ as an additional ninth free parameter. Because the concentration of subhalos after infall into the MW is difficult to measure accurately in our simulations, we measure the concentration at the time of accretion when implementing Equation C.6. We find that our galaxy–halo connection model constraints are largely unchanged in this case, although the upper limits on $\sigma_{\log R}$ (0.88 dex) and n (1.7) are more stringent than in our fiducial model. We find that the amplitude of the size relation is degenerate with γ , and our analysis does not place an upper limit on \mathcal{A} in this case. Here γ itself is constrained to lie between -1.5 and -0.2 at 95% confidence. The predicted luminosity functions and size distributions are nearly identical to those from our fiducial analysis (we illustrate the size distribution for our fiducial posterior evaluated with $\gamma = -0.7$ in Figure C.2).

C.1.6 Orphan Satellite Contribution

To test the importance of orphan satellites, we refit the DES and PS1 satellite populations with $\mathcal{O} = 0$, which adds zero orphans to our fiducial subhalo populations and effectively assumes that there is no artificial subhalo disruption in our simulations. Our constraints are virtually unaffected by this extreme variation in \mathcal{O} . In particular, the 95% confidence upper limit on \mathcal{M}_{50} increases by less than 1σ to $10^8 M_\odot$, and the rest of our galaxy–halo connection model constraints are also not significantly affected. The total number of MW satellites with $M_V < 0$ mag and $r_{1/2} > 10$ pc decreases to 190 ± 50 , as expected from the absence of an orphan satellite population. Thus, $\sim 15\%$ of the systems in our best-fit model are orphan satellites; these satellites might be associated with heavily stripped or disrupting subhalos.

C.2 Statistical Framework Details

Next, we provide additional details on our statistical framework, and we discuss several caveats.

C.2.1 Poisson Process Likelihood

The Poisson point process likelihood in our statistical comparison to observed satellites is implemented as follows. Suppose we observe n_i real satellites and $\hat{n}_{i,\nu}$ mock satellites in an absolute magnitude bin i , where $\nu = 1, \dots, \hat{N}$ runs over all model realizations, including different host halos, and draws from our stochastic galaxy–halo connection model. The likelihood of observing these satellites given our model realizations, which enters Equation 4.6, is then

$$P(n_i | \hat{n}_{i,1}, \dots, \hat{n}_{i,\hat{N}}) = \left(\frac{\hat{N}+1}{\hat{N}} \right)^{-(\hat{n}_{i,1} + \dots + \hat{n}_{i,\hat{N}} + 1)} \times (\hat{N}+1)^{-n_i} \frac{\Gamma(\hat{n}_{i,1} + \dots + \hat{n}_{i,\hat{N}} + n_i + 1)}{\Gamma(n_i + 1) \Gamma(\hat{n}_{i,1} + \dots + \hat{n}_{i,\hat{N}} + 1)}, \quad (\text{C.7})$$

where the dependence on galaxy–halo connection model parameters $\boldsymbol{\theta}$ is implicit, and we assumed (i) a flat prior on λ_i for $\lambda_i \geq 0$, and (ii) that n_i and all $\hat{n}_{i,\nu}$ are drawn from the same Poisson distribution with rate parameter λ_i . Note that our method yields noninteger numbers of mock satellites by counting each system as $p_{\text{detect}} \times (1 - p_{\text{disrupt}}) \times f_{\text{gal}}$ objects according to Equation 4.5, so we have replaced factorials in the Poisson likelihood with appropriate gamma functions. Our results are unaffected if we enforce integer satellite counts by performing a binary mock observation of each predicted satellite according to its detection probability.

C.2.2 Priors

We list the prior distributions used in our fiducial analysis in Table C.1, several of which are informed by previous work. The prior on the faint-end slope is a noninformative Jeffreys prior [241]. The upper limit on the luminosity scatter is chosen to be very conservative; for example, [288] found that abundance-matching scatter at the luminosity scale of the brightest systems used in our analysis is less than ~ 0.25 dex. For \mathcal{M}_{50} , we set the lower limit of the prior based on the resolution limit of our simulations, which is a maximally conservative choice from the perspective of the inferred upper limit on this quantity. In particular, while we can decrease the lower limit of this prior because the \mathcal{M}_{50} posterior is flat below $\sim 5 \times 10^7 M_{\odot}$ due to the limited sensitivity of the DES and PS1 satellite searches, doing so would artificially decrease the inferred 95% confidence upper limit.¹ Priors for \mathcal{B} and n are set based on studies that identify the preferred values of these parameters, and priors for σ_{gal} , \mathcal{A} , and $\sigma_{\log R}$ are chosen to be uniform with conservative upper bounds.

C.2.3 Bayes Factor Calculation

To calculate Bayes factors, we estimate the Bayesian evidence using the bounded harmonic mean method described in [336]. In particular, for a given posterior, we select samples of galaxy–halo

¹However, as noted in §4.8.4, our reported Bayes factors are independent of this choice.

connection model parameters $\boldsymbol{\theta}$ within a fixed Mahalanobis distance of a point $\boldsymbol{\theta}_0$ in a high-density region of the posterior. We then average the inverse of the posterior probabilities for these samples, and we normalize by the volume of the sampled region. We repeat this procedure for high-density regions that contain 10%–25% of the total number of MCMC samples, and we average over these percentiles to obtain the mean Bayesian evidence.

C.2.4 Caveats and Future Work

In this work, we fit to observed MW satellites in an observable parameter space x that consists of absolute magnitude and two large surface brightness bins. However, it would be more constraining to perform our inference in a higher-dimensional space that includes galactocentric distance. There are two main difficulties inherent in our statistical modeling.

1. We have binned observed and modeled satellites assuming that the unknown Poisson process rate in each bin is independent from the rate in other bins. This assumption is unphysical, as the rate should vary smoothly in observable parameter space.
2. As the number of bins increases, the number of satellites per bin decreases, which causes the uncertainty in the rate parameter to increase and our model to become increasingly unconstrained. This is a particularly challenging problem as we move to higher-dimensional parameter spaces, since the number of bins increases rapidly with dimensionality.

To address these issues, it is possible to connect rates in nearby regions of parameter space in an unbinned fashion using a correlated prior. This is equivalent to imposing that our galaxy–halo model should produce satellite abundances that vary smoothly as a function of observable quantities. We now lay out the mathematical formalism necessary for introducing this prior.

Our model of the distribution of satellites in observable space is an inhomogeneous Poisson process, where the number of “events” in any region \mathcal{T} of observable space x is given by a Poisson distribution with rate $\lambda_{\mathcal{T}} = \int_{\mathcal{T}} \lambda(x) dx$, where $\lambda(x)$ is referred to as the “rate function.” Given a rate function $\lambda(x)$, the likelihood of observing N events at a set of points $\{x_i\}_{i=1}^N$ is

$$p\left(\{x_i\}_{i=1}^N \mid \lambda\right) = \exp\left[-\int \lambda(x) dx\right] \prod_{i=1}^N \lambda(x_i), \quad (C.8)$$

where we suppressed the dependence of the rate on our galaxy–halo connection model parameters $\boldsymbol{\theta}$. In our case, the “events” $\{x_i\}_{i=1}^N$ are the locations of detected satellites in an observable parameter space. Note that in this formulation, there is no binning in x .

Calculating this likelihood exactly is challenging because, in order to compare observed and

modeled satellite populations, we must integrate over the unknown rate function λ ,

$$p(\{x_i\} | \{\hat{x}_j\}) = \frac{\int \mathcal{D}\lambda p(\{x_i\} | \lambda) p(\{\hat{x}_j\} | \lambda) p(\lambda)}{\int \mathcal{D}\lambda p(\{\hat{x}_j\} | \lambda) p(\lambda)}. \quad (\text{C.9})$$

Here, both the numerator and denominator contain functional integrals over the rate; these integrals are performed over an infinite-dimensional space consisting of the rate at each point in observable parameter space. Further, this rate is a stochastic function in our galaxy–halo connection model due to satellite luminosity and size scatter. This makes our model an inhomogeneous Poisson process with a stochastic rate function, which is known as a “Cox process.” The prior on the rate function, $p(\lambda)$, must admit only positive rates; one possible choice is to treat the logarithm of the rate as a Gaussian process. Models involving Cox processes are often termed “doubly intractable” due to the presence of intractable integrals over the rate function [331].

There are, however, several approaches to make Cox processes tractable. As noted above, we bin satellites in absolute magnitude and split the sample into two large surface brightness bins, so that our likelihood is over the number of counts in each bin, rather than the locations of the points. This is equivalent to assuming that the rate function is constant in each bin and leads to the likelihood

$$p\left(\{n_j\}_{j=1}^{N_{\text{bins}}} | \{\lambda_j\}\right) = \exp\left(-\sum_{\text{bins } j} \lambda_j \mathcal{V}_j\right) \prod_{\text{bins } k} \frac{\lambda_k^{n_k}}{n_k!}, \quad (\text{C.10})$$

where λ_j is the rate in bin j , \mathcal{V}_j is the volume of bin j , and n_j is the number of events in bin j . Binning turns the functional integral over $\lambda(x)$ in Equation C.9 into a finite-dimensional integral over the value of λ in each bin. Choosing Cartesian bins in observable parameter space then renders the problem tractable [178]. There also exist approaches that avoid binning the observable space [7, 245], which we intend to explore in future work.

C.3 Robustness to Observational Systematics

We now present a set of tests in order to verify the robustness of our key results to various observational systematics.

C.3.1 Kinematically Unconfirmed Satellites

To assess possible systematic uncertainties associated with the observed set of DES and PS1 satellites presented in Paper I, we rerun the entire analysis using only satellites that have are confirmed to exhibit dark matter-dominated internal kinematics. The candidate satellites excluded from this reanalysis are indicated in Table C.2. As shown in Figure C.3, our galaxy–halo connection model constraints are largely unaffected by refitting the DES and PS1 satellite populations under the

conservative assumption that all unconfirmed systems are star clusters. Most importantly, the upper limit on \mathcal{M}_{50} only increases by $\sim 1\sigma$, to $5 \times 10^8 M_\odot$ at 95% confidence, and the minimum halo mass increases to $6.5 \times 10^8 M_\odot$, similar to the minimum halo mass inferred from classical and SDSS satellites in [336]. In addition, the total predicted number of MW satellites decreases by $\sim 1\sigma$ to 150 ± 60 . These shifts are expected, since unconfirmed satellite candidates constitute many of the faintest systems in our fiducial sample. Thus, we conclude that our key constraints and predictions are not highly sensitive to the nature of kinematically unconfirmed satellite candidates.

C.3.2 Satellite Size Criterion

Next, we test whether a more conservative satellite size criterion impacts our results. For this test, we self-consistently exclude all observed and predicted satellites with $r_{1/2} > 20$ pc from our statistical inference, rather than the $r_{1/2} > 10$ pc cut used in our fiducial analysis. Our key constraints are not significantly affected; for example, the 95% confidence level upper limit on \mathcal{M}_{50} increases slightly, to $1.5 \times 10^8 M_\odot$. The upper limit on the amplitude of the galaxy–halo size relation, which was 110 pc in our fiducial analysis, increases to 220 pc, as we might expect from excluding small satellites in the fit.

C.3.3 Biases in Measured Satellite Properties

Finally, we test whether systematic offsets in measured satellite properties could affect our conclusions. In particular, we assume that every measured DES and PS1 satellite absolute magnitude is offset from the fiducial value listed in Table C.2 by $\Delta M_V = +1$ mag, which is similar to the width of the absolute magnitude bins used in our fiducial analysis. We rerun the entire analysis with these shifted magnitudes, and we repeat this procedure for $\Delta M_V = -1$ mag. In both cases, we still obtain a good joint fit to the DES and PS1 luminosity functions. As expected, the inferred faint-end slope is steeper (shallower) than that obtained from our fiducial analysis for $\Delta M_V = -1$ mag ($\Delta M_V = +1$ mag); however, the total predicted number of MW satellites with $M_V < 0$ mag and $r_{1/2} > 10$ pc and our 95% confidence upper limit on \mathcal{M}_{50} are not significantly affected in either case.

C.4 Resolution and Sample Variance

To assess the impact of resolution effects on our fiducial simulations and results, we compare the subhalo maximum circular velocity function, radial distribution, and size distribution from one of our fiducial host halos (excluding LMC satellites) to those from a higher-resolution resimulation of the same host. In particular, we resimulate this halo with a $4 \times 10^4 M_\odot h^{-1}$ high-resolution particle mass and an 85 pc h^{-1} minimum softening length. We find that the distributions of all relevant subhalo properties are not significantly affected above the resolution limit of our fiducial

simulations. Moreover, by rerunning our analysis, we find that none of our galaxy–halo connection model constraints are significantly affected when using a higher-resolution simulation.

We also assess the impact of sample variance on our fiducial subhalo and satellite populations, since the final positions of LMC satellites might be sensitive to the realizations of small-scale density fluctuations in our fiducial simulations. In particular, we resimulate both of our fiducial host halos at standard resolution with different random seeds for small-scale phases in the matter power spectrum below $60 \text{ kpc } h^{-1}$. We find that the properties of the MW host halo and LMC halo are not significantly affected in these resimulations, and that the resulting subhalo populations are nearly identical in terms of their distributions of M_{peak} , V_{peak} , halo size at accretion, and present-day heliocentric distance, implying that our results are robust to sample variance in the phases of the matter power spectrum on small scales.

C.5 Observed Satellite Data Vectors

The confirmed and candidate DES and PS1 satellites that pass the detection criteria defined in Paper I are listed in Table C.2. Note that, although Kim 2 (DES) and Laevens 1 (PS1) formally pass these detection criteria, we do not include these systems in our analysis or Table C.2 because they are suspected to be star clusters (Paper I).

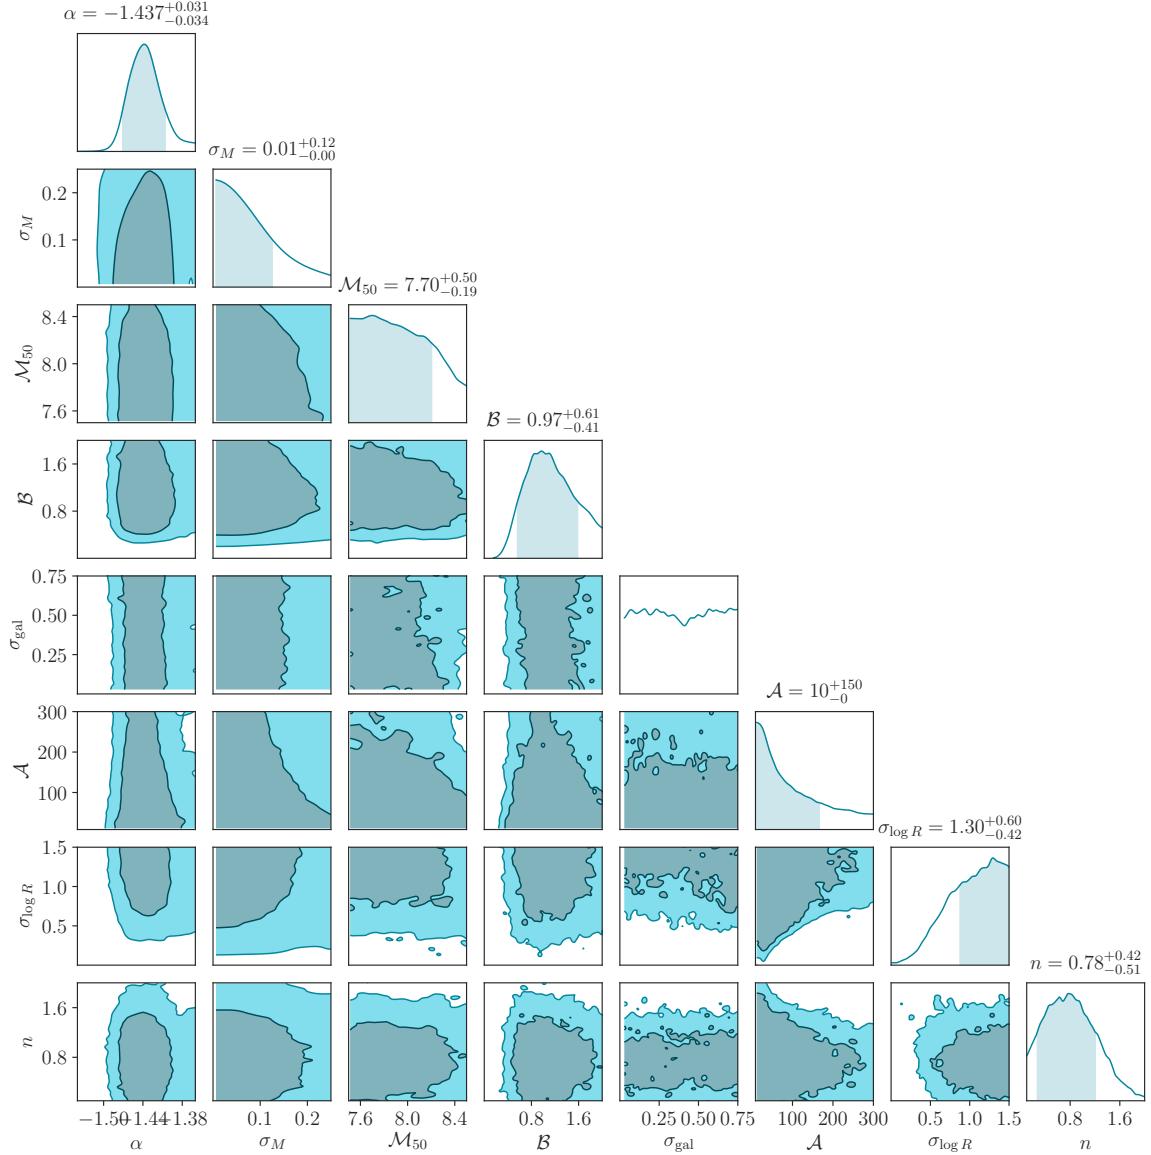


Figure C.3: Posterior distribution from our fit to the kinematically confirmed DES and PS1 satellite populations. Dark (light) shaded contours represent 68% (95%) confidence intervals. Shaded areas in the marginal distributions and parameter summaries correspond to 68% confidence intervals. Note that σ_M , σ_{gal} , and $\sigma_{\log R}$ are reported in dex, \mathcal{M}_{50} is reported as $\log(\mathcal{M}_{50}/M_\odot)$, and \mathcal{A} is reported in pc. Note that σ_{gal} is not constrained at 68% confidence in this fit.

Name	M_V (mag)	D (kpc)	$r_{1/2}$ (pc)
DES			
Fornax	-13.46	147	707
Sculptor	-10.82	84	223
Reticulum II	-3.88	30	31
Eridanus II ^a	-7.21	380	158
Tucana II	-3.8	58	165
Grus II*	-3.9	53	92
Horologium I	-3.55	79	31
Tucana III*	-2.4	25	44
Tucana IV	-3.5	48	128
Phoenix II	-3.30	83	21
Horologium II*	-2.6	78	33
Tucana V*	-1.6	55	16
Pictor I*	-3.45	114	18
Columba I*	-4.2	183	98
Cetus II*	0.02	30	17
Grus I*	-3.47	120	21
Reticulum III*	-3.31	92	64
PS1			
Leo I	-11.78	254	226
Leo II	-9.74	233	165
Draco	-8.71	76	180
Ursa Minor	-9.03	76	272
Sextans	-8.72	86	345
Canes Venatici I	-8.8	218	338
Boötes I	-6.02	66	160
Ursa Major II	-4.25	32	85
Coma Berenices	-4.38	44	57
Sagittarius II	-5.2	69	32
Willman 1	-2.53	38	20
Canes Venatici II	-5.17	160	55
Segue 1	-1.30	23	20
Segue 2*	-1.86	35	34
Crater II	-8.2	117	1066
Draco II*	-0.8	22	17
Triangulum II*	-1.60	30	13
Hercules	-5.83	132	120
Cetus II ^b	0.02	30	17

Table C.2: MW satellites used in our analysis. Properties of confirmed and candidate DES and PS1 satellites used in our analysis, listed in order of detection significance (Paper I). Asterisks mark kinematically unconfirmed systems. ^aEridanus II is not included because it lies outside our fiducial 300 kpc heliocentric distance cut. ^bCetus II is detected in both PS1 and DES; in our analysis, we only count this system in the observed DES population.

Appendix D

Warm, Interacting, & Fuzzy Dark Matter Constraints Supplement

D.1 Milky Way Satellite Model

Here, we provide a high-level overview of our MW satellite galaxy model and its associated uncertainties to supplement Chapter 7; we refer the reader to [336, 333] for a complete description. Our model of the MW satellite population is based on dark matter-only zoom-in simulations chosen to resemble the MW halo in terms of its mass, concentration, and formation history (namely, a major merger with mass ratio and infall time similar to observational estimates for the *Gaia*-Enceladus merger, and a quiescent mass accretion history thereafter). The MW host halos in these simulations respectively have virial masses of $1.57 \times 10^{12} M_{\odot}$ and $1.26 \times 10^{12} M_{\odot}$ and concentration values of 11.8 and 10.5. These concentration values are consistent within $\sim 1\sigma$ with the range derived in [102], and they are also consistent with the range inferred by [106] for the concentration of the MW halo before adiabatic contraction in the presence of baryons. This implies that our MW-like host halos are less concentrated than the real MW halo, which potentially impacts the efficiency of subhalo disruption beyond the leading-order effects captured by our model of subhalo disruption due to the MW disk, which is discussed below.

Crucially, each of these two simulations includes a realistic Large Magellanic Cloud analog system (i.e., a system with a total mass, infall time, and orbit consistent with observations of the LMC). As demonstrated in [333], the presence of a realistic LMC analog system is necessary to reproduce the observed anisotropy in the full-sky MW satellite population. We emphasize that our analysis conservatively marginalizes over the mass of the MW halo, which is the most important nuisance parameter governing the SHMF for halos that contain realistic LMC systems.

We combine these simulations with an empirical model for the galaxy–halo connection that is

specifically developed to model dwarf satellite galaxies. This model parameterizes the slope and scatter of the relation between satellite luminosity and peak halo maximum circular velocity, the amplitude, scatter, and power-law slope of the relation between satellite size and halo size, the fraction of dark matter halos occupied by observable dwarf galaxies, and the efficiency of subhalo disruption due to the MW disk. The corresponding eight galaxy–halo connection model parameters are shown in Table D.1. In this work, we also incorporate the thermal-relic WDM half-mode mass, M_{hm} (or the characteristic FDM SHMF mass suppression scale, M_0) into the model. Importantly, this empirical modeling framework allows us to parameterize and marginalize over uncertainties at the faint end of the galaxy–halo connection and the properties of the MW system in order to derive robust DM constraints.

To constrain the model using MW satellite observations, we follow the likelihood framework developed in [336, 333]. This framework compares the predicted surface brightness distribution of satellites weighted by their detection probability in the relevant photometric data to the observed count, assuming a Poisson likelihood and marginalizing over the underlying Poisson rate in each surface brightness bin. We calculate detection probabilities using the state-of-the-art observational selection functions derived in [149], which account for satellite detectability as a function of luminosity, size, distance, and sky position.

Although several recent studies reach similar conclusions about the galaxy–halo connection for MW satellites, our satellite model is the first to self-consistently include the population of LMC satellites. Moreover, our model comprehensively marginalizes over theoretical uncertainties in the faint-end galaxy–halo connection, including subhalo disruption efficiency and the satellite–halo size relation. In addition, the results of our MW satellite inference are consistent with predictions from hydrodynamic simulations [188, 30]. Nonetheless, there are several modeling uncertainties in our analysis that may impact the accuracy and precision of the resulting dark matter constraints. Most importantly, the limited number of existing high-resolution cosmological simulations with realistic LMC analogs is an important problem for any analysis that attempts to fit the full-sky MW satellite population. Fitting the observed MW satellite population using zoom-in halos without realistic LMC analogs results in a severe underabundance of predicted satellites in the DES footprint due to the lack of LMC-associated satellites [333]. Thus, redoing our analysis without specifically-selected MW host halos would yield biased (and more stringent) dark matter constraints than those presented here.

D.2 Fitting Procedure Details and Posterior Distributions

Our DM limits are derived by running 10^5 iterations of the MCMC sampler `emcee` [180] to sample the eight galaxy–halo connection model parameters described in [333], plus the DM model parameter of interest (i.e., M_{hm} for our thermal relic WDM fit and M_0 for our FDM fit), using 36

Parameter	Physical Interpretation	95% confidence interval
α	Power-law slope of satellite luminosity function	$-1.46 < \alpha < -1.38$
σ_M	Scatter in satellite luminosity at fixed halo V_{peak}	$0 \text{ dex}^* < \sigma_M < 0.2 \text{ dex}$
\mathcal{M}_{50}	Peak mass at which 50% of halos host galaxies	$7.5^* < \log(\mathcal{M}_{50}/M_\odot) < 8.0$
\mathcal{B}	Subhalo disruption efficiency relative to FIRE sims	$0.2 < \mathcal{B} < 1.9$
σ_{gal}	Width of the galaxy occupation fraction	$0 \text{ dex}^* < \sigma_{\text{gal}} < 0.66 \text{ dex}$
\mathcal{A}	Amplitude of galaxy–halo size relation	$0 \text{ pc}^* < \mathcal{A} < 90 \text{ pc}$
$\sigma_{\log R}$	Scatter in galaxy size at fixed halo properties	$0.1 \text{ dex}^* < \sigma_{\log R} < 1.1 \text{ dex}$
n	Power-law slope of galaxy–halo size relation	$0^* < n < 1.9$
M_{hm}	WDM SHMF suppression mass scale (Equation 7.5)	$7.0^* < \log(\mathcal{M}_{\text{hm}}/M_\odot) < 7.5$
M_0	FDM SHMF suppression mass scale (Equation 7.7)	$7.0^* < \log(\mathcal{M}_0/M_\odot) < 8.1$

Table D.1: Galaxy–halo connection and DM model parameters varied in our thermal relic WDM and FDM fits to the MW satellite population. Note that M_0 is constrained in a separate fit that yields similar confidence intervals for the eight galaxy–halo connection parameters. Asterisks mark prior-driven constraints. See [333] for details on our galaxy–halo connection model.

walkers. The eight galaxy–halo connection model parameters are shown in Table D.1 and described in detail by [333]. For both our thermal relic WDM and FDM fits, we discard a generous burn-in period of 2×10^4 steps, corresponding to ~ 20 autocorrelation lengths. We use the Python package `ChainConsumer` [219] to visualize the posterior distributions and calculate confidence intervals.

The posterior distributions over galaxy–halo connection and DM model parameters for our thermal relic WDM and FDM analyses are shown in Figures D.1 and D.2, respectively. Our IDM constraints are derived using the M_{hm} limit from our thermal relic WDM fit; thus, we do not show a separate posterior for the IDM analysis.

D.3 Resonantly-Produced Sterile Neutrino Constraints

To translate our upper bound on the mass of thermal relic WDM into constraints on resonantly-produced Shi–Fuller sterile neutrinos, we follow the procedure in [414, 308]. In particular, we use the sterile neutrino transfer functions generated by [111] using `CLASS` for a grid of sterile neutrino masses and mixing angles. We then compare these transfer functions to the $m_{\text{WDM}} = 6.5 \text{ keV}$ thermal relic transfer function that is ruled out at 95% confidence by our analysis. We derive the limits in Figure 7.2 by finding the combinations of sterile neutrino mass and mixing angle between the “DM Underproduction” and “DM Overproduction” lines in Figure 7.2 that yield transfer functions which are strictly more suppressed than the ruled-out thermal relic transfer function. The overproduction (underproduction) boundaries correspond to sterile neutrino models with zero (maximal) lepton asymmetry in the Neutrino Minimal Standard Model [414].

We benchmark our sterile neutrino limits using the recent estimate of SHMF suppression from [297], which is appropriate for a 7 keV resonantly-produced sterile neutrino with various lepton

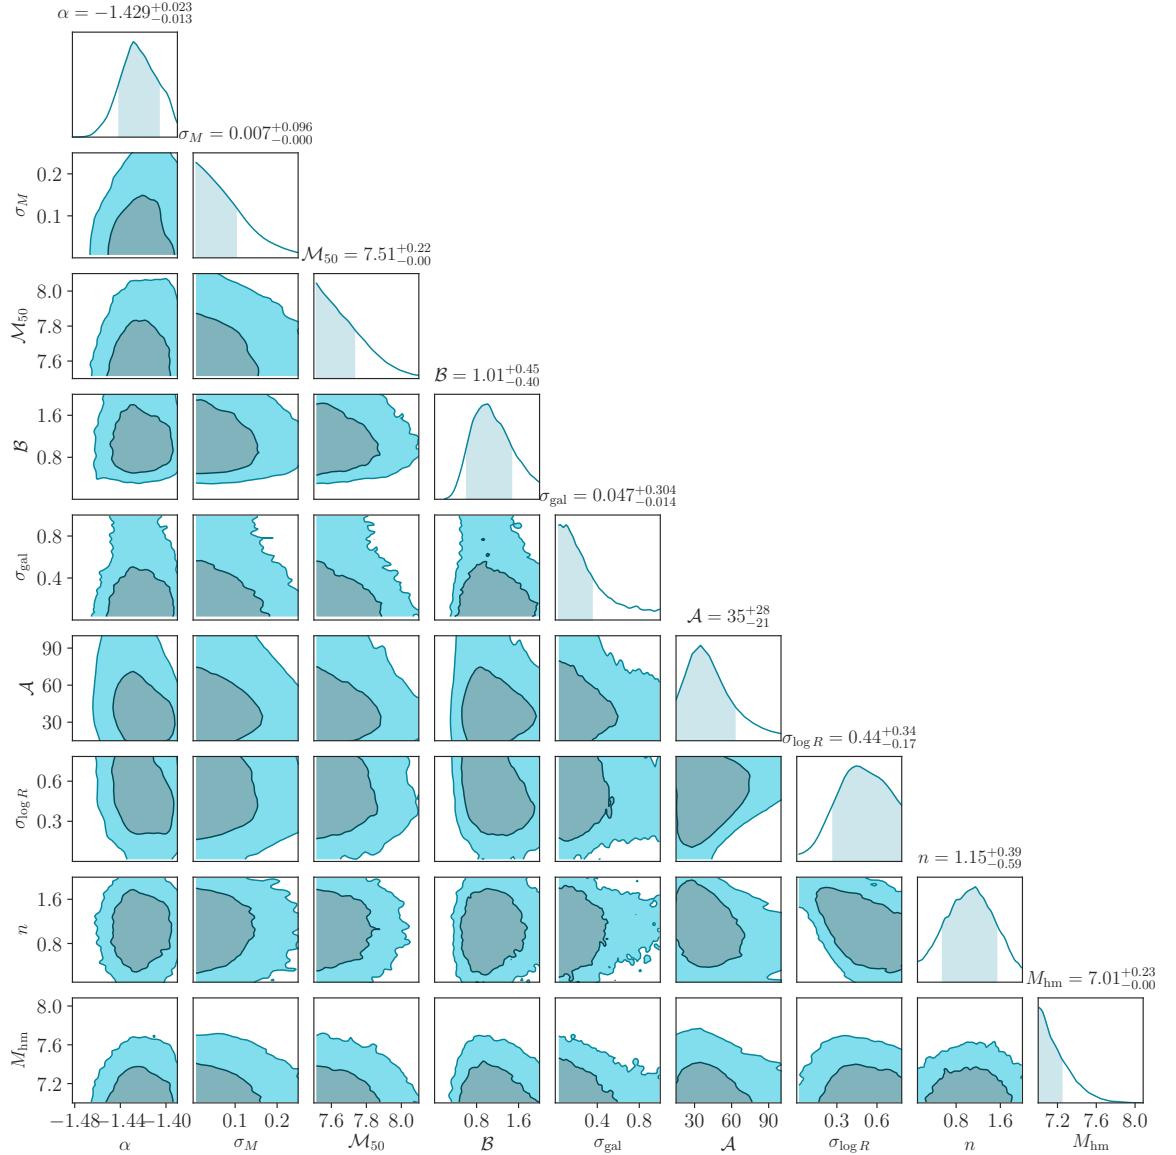


Figure D.1: Posterior distribution from our fit to the DES and PS1 satellite populations for thermal relic WDM. Dark (light) shaded contours represent 68% (95%) confidence intervals. Shaded areas in the marginal distributions and parameter summaries correspond to 68% confidence intervals. Note that σ_M , σ_{gal} , and $\sigma_{\log R}$ are reported in dex, \mathcal{M}_{50} and M_{hm} are reported as $\log(\mathcal{M}_{50}/M_{\odot})$ and $\log(M_{\text{hm}}/M_{\odot})$, \mathcal{A} is reported in pc, and α , \mathcal{B} , and n are dimensionless.

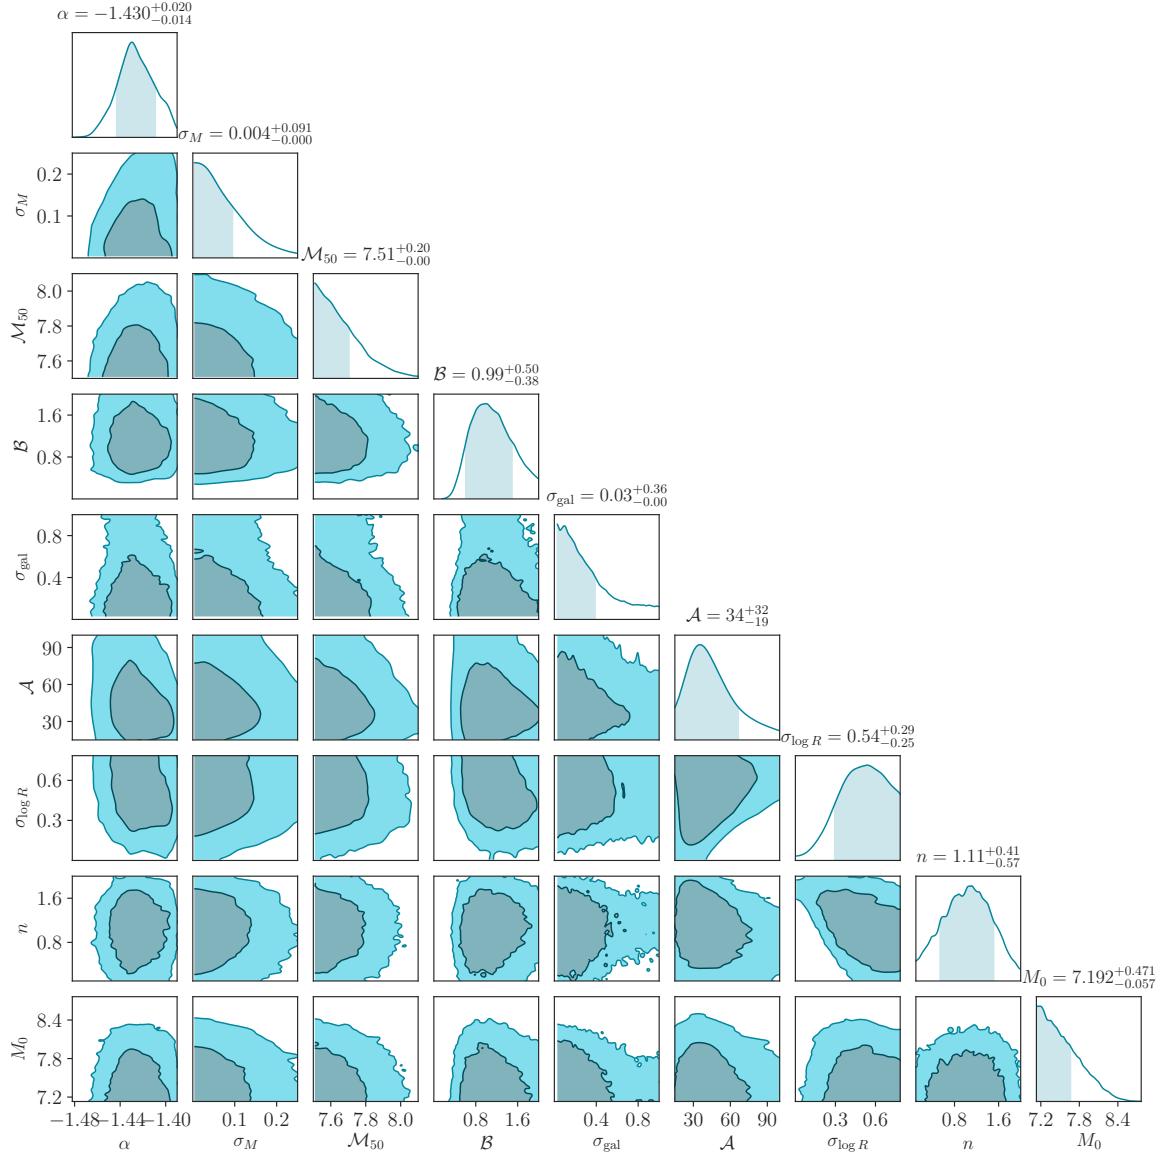


Figure D.2: Same as Figure D.1, but for our FDM fit.

asymmetry (or mixing angle) values. This SHMF suppression corresponds to Equation 7.4 with $\alpha = 4.2$, $\beta = 2.5$, $\gamma = -0.2$, and the relation between M_{hm} and lepton asymmetry L_6 is given for several models in [297]. Using this SHMF in our fitting procedure, we find $M_{\text{hm}} < 5.9 \times 10^7 \text{ M}_\odot$ at 95% confidence. Applying linear scaling with MW halo mass to the result of the joint fit yields $M_{\text{hm}} < 7.6 \times 10^7 \text{ M}_\odot$, which rules out the coldest sterile neutrino model presented in [297]—corresponding to $m_s = 7 \text{ keV}$ and $L_6 = 8$ —at $\gg 95\%$ confidence, consistent with our limit in Figure 7.2.

D.4 FDM Subhalo Mass Functions

Due to the difficulties of simulating non-linear structure formation in FDM, no consensus exists for a quantitative description of the suppression of the SHMF in this model. We therefore implemented two popular forms of the FDM SHMF to assess this uncertainty. The nominal model described in the text is the semi-analytic model derived in [153]. Our fit to this function is given by Equation 7.6, which is valid for $M \gtrsim 3 \times 10^8 \text{ M}_\odot$ and $m_\phi \gtrsim 10^{-21} \text{ eV}$ with

$$\tilde{\beta}(m_\phi) = \exp \left[- \left(\frac{m_\phi}{13.7 \times 10^{-22} \text{ eV}} \right)^{0.6} \right] + 0.77 \quad (\text{D.1})$$

$$\tilde{\gamma}(m_\phi) = 0.22 \log \left[\left(\frac{m_\phi}{10^{-22} \text{ eV}} \right)^{0.45} \right] - 0.78. \quad (\text{D.2})$$

An alternative model for the suppression of the halo mass function is derived from the “wave dark matter” simulations in [413], which corresponds to Equation 7.6 with $\tilde{\beta} = 1.1$ and $\tilde{\gamma} = -2.2$. This mass function was estimated using high-redshift ($z > 4$) simulation outputs and is systematically more suppressed than that derived semi-analytically in [153]. Adopting this alternative SHMF in our fitting procedure and accounting for the uncertainty in MW halo mass yields $M_0 < 3.4 \times 10^7 \text{ M}_\odot$ at 95% confidence, corresponding to $m_\phi > 9.1 \times 10^{-21} \text{ eV}$. Thus, the current FDM SHMF uncertainty results in roughly a factor of three difference relative to our fiducial $m_\phi > 2.9 \times 10^{-21} \text{ eV}$ constraint.

We caution that these uncertainties underlie FDM predictions from *both* semi-analytic models and simulations. For example, [413] simulate CDM-like particles with initial conditions appropriate for FDM, and thus do not solve the Schrödinger–Poisson system that governs FDM. This is an important caveat, because interference patterns on scales comparable to the de Broglie wavelength can potentially affect structure formation. Meanwhile, the semi-analytic treatment in [153] does not explicitly account for the “quantum pressure” term in the Madelung transformation of the Schrödinger–Poisson system, and makes several assumptions about the tidal evolution of subhalos with solitonic density profiles. The derivation of robust, quantitative predictions for the FDM SHMF represents an active area of theoretical and computational study.

Appendix E

SIDM Resolution Tests

This Appendix supplements the SIDM simulations presented in Chapter 9.

E.1 Simulation Resolution Tests

To test for convergence, we rerun our CDM and w_{500} simulations at higher resolution. In particular, these resimulations are run with a $4.0 \times 10^4 M_{\odot} h^{-1}$ high-resolution particle mass and an $85 \text{ pc } h^{-1}$ minimum softening length, corresponding to a factor of eight increase in mass resolution and a factor of two decrease in softening length relative to our fiducial simulations.

In Figure E.1, we compare the subhalo V_{peak} functions and radial distributions from these high-resolution resimulations to our fiducial results. We find fairly good agreement, at the $\sim 15\%$ level, between the standard and high-resolution results for subhalos above our fiducial V_{peak} thresholds. Interestingly, we find that there are *more* resolved subhalos above the relevant V_{peak} threshold in the lower-resolution simulations, which might be due to the fact that early interactions within subhalos are better resolved in the higher-resolution simulations, leading to more efficient coring. However, we note that there are also more resolved subhalos in the lower-resolution CDM simulation, so this also might indicate that the V_{peak} distributions are slightly different in the fiducial and high-resolution simulations; we have not attempted to match subhalo populations precisely for this comparison, since the level of agreement is reasonable. These findings suggest that artificial subhalo disruption is not a large effect relative to the physical disruption (or stripping below the resolution limit) of subhalos in our SIDM simulations. Thus, we conclude that our main results, including the amount of subhalo disruption in our SIDM simulations, are not highly sensitive to the resolution threshold of our fiducial simulations.

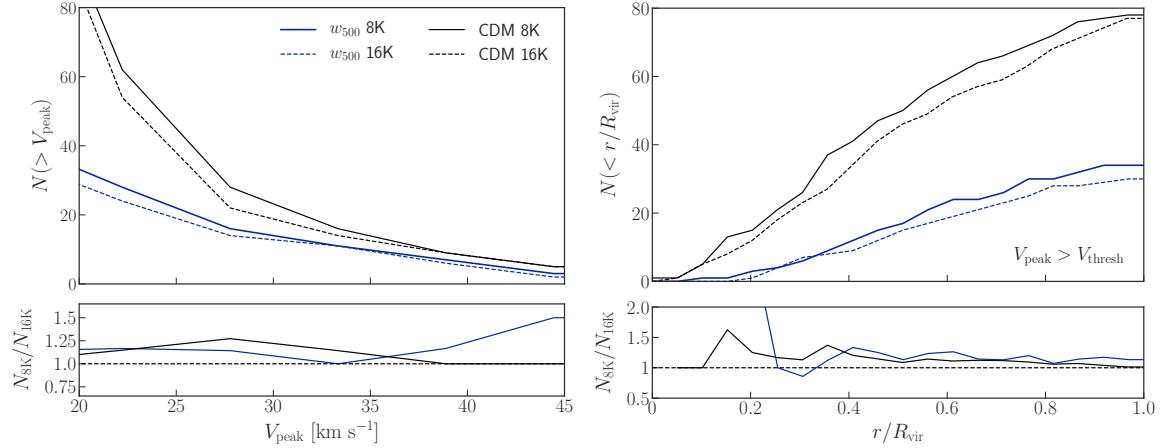


Figure E.1: Simulation resolution study. Left panel: peak velocity functions of subhalos for our high-resolution CDM and w_{500} resimulations (labeled “16K”), compared to those from our fiducial simulations (labeled “8K”). Right panel: corresponding radial subhalo distributions in units of the host halo virial radius in each simulation.

E.2 Subhalo Resolution Threshold

In our fiducial analysis, we employed a conservative V_{peak} resolution cut of 20 km s^{-1} for our CDM simulation, and we matched the number of surviving-plus-disrupted subhalos in CDM and each SIDM model variant using variable V_{peak} thresholds, denoted by V_{thresh} . In this appendix, we show that this statistical subhalo matching method is necessary given the scatter in the V_{peak} distributions measured by our halo finder. We then demonstrate that our matched subhalo populations are well converged given our fiducial choice of $V_{\text{thresh}} = 20 \text{ km s}^{-1}$ in CDM. Finally, we summarize our main results for a lower V_{peak} threshold.

E.2.1 Subhalo Abundance for a Fixed V_{peak} Threshold

Figure E.2 shows the number of surviving, disrupted, and surviving-plus-disrupted subhalos in SIDM relative to that in CDM if the *same* value of V_{thresh} is used for all of the simulations. The scatter among the SIDM model variants in all three panels of Figure E.2 strongly suggests that a variable V_{peak} threshold must be adopted in order to match subhalo populations in our CDM and SIDM simulations. Physically, this is reasonable because of the differences in subhalo assembly on an object-by-object basis caused by early self-interactions described in Section 9.6.1; however, the differences could also be driven by halo finder issues. Note that the scatter among our SIDM model variants increases for lower values of V_{thresh} , making precise subhalo population matching more difficult near the CDM resolution limit of $V_{\text{peak}} \approx 10 \text{ km s}^{-1}$.

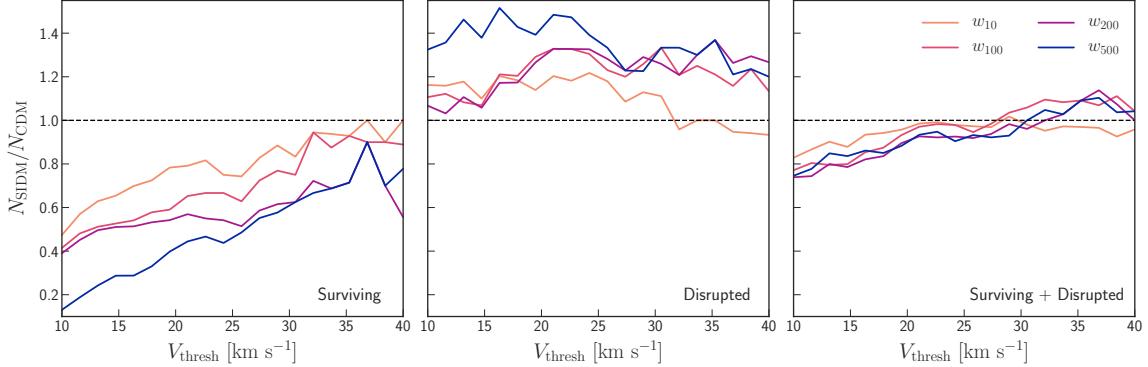


Figure E.2: Subhalo resolution threshold study. Left panel: total number of surviving subhalos in each SIDM model variant above a fixed V_{peak} threshold divided by the corresponding number of surviving subhalos in our CDM simulation. Middle panel: same as the left panel, but for the number of disrupted subhalos. Right panel: same as the previous panels, but for the number of surviving-plus-disrupted subhalos.

E.2.2 Choice of Fiducial V_{peak} Threshold

We observe that the *total* number of subhalos is stable at the $\sim 10\%$ level for $V_{\text{thresh}} \gtrsim 20 \text{ km s}^{-1}$, lending confidence to our fiducial subhalo resolution threshold. The residual differences relative to CDM for large values of V_{thresh} are likely due to a combination of physical effects (i.e., true differences in subhalo assembly in the presence of self-interactions) and numerical artifacts (e.g., uncertainties inherent to using halo finders optimized for CDM on our SIDM simulations, or artificial subhalo disruption). Disentangling these effects is beyond the scope of this paper; however, as noted in Section 9.8, this $\sim 10\%$ effect must be accounted for as a theoretical systematic until the cause of these differences is well understood.

E.2.3 Severity of Subhalo Disruption for a Less Conservative V_{peak} Threshold

Finally, we estimate the severity of subhalo disruption in our SIDM simulations for a lower V_{peak} threshold. In general, we expect enhanced disruption for subhalos with $V_{\text{peak}} \lesssim 20 \text{ km s}^{-1}$ because these systems are even more susceptible to tidal disruption than those studied in our fiducial analysis. Indeed, we observe that the fraction of surviving subhalos in SIDM falls off sharply at low V_{peak} thresholds, particularly for $V_{\text{thresh}} \lesssim 15 \text{ km s}^{-1}$, and we have confirmed that this behavior persists in our higher-resolution w_{500} resimulation. We also note that, in SIDM models with large self-interaction cross sections at low relative velocities (e.g., w_{10}), self-interactions *within* subhalos can significantly affect their density profiles, again making them more susceptible to tidal disruption. There is a hint of this effect in the left panel of Figure E.2, where we observe the steepest downturn

in the number of surviving subhalos relative to CDM for w_{10} near the CDM resolution limit of $V_{\text{peak}} \approx 10 \text{ km s}^{-1}$.

It is unclear whether the downturn in the fraction of surviving subhalos in our SIDM simulations is a consequence of physical disruption, artificial disruption, and/or halo finder issues in this regime. However, for completeness, we calculate the fraction of surviving subhalos in SIDM relative to that in CDM using our statistical subhalo matching procedure for $V_{\text{thresh}} = 15 \text{ km s}^{-1}$ in CDM. This choice yields $V_{\text{thresh}} = [14.36, 13.93, 13.84, 14.07] \text{ km s}^{-1}$ and a surviving subhalo fraction of $N_{\text{SIDM}}/N_{\text{CDM}} = [0.72, 0.63, 0.65, 0.31]$ for $[w_{10}, w_{100}, w_{200}, w_{500}]$, respectively. These surviving fractions are nearly identical to our fiducial results for w_{100} and w_{200} , and they are consistent at the $\sim 10\%$ (30%) level for w_{10} (w_{500}). In addition, the V_{peak} distributions of surviving-plus-disrupted subhalos remain consistent among the simulations using these V_{thresh} values.

E.3 Subhalos near the Host Center

As noted in Section 9.6.3, there is a curious “spike” in the inner radial distribution of surviving subhalos for w_{10} and w_{200} , shown in the right panel of Figure 9.6. The spike is particularly significant for w_{10} , where the number of subhalos in one of the inner radial bins is *enhanced* by a factor of two with respect to CDM. We emphasize that these spikes are not statistically significant: they only correspond to one or two additional subhalos near the center of the host in the SIDM simulations. Given the increased efficiency of tidal disruption in SIDM, it might seem surprising that these subhalos survive. While the survival (relative to our other SIDM model variants) and sinking (relative to CDM) of these objects are plausible consequences of self-interactions, we cannot distinguish the presence of these subhalos at small radii from statistical fluctuations. In particular, it is possible that pericentric passages at $z \approx 0$ determined by the orbital phases of the subhalos in w_{10} and w_{200} happen to align more closely with the $z = 0$ snapshot than for subhalos in CDM or our other SIDM model variants. On the other hand, if this behavior persists in a larger suite of SIDM simulations, it may be a physical consequence of increased drag due to self-interactions.

Appendix F

Probe Combination Details

This Appendix supplements the Milky Way satellite plus strong lensing probe combination presented in Chapter 11.

F.1 Milky Way Zoom-in Simulations

Our realistic MW-like simulations are drawn from the suite of 45 zoom-in simulations presented in [315], which have host halo virial masses between 1.2 and $1.6 \times 10^{12} M_\odot$.¹ The highest-resolution particles in these simulations have a mass of $3 \times 10^5 M_\odot h^{-1}$, and the softening length in the highest-resolution regions is $170 \text{ pc } h^{-1}$. Subhalos in these simulations are well resolved down to a present-day maximum circular velocity of $V_{\text{max}} \approx 9 \text{ km s}^{-1}$ [315], and halo catalogs and merger trees were generated using the ROCKSTAR halo finder and the CONSISTENT-TREES merger code [47, 48].

To account for the limited resolution of these simulations, we only analyze subhalos with maximum circular velocity $V_{\text{max}} > 9 \text{ km s}^{-1}$ and peak maximum circular velocity $V_{\text{peak}} > 10 \text{ km s}^{-1}$, which are typically resolved with $\gtrsim 100$ particles at the time V_{peak} is achieved. In addition, because we construct the $\mathcal{B}-\Sigma_{\text{sub}}$ relation by conservatively matching the MW zoom-in and **Galacticus**-predicted SHMFs down to the minimum halo mass scale, which corresponds to $M_{\text{peak}} > 2.5 \times 10^8 M_\odot$ or $V_{\text{peak}} > 19 \text{ km s}^{-1}$ before accounting for MW host halo mass uncertainty [333], only subhalos resolved with greater than ~ 600 particles at the time V_{peak} is achieved directly influence our results.

As noted in Section 11.6, we include disrupted orphan subhalos in our predictions using the model presented in Chapter 3 [336]. This model semianalytically tracks the orbital evolution of subhalos after disruption while accounting for tidal stripping and the evolving potential of the host halo, and it is calibrated by comparing to higher-resolution versions of halos from the [315] zoom-in simulation suite. We used a higher-resolution resimulation of one of our realistic MW-like halos

¹We define virial quantities according to the [93] virial definition, with overdensity $\Delta_{\text{vir}} \simeq 99.2$ in units of the critical density as appropriate for our fiducial cosmological parameters.

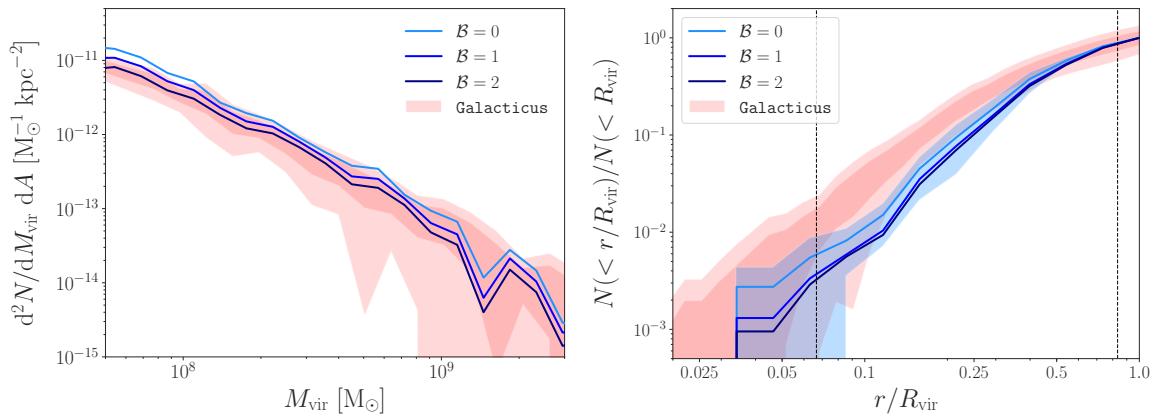


Figure F.1: Left panel: projected subhalo mass function for MW-like host halos as a function of present-day subhalo virial mass. Blue lines show results from the zoom-in simulations used in our MW satellite inference for several values of the baryonic disruption efficiency parameter \mathcal{B} (models with more efficient subhalo disruption are shown in darker colors). Right panel: average radial subhalo distribution in units of the host halo virial radius for our MW-like simulations (blue). Dashed vertical lines approximately mark the radial range of observed MW satellite galaxies used in our analysis. In both panels, dark (light) red contours show 68% (95%) confidence intervals from **Galacticus** for a sample of halos with characteristics matched to our MW-like simulations (see Appendix F.2 for details). To calculate the **Galacticus** radial distributions, we only consider halos with $M_{\text{peak}} > 10^8 M_{\odot}$ in addition to the V_{peak} and V_{max} cuts described in Appendix F.1 to facilitate a direct comparison to our simulation results.

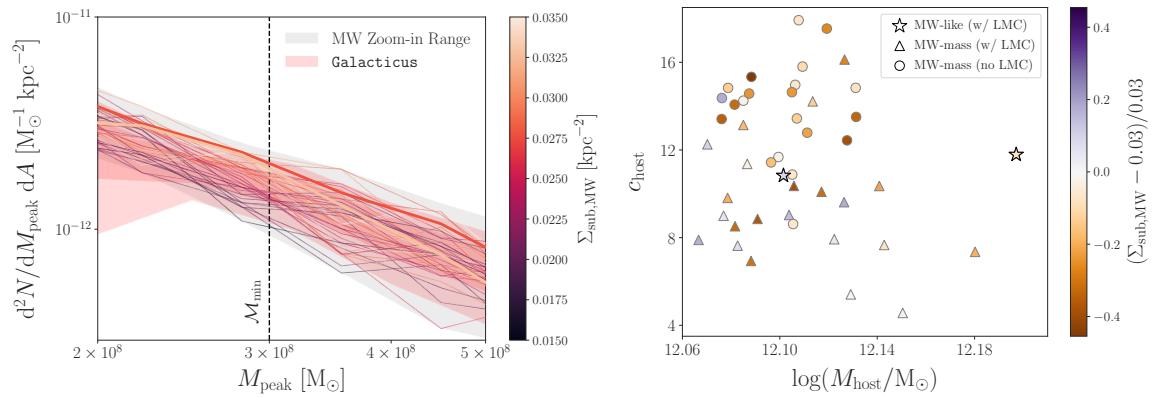


Figure F.2: Left panel: projected subhalo mass functions versus peak subhalo virial mass. Lines correspond to individual zoom-in simulations from the [315] suite of MW-mass host halos and are colored according to their projected subhalo number density $\Sigma_{\text{sub},\text{MW}}$; the two thickest lines correspond to the MW-like simulations used in our analysis. The gray band indicates the range of SHMFs from these simulations, and dark (light) red contours show 68% (95%) confidence intervals from **Galacticus** for a sample of halos with characteristics matched to our MW-like simulations. Right panel: relation between host halo mass, concentration, and $\Sigma_{\text{sub},\text{MW}}$ for the same suite of zoom-in simulations. Stars show the two MW-like simulations used in our analysis, which include realistic LMC analog systems, triangles show simulations from this suite that have an LMC analog (i.e., a subhalo with $V_{\text{max}} > 55 \text{ km s}^{-1}$) anywhere within their virial radius, and circles show simulations that do not have an LMC analog. Colors indicate fractional differences relative to the average value of $\Sigma_{\text{sub},\text{MW}}$ from the two MW-like simulations.

(described in [333]), which is run with high-resolution particles of $4 \times 10^4 M_\odot$ and an $85 \text{ pc } h^{-1}$ minimum softening length, to check that our SHMF predictions and the resulting \mathcal{B} – Σ_{sub} relation are numerically converged when including orphans.

As demonstrated in Chapter 3 [336], orphans contribute to the subhalo population at our fiducial zoom-in resolution at the $\sim 10\%$ level. The orphan contribution is roughly mass independent, increases at small Galactocentric radii, and is not highly degenerate with WDM physics, which suppresses low-mass subhalos in a radially independent manner. Furthermore, [333] show that the addition of orphans does not significantly affect galaxy–halo connection constraints derived from DES and PS1 data. The development of a self-consistent orphan model that can be applied to both N -body simulations and **Galacticus** predictions is left to future work.

F.2 Comparing Milky Way Zoom-in Simulations to **Galacticus**

We construct **Galacticus** predictions corresponding to our realistic MW-like simulations by generating host halos from a mass range corresponding to the simulations described in Appendix F.1. Host halo concentrations are generated using the [138] mass–concentration relation with 0.16 dex scatter and span the concentration values of the hosts in our MW-like zoom-in simulations. Out of these runs, we select halos that satisfy:

1. Host halo NFW concentration of $7 < c_{\text{host}} < 16$;
2. A realistic LMC analog system that accretes within the last 2 Gyr, has a present-day maximum circular velocity of $V_{\text{max}} > 55 \text{ km s}^{-1}$, Galactocentric distance of $40 \text{ kpc} < D < 60 \text{ kpc}$, and Galactocentric velocity of $267 \text{ km s}^{-1} < V < 375 \text{ km s}^{-1}$;
3. A Gaia-Enceladus-like accretion event, i.e., a merger with a satellite-to-host mass ratio in the range $[0.15, 0.25]$ in the redshift range $1 < z < 2$.

These criteria are chosen to match those imposed on our realistic MW-like simulations [333]. Note that we used Galactocentric distance to define LMC properties rather than heliocentric distance as in [333], but we do not expect this choice to impact our results. With the above criteria, roughly 0.1% of **Galacticus** runs in the relevant host halo mass range are accepted and we are left with 14 independent realizations. For computational efficiency, we ignore all subhalos accreted earlier than $z = 5$ when generating **Galacticus** predictions. A negligible fraction of halos that accrete earlier than $z = 5$ survive in our N -body simulations, implying that this choice does not impact our comparisons. Furthermore, we self-consistently employ the subhalo V_{peak} and V_{max} cuts described in Appendix F.1 when comparing to our simulation results.

The left panel of Figure F.1 shows that, in addition to the agreement among the SHMFs as a function of M_{peak} demonstrated in Figure 11.2, **Galacticus** predictions are consistent with our zoom-in simulations for SHMFs evaluated using present-day subhalo virial mass. This indicates

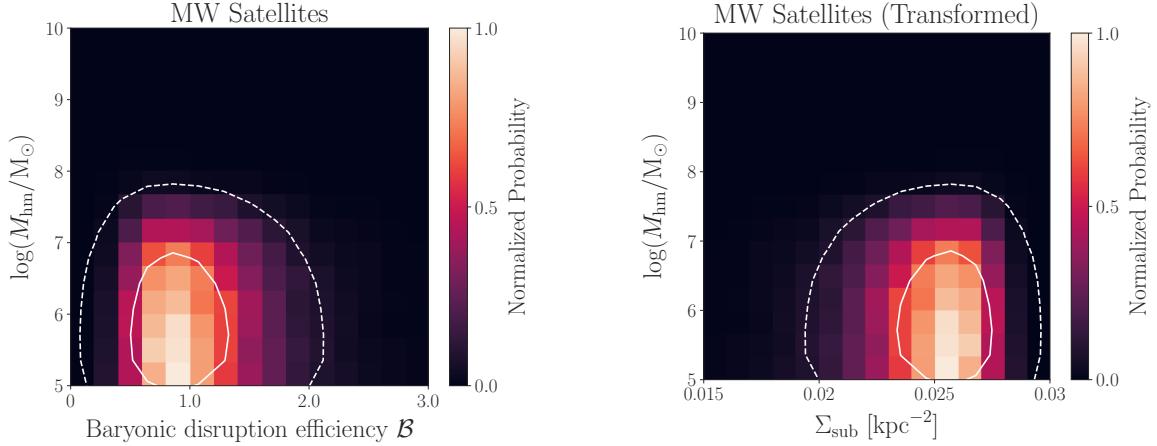


Figure F.3: Left panel: posterior distribution of WDM half-mode mass versus baryonic disruption efficiency from our analysis of the MW satellite population using a lognormal prior on \mathcal{B} , rather than a uniform prior as in Figure 11.1. Right panel: corresponding posterior distribution of WDM half-mode mass versus projected subhalo number density at the strong lensing scale, inferred according to the procedure in Section 11.6 with $q = 1$ (see the left panel of Figure 11.3 for comparison). In both panels, color maps show the probability density normalized to its maximum value in each parameter space, and solid (dashed) white lines indicate 1σ (2σ) contours for a two-dimensional Gaussian distribution.

that the amount of stripping experienced by subhalos in our N -body simulations is well captured by the `Galacticus` model, on average. However, as shown in the right panel of Figure F.1, our zoom-in simulations yield radial subhalo distributions that are slightly less concentrated than those predicted by `Galacticus`. This discrepancy is unchanged when comparing to the higher-resolution version of one of our simulations described above. We note that increasing the radial concentration of the subhalo distribution predicted by our simulations at fixed subhalo abundance would further strengthen our minimum halo mass and WDM constraints [333]. Because the radial distribution in our simulations and `Galacticus` are respectively subject to subtle numerical uncertainties including artificial subhalo disruption and semianalytic modeling of dynamical friction, we plan to explore this discrepancy systematically in future work.

F.3 Host-to-host, Poisson, and projection scatter in $\Sigma_{\text{sub},\text{MW}}$

As we have emphasized, it is challenging to accurately estimate the host-to-host scatter in $\Sigma_{\text{sub},\text{MW}}$ given the requirements we place on our realistic MW-like simulations. In this appendix, we take a very conservative approach by quantifying the scatter in $\Sigma_{\text{sub},\text{MW}}$ for the entire suite of simulations from [315] described in Appendix F.1. In particular, the left panel of Figure F.2 shows the projected SHMF for all 45 of the [315] simulations, shaded by their $\Sigma_{\text{sub},\text{MW}}$ calculated according to the

procedure in Section 11.6.2, and the right panel of Figure F.2 shows the dependence of $\Sigma_{\text{sub},\text{MW}}$ on host halo mass and concentration. These panels illustrate that the scatter in $\Sigma_{\text{sub},\text{MW}}$ is at most $\sim 40\%$ toward smaller values of $\Sigma_{\text{sub},\text{MW}}$ than inferred from our realistic MW-like simulations, and at most $\sim 20\%$ toward larger values of $\Sigma_{\text{sub},\text{MW}}$. The dependence of subhalo abundance on host halo properties among these zoom-in simulations is studied in detail by [315] and [173].

Although the host-to-host uncertainty quoted above is not small in an absolute sense, the scatter in either direction is overshadowed by the factor of 2 uncertainty introduced by q in the translation from $\Sigma_{\text{sub},\text{MW}}$ to Σ_{sub} . Moreover, scatter toward lower values of $\Sigma_{\text{sub},\text{MW}}$ (which is more common) would further strengthen our joint WDM constraints as described in Section 11.7.3. Furthermore, this estimate of the host-to-host scatter using the entire zoom-in simulation suite is an overestimate because it does not leverage additional information about the properties of the MW halo and because the [315] hosts were chosen to span a cosmologically representative range of formation histories rather than being selected uniformly in host halo mass. We therefore regard our current analysis to be conservative because it accounts for the dominant uncertainties (i.e., M_{MW} and q), and we plan to simultaneously infer q , M_{MW} , and $\Sigma_{\text{sub},\text{MW}}$ along with their associated uncertainties in future work.

The Poisson scatter in the projected SHMFs predicted by our simulations near the minimum halo mass is also moderate compared to the other systematic uncertainties discussed above. In particular, given our fiducial binning scheme, there are ~ 50 subhalos per M_{peak} bin near $\mathcal{M}_{\text{min}} = 3 \times 10^8 M_{\odot}$, corresponding to $\sim 15\%$ Poisson scatter, which is again relatively minor compared to uncertainties in q and M_{MW} . We refer the reader to [315] for a detailed study of these subhalo populations that justifies the use of a Poisson distribution to describe their scatter.

Finally, we note that the scatter in the projected subhalo mass function induced by different orientations for the projection of the MW subhalo population is also small compared to the other sources of uncertainty we have discussed. For example, the subhalo population projected with half of the virial radius in our MW-like simulations varies at the percent level for different orientations.

F.4 Subhalo Disruption Efficiency Prior

To formulate our probe combination in a statistically consistent way, we reran the Chapter 7 [332] MW satellite analysis with a uniform prior on the baryonic subhalo disruption \mathcal{B} as described in Section 11.4. However, the fiducial [332] model assumes a lognormal prior on this quantity centered around $\mathcal{B} = 1$ (i.e., the expectation for the efficiency of subhalo disruption from hydrodynamic simulations of MW-mass halos; also see Chapters 3-4 [336, 333]). In this appendix, we explore the effects of performing the probe combination using this lognormal prior.

In particular, Figure F.3 shows the posterior from the MW satellite analysis in the $\mathcal{B}-M_{\text{hm}}$ and $\Sigma_{\text{sub}}-M_{\text{hm}}$ parameter spaces, translated according to Equation 11.3 with $q = 1$, assuming the fiducial [332] prior of $\ln \mathcal{B} \sim \mathcal{N}(\mu = 1, \sigma = 0.5)$. It is visually evident that this prior favors a narrower range

of Σ_{sub} , as expected. Using this alternative prior and setting $q = 1$ does not change the results of our joint analysis, with $M_{\text{hm}} < 10^{7.0} M_{\odot}$ ($m_{\text{WDM}} > 9.7$ keV) at 95% confidence and $M_{\text{hm}} = 10^{7.5} M_{\odot}$ ($m_{\text{WDM}} = 6.9$ keV) disfavored with a 20:1 marginal likelihood ratio. This is due to a cancellation of effects: using a lognormal prior on \mathcal{B} slightly strengthens our MW satellite constraint on M_{hm} (compare the left panel of Figure F.3 to Figure 11.1), but also removes the low- Σ_{sub} tail of the MW satellite posterior (compare the right panel of Figure F.3 to the left panel of Figure 11.4). Because larger values of Σ_{sub} lead to weaker joint constraints as described in Section 11.7.3, these effects push our joint WDM constraints in opposite directions and happen to be roughly equal in magnitude.

Bibliography

- [1] M.G. Aartsen et al. Searches for Sterile Neutrinos with the IceCube Detector. *Phys. Rev. Lett.*, 117(7):071801, 2016.
- [2] M.G. Aartsen et al. Search for sterile neutrino mixing using three years of IceCube DeepCore data. *Phys. Rev. D*, 95(11):112002, 2017.
- [3] Kevork N. Abazajian. Resonantly Produced 7 keV Sterile Neutrino Dark Matter Models and the Properties of Milky Way Satellites. *Phys. Rev. Lett.*, 112(16):161303, April 2014.
- [4] Kevork N. Abazajian. Sterile neutrinos in cosmology. *Phys. Rep.*, 711:1–28, November 2017.
- [5] T. M. C. Abbott, F. B. Abdalla, S. Allam, et al. The Dark Energy Survey: Data Release 1. *ApJ Suppl.*, 239(2):18, Dec 2018.
- [6] Mario A. Acero and Julien Lesgourgues. Cosmological constraints on a light nonthermal sterile neutrino. *Phys. Rev. D*, 79:045026, Feb 2009.
- [7] Ryan Prescott Adams, Iain Murray, and David J. C. MacKay. Tractable nonparametric bayesian inference in poisson processes with gaussian process intensities. In *Proceedings of the 26th Annual International Conference on Machine Learning*, ICML '09, page 9–16, New York, NY, USA, 2009. Association for Computing Machinery.
- [8] P. Adamson, F. P. An, I. Anghel, A. Aurisano, A. B. Balantekin, H. R. Band, G. Barr, M. Bishai, A. Blake, S. Blyth, et al. Improved constraints on sterile neutrino mixing from disappearance searches in the minos, MINOS+, daya bay, and bugey-3 experiments. *Phys. Rev. Lett.*, 125:071801, Aug 2020.
- [9] P. Adamson et al. Search for active-sterile neutrino mixing using neutral-current interactions in NOvA. *Phys. Rev. D*, 96(7):072006, 2017.
- [10] P. A. R. Ade et al. Planck 2015 results. XIII. Cosmological parameters. *A&A*, 594:A13, 2016.

- [11] R. Adhikari, M. Agostini, N. Anh Ky, T. Araki, M. Archidiacono, M. Bahr, J. Baur, J. Behrens, F. Bezrukov, P. S. Bhupal Dev, D. Borah, A. Boyarsky, A. de Gouvea, C. A. de S. Pires, H. J. de Vega, A. G. Dias, P. Di Bari, Z. Djurcic, K. Dolde, H. Dorrer, M. Durero, O. Dragoun, M. Drewes, G. Drexlin, Ch. E. Düllmann, K. Eberhardt, S. Eliseev, C. Enss, N. W. Evans, A. Faessler, P. Filianin, V. Fischer, A. Fleischmann, J. A. Formaggio, J. Franse, F. M. Fraenkle, C. S. Frenk, G. Fuller, L. Gastaldo, A. Garzilli, C. Giunti, F. Glück, M. C. Goodman, M. C. Gonzalez-Garcia, D. Gorbunov, J. Hamann, V. Hannen, S. Hannestad, S. H. Hansen, C. Hassel, J. Heeck, F. Hofmann, T. Houdy, A. Huber, D. Iakubovskyi, A. Ianni, A. Ibarra, R. Jacobsson, T. Jeltema, J. Jochum, S. Kempf, T. Kieck, M. Korzeczek, V. Kornoukhov, T. Lachenmaier, M. Laine, P. Langacker, T. Lasserre, J. Lesgourgues, D. Lhuillier, Y. F. Li, W. Liao, A. W. Long, M. Maltoni, G. Mangano, N. E. Mavromatos, N. Menci, A. Merle, S. Mertens, A. Mirizzi, B. Monreal, A. Nozik, A. Neronov, V. Niro, Y. Novikov, L. Oberauer, E. Otten, N. Palanque-Delabrouille, M. Pallavicini, V. S. Pantuev, E. Papastergis, S. Parke, S. Pascoli, S. Pastor, A. Patwardhan, A. Pilaftsis, D. C. Radford, P. C. O. Ranitzsch, O. Rest, D. J. Robinson, P. S. Rodrigues da Silva, O. Ruchayskiy, N. G. Sanchez, M. Sasaki, N. Saviano, A. Schneider, F. Schneider, T. Schwetz, S. Schönert, S. Scholl, F. Shankar, R. Shrock, N. Steinbrink, L. Strigari, F. Suekane, B. Suerfu, R. Takahashi, N. Thi Hong Van, I. Tkachev, M. Totzauer, Y. Tsai, C. G. Tully, K. Valerius, J. W. F. Valle, D. Venos, M. Viel, M. Vivier, M. Y. Wang, C. Weinheimer, K. Wendt, L. Winslow, J. Wolf, M. Wurm, Z. Xing, S. Zhou, and K. Zuber. A White Paper on keV sterile neutrino Dark Matter. *JCAP*, 1:025, January 2017.
- [12] Susmita Adhikari, Neal Dalal, and Robert T. Chamberlain. Splashback in accreting dark matter halos. *JCAP*, 2014(11):019, Nov 2014.
- [13] S. Agarwal, P. S. Corasaniti, S. Das, and Y. Rasera. Small scale clustering of late forming dark matter. *Phys. Rev. D*, 92(6):063502, September 2015.
- [14] Oscar Agertz, Andrew Pontzen, Justin I. Read, Martin P. Rey, Matthew Orkney, Joakim Rosdahl, Romain Teyssier, Robbert Verbeke, Michael Kretschmer, and Sarah Nickerson. EDGE: the mass-metallicity relation as a critical test of galaxy formation physics. *MNRAS*, 491(2):1656–1672, Jan 2020.
- [15] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, S. Basak, et al. Planck 2018 results. VI. Cosmological parameters. *A&A*, 641:A6, September 2020.
- [16] N. Aghanim, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, S. Basak, et al. Planck intermediate results. XLVI. Reduction of large-scale systematic effects in HFI polarization maps and estimation of the reionization optical

depth. *A&A*, 596:A107, December 2016.

[17] A. A. Aguilar-Arevalo, C. E. Anderson, S. J. Brice, B. C. Brown, L. Bugel, J. M. Conrad, R. Dharmapalan, Z. Djurcic, B. T. Fleming, R. Ford, et al. Event excess in the miniboone search for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations. *Phys. Rev. Lett.*, 105:181801, Oct 2010.

[18] A. A. Aguilar-Arevalo, B. C. Brown, L. Bugel, G. Cheng, E. D. Church, J. M. Conrad, R. Dharmapalan, Z. Djurcic, D. A. Finley, R. Ford, et al. Improved search for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations in the miniboone experiment. *Phys. Rev. Lett.*, 110:161801, Apr 2013.

[19] A. A. Aguilar-Arevalo, B. C. Brown, L. Bugel, G. Cheng, J. M. Conrad, R. L. Cooper, R. Dharmapalan, A. Diaz, Z. Djurcic, D. A. Finley, et al. Significant excess of electron-like events in the miniboone short-baseline neutrino experiment. *Phys. Rev. Lett.*, 121:221801, Nov 2018.

[20] Kyungjin Ahn and Paul R. Shapiro. Formation and evolution of self-interacting dark matter haloes. *MNRAS*, 363(4):1092–1110, Nov 2005.

[21] S. Alam, F. D. Albareti, C. Allende Prieto, F. Anders, S. F. Anderson, T. Anderton, B. H. Andrews, E. Armengaud, É. Aubourg, S. Bailey, and et al. The Eleventh and Twelfth Data Releases of the Sloan Digital Sky Survey: Final Data from SDSS-III. *ApJs*, 219:12, July 2015.

[22] Marcelo A. Alvarez, Michael Busha, Tom Abel, and Risa H. Wechsler. Connecting Reionization to the Local Universe. *ApJl*, 703(2):L167–L171, Oct 2009.

[23] James Alvey, Nashwan Sabti, Victoria Tiki, Diego Blas, Kyrylo Bondarenko, Alexey Boyarsky, Miguel Escudero, Malcolm Fairbairn, Matthew Orkney, and Justin I. Read. New constraints on the mass of fermionic dark matter from dwarf spheroidal galaxies. *MNRAS*, 501(1):1188–1201, February 2021.

[24] Nicola C. Amorisco. Giant cold satellites from low-concentration haloes. *MNRAS*, 489(1):L22–L27, October 2019.

[25] F.P. An et al. Observation of electron-antineutrino disappearance at Daya Bay. *Phys. Rev. Lett.*, 108:171803, 2012.

[26] Donnino Anderhalden, Aurel Schneider, Andrea V. Macciò, Juerg Diemand, and Gianfranco Bertone. Hints on the nature of dark matter from the properties of Milky Way satellites. *JCAP*, 03(03):014, March 2013.

[27] G. Angloher, P. Bauer, A. Bento, C. Bucci, L. Canonica, X. Defay, A. Erb, F. v. Feilitzsch, N. Ferreiro Iachellini, P. Gorla, A. Gütlein, D. Hauff, J. Jochum, M. Kiefer, H. Kluck, H. Kraus, J.-C. Lanfranchi, A. Langenkämper, J. Loebell, M. Mancuso, E. Mondragon, A. Münster,

L. Oberauer, C. Pagliarone, F. Petricca, W. Potzel, F. Pröbst, R. Puig, F. Reindl, J. Rothe, K. Schäffner, J. Schieck, S. Schönert, W. Seidel, M. Stahlberg, L. Stodolsky, C. Strandhagen, R. Strauss, A. Tanzke, H. H. Trinh Thi, C. Türkoglu, M. Uffinger, A. Ulrich, I. Usherov, S. Wawoczny, M. Willers, M. Wüstrich, and A. Zöller. Results on mev-scale dark matter from a gram-scale cryogenic calorimeter operated above ground. *EPJC*, 77(9):637, Sep 2017.

[28] Raul E. Angulo, Oliver Hahn, and Tom Abel. The warm dark matter halo mass function below the cut-off scale. *Mon. Not. R. Astron. Soc.*, 434(4):3337–3347, October 2013.

[29] T. Antoja, C. Mateu, L. Aguilar, F. Figueiras, E. Antiche, F. Hernández-Pérez, A. G. A. Brown, O. Valenzuela, A. Aparicio, S. Hidalgo, and H. Velázquez. Detection of satellite remnants in the Galactic halo with Gaia- III. Detection limits for ultrafaint dwarf galaxies. *MNRAS*, 453:541–560, October 2015.

[30] Elaad Applebaum, Alyson M. Brooks, Charlotte R. Christensen, Ferah Munshi, Thomas R. Quinn, Sijing Shen, and Michael Tremmel. Ultrafaint Dwarfs in a Milky Way Context: Introducing the Mint Condition DC Justice League Simulations. *ApJ*, 906(2):96, January 2021.

[31] E. Aprile, J. Aalbers, F. Agostini, M. Alfonsi, F. D. Amaro, M. Anthony, F. Arneodo, P. Barrow, L. Baudis, B. Bauermeister, M. L. Benabderrahmane, T. Berger, P. A. Breur, A. Brown, A. Brown, E. Brown, S. Bruenner, G. Bruno, R. Budnik, L. Bütkofer, J. Calvén, J. M. R. Cardoso, M. Cervantes, D. Cichon, D. Coderre, A. P. Colijn, J. Conrad, J. P. Cussonneau, M. P. Decowski, P. de Perio, P. Di Gangi, A. Di Giovanni, S. Diglio, G. Eurin, J. Fei, A. D. Ferella, A. Fieguth, W. Fulgione, A. Gallo Rosso, M. Galloway, F. Gao, M. Garbini, R. Gardner, C. Geis, L. W. Goetzke, L. Grandi, Z. Greene, C. Grignon, C. Hasterok, E. Hogenbirk, J. Howlett, R. Itay, B. Kaminsky, S. Kazama, G. Kessler, A. Kish, H. Landsman, R. F. Lang, D. Lellouch, L. Levinson, Q. Lin, S. Lindemann, M. Lindner, F. Lombardi, J. A. M. Lopes, A. Manfredini, I. Mariş, T. Marrodán Undagoitia, J. Masbou, F. V. Massoli, D. Masson, D. Mayani, M. Messina, K. Micheneau, A. Molinario, K. Morå, M. Murra, J. Naganoma, K. Ni, U. Oberlack, P. Pakarha, B. Pelssers, R. Persiani, F. Piastra, J. Pienaar, V. Pizzella, M.-C. Piro, G. Plante, N. Priel, L. Rauch, S. Reichard, C. Reuter, B. Riedel, A. Rizzo, S. Rosendahl, N. Rupp, R. Saldanha, J. M. F. dos Santos, G. Sartorelli, M. Scheibelhut, S. Schindler, J. Schreiner, M. Schumann, L. Scotto Lavina, M. Selvi, P. Shagin, E. Shockley, M. Silva, H. Simgen, M. v. Sivers, A. Stein, S. Thapa, D. Thers, A. Tiseni, G. Trinchero, C. Tunnell, M. Vargas, N. Upole, H. Wang, Z. Wang, Y. Wei, C. Weinheimer, J. Wulf, J. Ye, Y. Zhang, and T. Zhu. First dark matter search results from the xenon1t experiment. *Phys. Rev. Lett.*, 119:181301, Oct 2017.

- [32] M. Arik, S. Aune, K. Barth, A. Belov, H. Bräuninger, J. Bremer, V. Burwitz, G. Cantatore, J. M. Carmona, S. A. Cetin, et al. New solar axion search using the CERN Axion Solar Telescope with ${}^4\text{He}$ filling. *Phys. Rev. D*, 92(2):021101, July 2015.
- [33] Eric Armengaud, Nathalie Palanque-Delabrouille, Christophe Yèche, David J. E. Marsh, and Julien Baur. Constraining the mass of light bosonic dark matter using SDSS Lyman- α forest. *MNRAS*, 471:4606–4614, November 2017.
- [34] Takehiko Asaka, Steve Blanchet, and Mikhail Shaposhnikov. The ν MSM, dark matter and neutrino masses [rapid communication]. *Physics Letters B*, 631(4):151–156, December 2005.
- [35] Metin Ata et al. The clustering of the SDSS-IV extended Baryon Oscillation Spectroscopic Survey DR14 quasar sample: first measurement of baryon acoustic oscillations between redshift 0.8 and 2.2. *MNRAS*, 473(4):4773–4794, 2018.
- [36] M. W. Auger, T. Treu, A. S. Bolton, R. Gavazzi, L. V. E. Koopmans, P. J. Marshall, L. A. Moustakas, and S. Burles. The Sloan Lens ACS Survey. X. Stellar, Dynamical, and Total Mass Correlations of Massive Early-type Galaxies. *ApJ*, 724(1):511–525, November 2010.
- [37] Mariah Baker and Beth Willman. Charting Unexplored Dwarf Galaxy Territory with RR Lyrae. *AJ*, 150(5):160, November 2015.
- [38] Shmuel Balberg, Stuart L. Shapiro, and Shogo Inagaki. Self-Interacting Dark Matter Halos and the Gravothermal Catastrophe. *ApJ*, 568(2):475–487, Apr 2002.
- [39] Arka Banerjee, Susmita Adhikari, Neal Dalal, Surhud More, and Andrey Kravtsov. Signatures of self-interacting dark matter on cluster density profile and subhalo distributions. *JCAP*, 2020(2):024, February 2020.
- [40] Nilanjan Banik, Jo Bovy, Gianfranco Bertone, Denis Erkal, and T. J. L. de Boer. Novel constraints on the particle nature of dark matter from stellar streams. *arXiv e-prints*, page arXiv:1911.02663, November 2019.
- [41] Nilanjan Banik, Jo Bovy, Gianfranco Bertone, Denis Erkal, and T. J. L. de Boer. Evidence of a population of dark subhaloes from Gaia and Pan-STARRS observations of the GD-1 stream. *MNRAS*, 502(2):2364–2380, April 2021.
- [42] K. Bechtol, A. Drlica-Wagner, E. Balbinot, A. Pieres, J. D. Simon, B. Yanny, B. Santiago, R. H. Wechsler, J. Frieman, A. R. Walker, P. Williams, E. Rozo, E. S. Rykoff, A. Queiroz, E. Luque, A. Benoit-Lévy, D. Tucker, I. Sevilla, R. A. Gruendl, L. N. da Costa, A. Fausti Neto, M. A. G. Maia, T. Abbott, S. Allam, R. Armstrong, A. H. Bauer, G. M. Bernstein, R. A. Bernstein, E. Bertin, D. Brooks, E. Buckley-Geer, D. L. Burke, A. Carnero Rosell, F. J. Castander, R. Covarrubias, C. B. D’Andrea, D. L. DePoy, S. Desai, H. T. Diehl, T. F. Eifler,

J. Estrada, A. E. Evrard, E. Fernandez, D. A. Finley, B. Flaugher, E. Gaztanaga, D. Gerdes, L. Girardi, M. Gladders, D. Gruen, G. Gutierrez, J. Hao, K. Honscheid, B. Jain, D. James, S. Kent, R. Kron, K. Kuehn, N. Kuropatkin, O. Lahav, T. S. Li, H. Lin, M. Makler, M. March, J. Marshall, P. Martini, K. W. Merritt, C. Miller, R. Miquel, J. Mohr, E. Neilsen, R. Nichol, B. Nord, R. Ogando, J. Peoples, D. Petravick, A. A. Plazas, A. K. Romer, A. Roodman, M. Sako, E. Sanchez, V. Scarpine, M. Schubnell, R. C. Smith, M. Soares-Santos, F. Sobreira, E. Suchyta, M. E. C. Swanson, G. Tarle, J. Thaler, D. Thomas, W. Wester, J. Zuntz, and DES Collaboration. Eight New Milky Way Companions Discovered in First-year Dark Energy Survey Data. *ApJ*, 807:50, July 2015.

[43] Keith Bechtol, Alex Drlica-Wagner, Kevork N. Abazajian, Muntazir Abidi, Susmita Adhikari, Yacine Ali-Haïmoud, James Annis, Behzad Ansarinejad, Robert Armstrong, Jacobo Asorey, Carlo Baccigalupi, Arka Banerjee, Nilanjan Banik, Charles Bennett, Florian Beutler, Simeon Bird, Simon Birrer, Rahul Biswas, Andrea Biviano, Jonathan Blazek, Kimberly K. Boddy, Ana Bonaca, Julian Borrill, Sownak Bose, Jo Bovy, Brenda Frye, Alyson M. Brooks, Matthew R. Buckley, Elizabeth Buckley-Geer, Esra Bulbul, Patricia R. Burchat, Cliff Burgess, Francesca Calore, Regina Caputo, Emanuele Castorina, Chihway Chang, George Chapline, Eric Charles, Xingang Chen, Douglas Clowe, Johann Cohen-Tanugi, Johan Comparat, Rupert A. C. Croft, Alessandro Cuoco, Francis-Yan Cyr-Racine, Guido D'Amico, Tamara M. Davis, William A. Dawson, Axel de la Macorra, Eleonora Di Valentino, Ana Díaz Rivero, Seth Digel, Scott Dodelson, Olivier Doré, Cora Dvorkin, Christopher Eckner, John Ellison, Denis Erkal, Arya Farahi, Christopher D. Fassnacht, Pedro G. Ferreira, Brenna Flaugher, Simon Foreman, Oliver Friedrich, Joshua Frieman, Juan García-Bellido, Eric Gawiser, Martina Gerbino, Maurizio Giannotti, Mandeep S. S. Gill, Vera Gluscevic, Nathan Golovich, Satya Gontcho A. Gontcho, Alma X. González-Morales, Daniel Grin, Daniel Gruen, Andrew P. Hearin, David Hendel, Yashar D. Hezaveh, Christopher M. Hirata, Renee Hložek, Shunsaku Horiuchi, Bhuvnesh Jain, M. James Jee, Tesla E. Jeltema, Marc Kamionkowski, Manoj Kaplinghat, Ryan E. Keeley, Charles R. Keeton, Rishi Khatri, Sergey E. Koposov, Savvas M. Koushiappas, Ely D. Kovetz, Ofer Lahav, Casey Lam, Chien-Hsiu Lee, Ting S. Li, Michele Liguori, Tongyan Lin, Mariangela Lisanti, Marilena LoVerde, Jessica R. Lu, Rachel Mandelbaum, Yao-Yuan Mao, Samuel D. McDermott, Mitch McNanna, Michael Medford, P. Daniel Meerburg, Manuel Meyer, Mehrdad Mirbabayi, Siddharth Mishra-Sharma, Moniez Marc, Surhud More, John Moustakas, Julian B. Muñoz, Simona Murgia, Adam D. Myers, Ethan O. Nadler, Lina Necib, Laura Newburgh, Jeffrey A. Newman, Brian Nord, Erfan Nourbakhsh, Eric Nuss, Paul O'Connor, Andrew B. Pace, Hamsa Padmanabhan, Antonella Palmese, Hiranya V. Peiris, Annika H. G. Peter, Francesco Piacentini, Andrés Plazas, Daniel A. Polin, Abhishek Prakash, Chanda Prescod-Weinstein, Justin I. Read, Steven Ritz, Brant E. Robertson, Benjamin Rose, Rogerio Rosenfeld, Graziano Rossi, Lado Samushia, Javier Sánchez, Miguel A. Sánchez-Conde, Emmanuel Schaan, Neelima

Sehgal, Leonardo Senatore, Hee-Jong Seo, Arman Shafieloo, Huanyuan Shan, Nora Shipp, Joshua D. Simon, Sara Simon, Tracy R. Slatyer, Anže Slosar, Srivatsan Sridhar, Albert Stebbins, Oscar Straniero, Louis E. Strigari, Tim M. P. Tait, Erik Tollerud, M. A. Troxel, J. Anthony Tyson, Cora Uhlemann, L. Arturo Urenña-López, Aprajita Verma, Ricardo Vilalta, Christopher W. Walter, Mei-Yu Wang, Scott Watson, Risa H. Wechsler, David Wittman, Weishuang Xu, Brian Yanny, Sam Young, Hai-Bo Yu, Gabrijela Zaharijas, Andrew R. Zentner, and Joe Zuntz. Dark Matter Science in the Era of LSST. *BAAS*, 51(3):207, May 2019.

[44] P. Behroozi, R. Wechsler, A. Hearin, and C. Conroy. UniverseMachine: The Correlation between Galaxy Growth and Dark Matter Halo Assembly from $z=0-10$. *ArXiv e-prints*, June 2018.

[45] P. S. Behroozi, R. H. Wechsler, and C. Conroy. The Average Star Formation Histories of Galaxies in Dark Matter Halos from $z = 0-8$. *ApJ*, 770:57, June 2013.

[46] P. S. Behroozi, R. H. Wechsler, Y. Lu, O. Hahn, M. T. Busha, A. Klypin, and J. R. Primack. Mergers and mass accretion for infalling halos both end well outside cluster virial radii. *ApJ*, 787:156, jun 2014.

[47] P. S. Behroozi, R. H. Wechsler, and H.-Y. Wu. The rockstar phase-space temporal halo finder and the velocity offsets of cluster cores. *ApJ*, 762:109, jan 2013.

[48] P. S. Behroozi, R. H. Wechsler, H.-Y. Wu, M. T. Busha, A. A. Klypin, and J. R. Primack. Gravitationally consistent halo catalogs and merger trees for precision cosmology. *ApJ*, 763:18, jan 2013.

[49] Nicole F. Bell, Elena Pierpaoli, and Kris Sigurdson. Cosmological signatures of interacting neutrinos. *Phys. Rev. D*, 73:063523, Mar 2006.

[50] V. Belokurov, M. G. Walker, N. W. Evans, D. C. Faria, G. Gilmore, M. J. Irwin, S. Koposov, M. Mateo, E. Olszewski, and D. B. Zucker. Leo V: A Companion of a Companion of the Milky Way Galaxy? *ApJl*, 686:L83–L86, October 2008.

[51] V. Belokurov, M. G. Walker, N. W. Evans, G. Gilmore, M. J. Irwin, D. Just, S. Koposov, M. Mateo, E. Olszewski, L. Watkins, and L. Wyrzykowski. Big Fish, Little Fish: Two New Ultra-faint Satellites of the Milky Way. *ApJl*, 712:L103–L106, March 2010.

[52] V. Belokurov, M. G. Walker, N. W. Evans, G. Gilmore, M. J. Irwin, M. Mateo, L. Mayer, E. Olszewski, J. Bechtold, and T. Pickering. The discovery of Segue 2: a prototype of the population of satellites of satellites. *MNRAS*, 397:1748–1755, August 2009.

[53] V. Belokurov, D. B. Zucker, N. W. Evans, J. T. Kleyna, S. Koposov, S. T. Hodgkin, M. J. Irwin, G. Gilmore, M. I. Wilkinson, M. Fellhauer, D. M. Bramich, P. C. Hewett, S. Vidrih,

J. T. A. De Jong, J. A. Smith, H.-W. Rix, E. F. Bell, R. F. G. Wyse, H. J. Newberg, P. A. Mayeur, B. Yanny, C. M. Rockosi, O. Y. Gnedin, D. P. Schneider, T. C. Beers, J. C. Barentine, H. Brewington, J. Brinkmann, M. Harvanek, S. J. Kleinman, J. Krzesinski, D. Long, A. Nitta, and S. A. Snedden. Cats and Dogs, Hair and a Hero: A Quintet of New Milky Way Companions. *ApJ*, 654:897–906, January 2007.

[54] V. Belokurov, D. B. Zucker, N. W. Evans, M. I. Wilkinson, M. J. Irwin, S. Hodgkin, D. M. Bramich, J. M. Irwin, G. Gilmore, B. Willman, S. Vidrih, H. J. Newberg, R. F. G. Wyse, M. Fellhauer, P. C. Hewett, N. Cole, E. F. Bell, T. C. Beers, C. M. Rockosi, B. Yanny, E. K. Grebel, D. P. Schneider, R. Lupton, J. C. Barentine, H. Brewington, J. Brinkmann, M. Harvanek, S. J. Kleinman, J. Krzesinski, D. Long, A. Nitta, J. A. Smith, and S. A. Snedden. A Faint New Milky Way Satellite in Bootes. *ApJl*, 647:L111–L114, August 2006.

[55] P. Bennet, D. J. Sand, D. Crnojević, K. Spekkens, A. Karunakaran, D. Zaritsky, and B. Mutlu-Pakdil. The M101 Satellite Luminosity Function and the Halo–Halo Scatter among Local Volume Hosts. *ApJ*, 885(2):153, Nov 2019.

[56] Andrew J. Benson. G ALACTICUS: A semi-analytic model of galaxy formation. *New A*, 17(2):175–197, February 2012.

[57] Andrew J. Benson. The normalization and slope of the dark matter (sub-)halo mass function on sub-galactic scales. *MNRAS*, 493(1):1268–1276, March 2020.

[58] T Bernal, L M Fernández-Hernández, T Matos, and M A Rodríguez-Meza. Rotation curves of high-resolution LSB and SPARC galaxies with fuzzy and multistate (ultralight boson) scalar field dark matter. *Mon. Not. R. Astron. Soc.*, 475(2):1447–1468, 12 2018.

[59] Gianfranco Bertone, Dan Hooper, and Joseph Silk. Particle dark matter: evidence, candidates and constraints. *Phys. Rep.*, 405(5-6):279–390, January 2005.

[60] Gurtina Besla. The Orbits and Total Mass of the Magellanic Clouds. *arXiv e-prints*, page arXiv:1511.03346, Nov 2015.

[61] Amit Bhoonah, Joseph Bramante, Fatemeh Elahi, and Sarah Schon. Calorimetric Dark Matter Detection with Galactic Center Gas Clouds. *Phys. Rev. Lett.*, 121:131101, Sep 2018.

[62] Simon Birrer, Adam Amara, and Alexandre Refregier. Lensing substructure quantification in RXJ1131-1231: a 2 keV lower bound on dark matter thermal relic mass. *JCAP*, 2017(5):037, May 2017.

[63] J. Bland-Hawthorn and O. Gerhard. The Galaxy in Context: Structural, Kinematic, and Integrated Properties. *ARA&A*, 54:529–596, September 2016.

- [64] J. Bland-Hawthorn, R. Sutherland, and D. Webster. Ultrafaint Dwarf Galaxies—the Lowest-mass Relics from Before Reionization. *ApJ*, 807:154, July 2015.
- [65] N. Boardman, G. Zasowski, J. A. Newman, B. Andrews, C. Fielder, M. Bershady, J. Brinkmann, N. Drory, D. Krishnarao, R. R. Lane, T. Mackereth, K. Masters, and G. S. Stringfellow. Are the Milky Way and Andromeda unusual? A comparison with Milky Way and Andromeda analogues. *MNRAS*, 498(4):4943–4954, September 2020.
- [66] Kimberly K. Boddy, Jonathan L. Feng, Manoj Kaplinghat, and Tim M. P. Tait. Self-interacting dark matter from a non-Abelian hidden sector. *Phys. Rev. D*, 89(11):115017, Jun 2014.
- [67] Kimberly K. Boddy and Vera Gluscevic. First cosmological constraint on the effective theory of dark matter-proton interactions. *Phys. Rev. D*, 98:083510, Oct 2018.
- [68] Kimberly K. Boddy, Vera Gluscevic, Vivian Poulin, Ely D. Kovetz, Marc Kamionkowski, and Rennan Barkana. Critical assessment of CMB limits on dark matter-baryon scattering: New treatment of the relative bulk velocity. *Phys. Rev. D*, 98:123506, December 2018.
- [69] Paul Bode, Jeremiah P. Ostriker, and Neil Turok. Halo Formation in Warm Dark Matter Models. *ApJ*, 556:93–107, July 2001.
- [70] C. Boehm, P. Fayet, and R. Schaeffer. Constraining dark matter candidates from structure formation. *Phys. Lett. B*, 518(1-2):8–14, Oct 2001.
- [71] C. Boehm and R. Schaeffer. Constraints on Dark Matter interactions from structure formation: damping lengths. *A&A*, 438(2):419–442, Aug 2005.
- [72] C. Boehm, J. A. Schewtschenko, R. J. Wilkinson, C. M. Baugh, and S. Pascoli. Using the Milky Way satellites to study interactions between cold dark matter and radiation. *MNRAS*, 445:L31–L35, November 2014.
- [73] Céline Boehm, Matthew J. Dolan, and Christopher McCabe. A lower bound on the mass of cold thermal dark matter from Planck. *JCAP*, 08(8):041, August 2013.
- [74] Céline Boehm, Alain Riazuelo, Steen H. Hansen, and Richard Schaeffer. Interacting dark matter disguised as warm dark matter. *Phys. Rev. D*, 66(8):083505, October 2002.
- [75] Sebastian Bohr, Jesús Zavala, Francis-Yan Cyr-Racine, Mark Vogelsberger, Torsten Bringmann, and Christoph Pfrommer. ETHOS - an effective parametrization and classification for structure formation: the non-linear regime at $z \gtrsim 5$. *MNRAS*, 498(3):3403–3419, August 2020.

- [76] Ana Bonaca, David W. Hogg, Adrian M. Price-Whelan, and Charlie Conroy. The Spur and the Gap in GD-1: Dynamical Evidence for a Dark Substructure in the Milky Way Halo. *ApJ*, 880(1):38, July 2019.
- [77] S. Bose, W. A. Hellwing, C. S. Frenk, A. Jenkins, M. R. Lovell, J. C. Helly, and B. Li. The Copernicus Complexio: statistical properties of warm dark matter haloes. *MNRAS*, 455:318–333, January 2016.
- [78] Sownak Bose, Alis J. Deason, Vasily Belokurov, and Carlos S. Frenk. The little things matter: relating the abundance of ultrafaint satellites to the hosts’ assembly history. *MNRAS*, 495(1):743–757, June 2020.
- [79] Sownak Bose, Alis J. Deason, and Carlos S. Frenk. The Imprint of Cosmic Reionization on the Luminosity Function of Galaxies. *ApJ*, 863(2):123, Aug 2018.
- [80] Sownak Bose, Wojciech A. Hellwing, Carlos S. Frenk, Adrian Jenkins, Mark R. Lovell, John C. Helly, Baojiu Li, Violeta Gonzalez-Perez, and Liang Gao. Substructure and galaxy formation in the Copernicus Complexio warm dark matter simulations. *Mon. Not. R. Astron. Soc.*, 464(4):4520–4533, Feb 2017.
- [81] Jo Bovy. Detecting the Disruption of Dark-Matter Halos with Stellar Streams. *Phys. Rev. Lett.*, 116(12):121301, Mar 2016.
- [82] A. Boyarsky, O. Ruchayskiy, D. Iakubovskyi, and J. Franse. Unidentified Line in X-Ray Spectra of the Andromeda Galaxy and Perseus Galaxy Cluster. *Phys. Rev. Lett.*, 113(25):251301, December 2014.
- [83] Alexey Boyarsky, Oleg Ruchayskiy, and Dmytro Iakubovskyi. A lower bound on the mass of dark matter particles. *JCAP*, 03(3):005, March 2009.
- [84] M. Boylan-Kolchin, J. S. Bullock, and M. Kaplinghat. Too big to fail? the puzzling darkness of massive milky way subhaloes. *MNRAS*, 415:L40–L44, jul 2011.
- [85] M. Boylan-Kolchin, J. S. Bullock, and M. Kaplinghat. The milky way’s bright satellites as an apparent failure of Λ cdm. *MNRAS*, 422:1203–1218, may 2012.
- [86] M. Boylan-Kolchin, V. Springel, S. D. M. White, and A. Jenkins. There’s no place like home? Statistics of Milky Way-mass dark matter haloes. *MNRAS*, 406:896–912, August 2010.
- [87] Michael Boylan-Kolchin, Daniel R. Weisz, Benjamin D. Johnson, James S. Bullock, Charlie Conroy, and Alex Fitts. The Local Group as a time machine: studying the high-redshift Universe with nearby galaxies. *MNRAS*, 453(2):1503–1512, Oct 2015.

- [88] Torsten Bringmann and Maxim Pospelov. Novel direct detection constraints on light dark matter. *Phys. Rev. Lett.*, 122:171801, May 2019.
- [89] Tom Broadhurst, Ivan De Martino, Hoang Nhan Luu, George F. Smoot, and S.-H. Henry Tye. Ghostly galaxies as solitons of bose-einstein dark matter. *Phys. Rev. D*, 101:083012, Apr 2020.
- [90] C. B. Brook and A. Di Cintio. Expanded haloes, abundance matching and too-big-to-fail in the local group. *MNRAS*, 450:3920–3934, jul 2015.
- [91] A. M. Brooks, M. Kuhlen, A. Zolotov, and D. Hooper. A baryonic solution to the missing satellites problem. *ApJ*, 765:22, mar 2013.
- [92] A. M. Brooks, E. Papastergis, C. R. Christensen, F. Governato, A. Stilp, T. R. Quinn, and J. Wadsley. How to Reconcile the Observed Velocity Function of Galaxies with Theory. *ApJ*, 850:97, November 2017.
- [93] Greg L. Bryan and Michael L. Norman. Statistical Properties of X-Ray Clusters: Analytic and Numerical Comparisons. *ApJ*, 495(1):80–99, Mar 1998.
- [94] D. Buch, E. O. Nadler, R. H. Wechsler, and Y.-Y Mao. Milky way-est simulation suite. *in prep.*, 2021.
- [95] Matthew R. Buckley and Annika H. G. Peter. Gravitational probes of dark matter physics. *Phys. Rep.*, 761:1–60, October 2018.
- [96] J. S. Bullock. Notes on the Missing Satellites Problem. *ArXiv e-prints*, September 2010.
- [97] J. S. Bullock and M. Boylan-Kolchin. Small-scale challenges to the Λ CDM paradigm. *ARA&A*, 55:343–387, aug 2017.
- [98] J. S. Bullock, A. V. Kravtsov, and D. H. Weinberg. Reionization and the abundance of galactic satellites. *ApJ*, 539:517–521, aug 2000.
- [99] J. S. Bullock, K. R. Stewart, M. Kaplinghat, E. J. Tollerud, and J. Wolf. Stealth Galaxies in the Halo of the Milky Way. *ApJ*, 717:1043–1053, July 2010.
- [100] M. T. Busha, P. J. Marshall, R. H. Wechsler, A. Klypin, and J. Primack. The Mass Distribution and Assembly of the Milky Way from the Properties of the Magellanic Clouds. *ApJ*, 743:40, December 2011.
- [101] Michael T. Busha, Marcelo A. Alvarez, Risa H. Wechsler, Tom Abel, and Louis E. Strigari. The Impact of Inhomogeneous Reionization on the Satellite Galaxy Population of the Milky Way. *ApJ*, 710(1):408–420, Feb 2010.

- [102] Thomas M. Callingham, Marius Cautun, Alis J. Deason, Carlos S. Frenk, Wenting Wang, Facundo A. Gómez, Robert J. J. Grand, Federico Marinacci, and Ruediger Pakmor. The mass of the Milky Way from satellite dynamics. *MNRAS*, page 365, Feb 2019.
- [103] D. J. R. Campbell, C. S. Frenk, A. Jenkins, V. R. Eke, J. F. Navarro, T. Sawala, M. Schaller, A. Fattahi, K. A. Oman, and T. Theuns. Knowing the unknowns: uncertainties in simple estimators of galactic dynamical masses. *MNRAS*, 469:2335–2360, August 2017.
- [104] Scott G. Carlsten, Jenny E. Greene, Annika H. G. Peter, Rachael L. Beaton, and Johnny P. Greco. Luminosity Functions and Host-to-host Scatter of Dwarf Satellite Systems in the Local Volume. *ApJ*, 908(1):109, February 2021.
- [105] Scott G. Carlsten, Jenny E. Greene, Annika H. G. Peter, Johnny P. Greco, and Rachael L. Beaton. Radial Distributions of Dwarf Satellite Systems in the Local Volume. *ApJ*, 902(2):124, October 2020.
- [106] Marius Cautun, Alejandro Benítez-Llambay, Alis J. Deason, Carlos S. Frenk, Azadeh Fattahi, Facundo A. Gómez, Robert J. J. Grand, Kyle A. Oman, Julio F. Navarro, and Christine M. Simpson. The milky way total mass profile as inferred from Gaia DR2. *MNRAS*, 494(3):4291–4313, May 2020.
- [107] Marius Cautun, Alis J. Deason, Carlos S. Frenk, and Stuart McAlpine. The aftermath of the Great Collision between our Galaxy and the Large Magellanic Cloud. *MNRAS*, 483(2):2185–2196, Feb 2019.
- [108] K. C. Chambers, E. A. Magnier, N. Metcalfe, H. A. Flewelling, M. E. Huber, C. Z. Waters, L. Denneau, P. W. Draper, D. Farrow, D. P. Finkbeiner, C. Holmberg, J. Koppenhoefer, P. A. Price, A. Rest, R. P. Saglia, E. F. Schlafly, S. J. Smartt, W. Sweeney, R. J. Wainscoat, W. S. Burgett, S. Chastel, T. Grav, J. N. Heasley, K. W. Hodapp, R. Jedicke, N. Kaiser, R. P. Kudritzki, G. A. Luppino, R. H. Lupton, D. G. Monet, J. S. Morgan, P. M. Onaka, B. Shiao, C. W. Stubbs, J. L. Tonry, R. White, E. Bañados, E. F. Bell, R. Bender, E. J. Bernard, M. Boegner, F. Boffi, M. T. Botticella, A. Calamida, S. Casertano, W. P. Chen, X. Chen, S. Cole, N. Deacon, C. Frenk, A. Fitzsimmons, S. Gezari, V. Gibbs, C. Goessl, T. Goggia, R. Gourgue, B. Goldman, P. Grant, E. K. Grebel, N. C. Hambly, G. Hasinger, A. F. Heavens, T. M. Heckman, R. Henderson, T. Henning, M. Holman, U. Hopp, W. H. Ip, S. Isani, M. Jackson, C. D. Keyes, A. M. Koekemoer, R. Kotak, D. Le, D. Liska, K. S. Long, J. R. Lucey, M. Liu, N. F. Martin, G. Masci, B. McLean, E. Mindel, P. Misra, E. Morganson, D. N. A. Murphy, A. Obaika, G. Narayan, M. A. Nieto-Santisteban, P. Norberg, J. A. Peacock, E. A. Pier, M. Postman, N. Primak, C. Rae, A. Rai, A. Riess, A. Riffeser, H. W. Rix, S. Röser, R. Russel, L. Rutz, E. Schilbach, A. S. B. Schultz, D. Scolnic, L. Strolger, A. Szalay, S. Seitz, E. Small, K. W. Smith, D. R. Soderblom, P. Taylor, R. Thomson, A. N. Taylor, A. R. Thakar,

J. Thiel, D. Thilker, D. Unger, Y. Urata, J. Valenti, J. Wagner, T. Walder, F. Walter, S. P. Watters, S. Werner, W. M. Wood-Vasey, and R. Wyse. The Pan-STARRS1 Surveys. *arXiv e-prints*, page arXiv:1612.05560, December 2016.

[109] T. K. Chan, D. Kereš, J. Oñorbe, P. F. Hopkins, A. L. Muratov, C.-A. Faucher-Giguère, and E. Quataert. The impact of baryonic physics on the structure of dark matter haloes: the view from the fire cosmological simulations. *MNRAS*, 454:2981–3001, dec 2015.

[110] S. Chandrasekhar. Dynamical Friction. I. General Considerations: the Coefficient of Dynamical Friction. *ApJ*, 97:255, March 1943.

[111] John F. Cherry and Shunsaku Horiuchi. Closing in on resonantly produced sterile neutrino dark matter. *Phys. Rev. D*, 95(8):083015, April 2017.

[112] Ioana Ciucă, Daisuke Kawata, Shin’ichiro Ando, Francesca Calore, Justin I. Read, and Cecilia Mateu. A Gaia DR2 search for dwarf galaxies towards Fermi-LAT sources: implications for annihilating dark matter. *MNRAS*, 480(2):2284–2291, October 2018.

[113] Pedro Colín, Vladimir Avila-Reese, Octavio Valenzuela, and Claudio Firmani. Structure and Subhalo Population of Halos in a Self-interacting Dark Matter Cosmology. *ApJ*, 581(2):777–793, Dec 2002.

[114] Thomas E. Collett. The Population of Galaxy-Galaxy Strong Lenses in Forthcoming Optical Imaging Surveys. *ApJ*, 811(1):20, September 2015.

[115] P. S. Corasaniti, S. Agarwal, D. J. E. Marsh, and S. Das. Constraints on dark matter scenarios from measurements of the galaxy luminosity function at high redshifts. *Phys. Rev. D*, 95:083512, Apr 2017.

[116] Benoit Côté, Devin W. Silvia, Brian W. O’Shea, Britton Smith, and John H. Wise. Validating Semi-analytic Models of High-redshift Galaxy Formation Using Radiation Hydrodynamical Simulations. *ApJ*, 859(1):67, May 2018.

[117] Peter Creasey, Omid Sameie, Laura V. Sales, Hai-Bo Yu, Mark Vogelsberger, and Jesús Zavala. Spreading out and staying sharp - creating diverse rotation curves via baryonic and self-interaction effects. *MNRAS*, 468(2):2283–2295, Jun 2017.

[118] CRESST Collaboration, F. Petricca, G. Angloher, P. Bauer, A. Bento, C. Bucci, L. Canonica, X. Defay, A. Erb, F. v. Feilitzsch, N. Ferreiro Iachellini, P. Gorla, A. Gütlein, D. Hauff, J. Jochum, M. Kiefer, H. Kluck, H. Kraus, J. C. Lanfranchi, A. Lagenkämper, J. Loebell, M. Mancuso, E. Mondragon, A. Münster, C. Pagliarone, W. Potzel, F. Pröbst, R. Puig, F. Reindl, J. Rothe, K. Schäffner, J. Schieck, S. Schönert, W. Seidel, M. Stahlberg, L. Stodolsky, C. Strand hagen, R. Strauss, A. Tanzke, H. H. Trinh Thi, C. Türkoğlu, A. Ulrich,

I. Usherov, S. Wawoczny, M. Willers, and M. Wüstrich. First results on low-mass dark matter from the CRESST-III experiment. *arXiv e-prints*, page arXiv:1711.07692, Nov 2017.

[119] Francis-Yan Cyr-Racine, Kris Sigurdson, Jesús Zavala, Torsten Bringmann, Mark Vogelsberger, and Christoph Pfrommer. ETHOS—an effective theory of structure formation: From dark particle physics to the matter distribution of the Universe. *Phys. Rev. D*, 93(12):123527, June 2016.

[120] Neal Dalal, Jo Bovy, Lam Hui, and Xinyu Li. Don't cross the streams: caustics from Fuzzy Dark Matter. *arXiv e-prints*, page arXiv:2011.13141, November 2020.

[121] Anirban Das, Basudeb Dasgupta, and Rishi Khatri. Ballistic dark matter oscillates above λ CDM. *JCAP*, 2019(04):018–018, apr 2019.

[122] Subinoy Das, Rajesh Mondal, Vikram Rentala, and Srikanth Suresh. On dark matter - dark radiation interaction and cosmic reionization. *JCAP*, 1808(08):045, 2018.

[123] Subinoy Das and Ethan O. Nadler. Constraints on the epoch of dark matter formation from milky way satellites. *Phys. Rev. D*, 103:043517, Feb 2021.

[124] Subinoy Das and Neal Weiner. Late forming dark matter in theories of neutrino dark energy. *Phys. Rev. D*, 84(12):123511, December 2011.

[125] Romeel Davé, David N. Spergel, Paul J. Steinhardt, and Benjamin D. Wandelt. Halo Properties in Cosmological Simulations of Self-interacting Cold Dark Matter. *ApJ*, 547(2):574–589, Feb 2001.

[126] W. J. G. de Blok. The core-cusp problem. *Advances in Astronomy*, 2010:789293, 2010.

[127] A. J. Deason, A. R. Wetzel, S. Garrison-Kimmel, and V. Belokurov. Satellites of LMC-mass dwarfs: close friendships ruined by Milky Way mass haloes. *MNRAS*, 453(4):3568–3574, Nov 2015.

[128] Alis J Deason, Carlos S Frenk, Marius Cautun, Thomas M Callingham, Wenting Wang, Facundo A Gómez, Ruediger Pakmor, Robert J J Grand, and Federico Marinacci. The mass of the Milky Way from satellite dynamics. *MNRAS*, 484(4):5453–5467, 02 2019.

[129] A. Del Popolo and M. Le Delliou. Small scale problems of the Λ cdm model: A short review. *Galaxies*, 5:17, feb 2017.

[130] DES Collaboration. The Dark Energy Survey: more than dark energy - an overview. *MNRAS*, 460:1270–1299, August 2016.

- [131] DESI Collaboration, A. Aghamousa, J. Aguilar, S. Ahlen, S. Alam, L. E. Allen, C. Allende Prieto, J. Annis, S. Bailey, C. Balland, and et al. The desi experiment part i: Science,targeting, and survey design. *ArXiv e-prints*, oct 2016.
- [132] Giulia Despali and Simona Vegetti. The impact of baryonic physics on the subhalo mass function and implications for gravitational lensing. *MNRAS*, 469(2):1997–2010, August 2017.
- [133] Giulia Despali, Simona Vegetti, Simon D. M. White, Carlo Giocoli, and Frank C. van den Bosch. Modelling the line-of-sight contribution in substructure lensing. *MNRAS*, 475(4):5424–5442, April 2018.
- [134] Christopher Dessert, Nicholas L. Rodd, and Benjamin R. Safdi. The dark matter interpretation of the 3.5-kev line is inconsistent with blank-sky observations. *Science*, 367(6485):1465–1467, 2020.
- [135] A. Di Cintio, C. B. Brook, A. A. Dutton, A. V. Macciò, G. S. Stinson, and A. Knebe. A mass-dependent density profile for dark matter haloes including the influence of galaxy formation. *MNRAS*, 441:2986–2995, jul 2014.
- [136] Arianna Di Cintio, Chris B. Brook, Andrea V. Macciò, Greg S. Stinson, Alexander Knebe, Aaron A. Dutton, and James Wadsley. The dependence of dark matter profiles on the stellar-to-halo mass ratio: a prediction for cusps versus cores. *MNRAS*, 437(1):415–423, Jan 2014.
- [137] J. Diemand, B. Moore, and J. Stadel. Earth-mass dark-matter haloes as the first structures in the early universe. *Nature*, 433:389 EP –, 01 2005.
- [138] Benedikt Diemer and Michael Joyce. An Accurate Physical Model for Halo Concentrations. *ApJ*, 871(2):168, February 2019.
- [139] Benedikt Diemer and Andrey V. Kravtsov. Dependence of the Outer Density Profiles of Halos on Their Mass Accretion Rate. *ApJ*, 789(1):1, July 2014.
- [140] Michael Dine, Willy Fischler, and Mark Srednicki. A simple solution to the strong CP problem with a harmless axion. *Phys. Lett. B*, 104(3):199–202, August 1981.
- [141] Scott Dodelson and Lawrence M. Widrow. Sterile neutrinos as dark matter. *Phys. Rev. Lett.*, 72(1):17–20, January 1994.
- [142] Elena D’Onghia and Andreas Burkert. The Failure of Self-Interacting Dark Matter to Solve the Overabundance of Dark Satellites and the Soft Core Question. *ApJ*, 586(1):12–16, Mar 2003.
- [143] Elena D’Onghia and George Lake. Small Dwarf Galaxies within Larger Dwarfs: Why Some Are Luminous while Most Go Dark. *ApJl*, 686(2):L61, Oct 2008.

- [144] Elena D'Onghia, Volker Springel, Lars Hernquist, and Dusan Keres. Substructure Depletion in the Milky Way Halo by the Disk. *ApJ*, 709(2):1138–1147, February 2010.
- [145] G. A. Dooley, A. H. G. Peter, J. L. Carlin, A. Frebel, K. Bechtol, and B. Willman. The predicted luminous satellite populations around SMC- and LMC-mass galaxies - a missing satellite problem around the LMC? *MNRAS*, 472:1060–1073, November 2017.
- [146] Gregory A. Dooley, Brendan F. Griffen, Phillip Zukin, Alexander P. Ji, Mark Vogelsberger, Lars E. Hernquist, and Anna Frebel. The Effects of Varying Cosmological Parameters on Halo Substructure. *ApJ*, 786(1):50, May 2014.
- [147] Gregory A. Dooley, Annika H. G. Peter, Mark Vogelsberger, Jesús Zavala, and Anna Frebel. Enhanced tidal stripping of satellites in the galactic halo from dark matter self-interactions. *MNRAS*, 461(1):710–727, Sep 2016.
- [148] A. Drlica-Wagner, K. Bechtol, S. Allam, D. L. Tucker, R. A. Gruendl, M. D. Johnson, A. R. Walker, D. J. James, D. L. Nidever, and K. A. G. Olsen. An Ultra-faint Galaxy Candidate Discovered in Early Data from the Magellanic Satellites Survey. *ApJL*, 833(1):L5, Dec 2016.
- [149] A. Drlica-Wagner, K. Bechtol, S. Mau, M. McNamara, E. O. Nadler, A. B. Pace, T. S. Li, A. Pieres, E. Rozo, J. D. Simon, A. R. Walker, R. H. Wechsler, T. M. C. Abbott, S. Allam, J. Annis, E. Bertin, D. Brooks, D. L. Burke, A. Carnero Rosell, M. Carrasco Kind, J. Carretero, M. Costanzi, L. N. da Costa, J. De Vicente, S. Desai, H. T. Diehl, P. Doel, T. F. Eifler, S. Everett, B. Flaugher, J. Frieman, J. García-Bellido, E. Gaztanaga, D. Gruen, R. A. Gruendl, J. Gschwend, G. Gutierrez, K. Honscheid, D. J. James, E. Krause, K. Kuehn, N. Kuropatkin, O. Lahav, M. A. G. Maia, J. L. Marshall, P. Melchior, F. Menanteau, R. Miquel, A. Palmese, A. A. Plazas, E. Sanchez, V. Scarpine, M. Schubnell, S. Serrano, I. Sevilla-Noarbe, M. Smith, E. Suchyta, G. Tarle, and DES Collaboration. Milky Way Satellite Census. I. The Observational Selection Function for Milky Way Satellites in DES Y3 and Pan-STARRS DR1. *ApJ*, 893(1):47, April 2020.
- [150] A. Drlica-Wagner, K. Bechtol, E. S. Rykoff, E. Luque, A. Queiroz, Y.-Y. Mao, R. H. Wechsler, J. D. Simon, B. Santiago, B. Yanny, E. Balbinot, S. Dodelson, A. Fausti Neto, D. J. James, T. S. Li, M. A. G. Maia, J. L. Marshall, A. Pieres, K. Stringer, A. R. Walker, T. M. C. Abbott, F. B. Abdalla, S. Allam, A. Benoit-Lévy, G. M. Bernstein, E. Bertin, D. Brooks, E. Buckley-Geer, D. L. Burke, A. Carnero Rosell, M. Carrasco Kind, J. Carretero, M. Crocce, L. N. da Costa, S. Desai, H. T. Diehl, J. P. Dietrich, P. Doel, T. F. Eifler, A. E. Evrard, D. A. Finley, B. Flaugher, P. Fosalba, J. Frieman, E. Gaztanaga, D. W. Gerdes, D. Gruen, R. A. Gruendl, G. Gutierrez, K. Honscheid, K. Kuehn, N. Kuropatkin, O. Lahav, P. Martini, R. Miquel, B. Nord, R. Ogando, A. A. Plazas, K. Reil, A. Roodman, M. Sako, E. Sanchez, V. Scarpine, M. Schubnell, I. Sevilla-Noarbe, R. C. Smith, M. Soares-Santos, F. Sobreira, E. Suchyta,

M. E. C. Swanson, G. Tarle, D. Tucker, V. Vikram, W. Wester, Y. Zhang, J. Zuntz, and DES Collaboration. Eight Ultra-faint Galaxy Candidates Discovered in Year Two of the Dark Energy Survey. *ApJ*, 813:109, November 2015.

[151] Alex Drlica-Wagner, Yao-Yuan Mao, Susmita Adhikari, Robert Armstrong, Arka Banerjee, Nilanjan Banik, Keith Bechtol, Simeon Bird, Kimberly K. Boddy, Ana Bonaca, Jo Bovy, Matthew R. Buckley, Esra Bulbul, Chihway Chang, George Chapline, Johann Cohen-Tanugi, Alessandro Cuoco, Francis-Yan Cyr-Racine, William A. Dawson, Ana Díaz Rivero, Cora Dvorkin, Denis Erkal, Christopher D. Fassnacht, Juan García-Bellido, Maurizio Giannotti, Vera Gluscevic, Nathan Golovich, David Hendel, Yashar D. Hezaveh, Shunsaku Horiuchi, M. James Jee, Manoj Kaplinghat, Charles R. Keeton, Sergey E. Koposov, Ting S. Li, Rachel Mandelbaum, Samuel D. McDermott, Mitch McNanna, Michael Medford, Manuel Meyer, Moniez Marc, Simona Murgia, Ethan O. Nadler, Lina Necib, Eric Nuss, Andrew B. Pace, Annika H. G. Peter, Daniel A. Polin, Chanda Prescod- Weinstein, Justin I. Read, Rogerio Rosenfeld, Nora Shipp, Joshua D. Simon, Tracy R. Slatyer, Oscar Straniero, Louis E. Strigari, Erik Tollerud, J. Anthony Tyson, Mei-Yu Wang, Risa H. Wechsler, David Wittman, Hai-Bo Yu, Gabrijela Zaharijas, Yacine Ali-Haïmoud, James Annis, Simon Birrer, Rahul Biswas, Jonathan Blazek, Alyson M. Brooks, Elizabeth Buckley-Geer, Regina Caputo, Eric Charles, Seth Digel, Scott Dodelson, Brenna Flaugher, Joshua Frieman, Eric Gawiser, Andrew P. Hearin, Renee Hložek, Bhuvnesh Jain, Tesla E. Jeltema, Savvas M. Koushiappas, Mariangela Lisanti, Marilena LoVerde, Siddharth Mishra-Sharma, Jeffrey A. Newman, Brian Nord, Erfan Nourbakhsh, Steven Ritz, Brant E. Robertson, Miguel A. Sánchez-Conde, Anže Slosar, Tim M. P. Tait, Aprajita Verma, Ricardo Vilalta, Christopher W. Walter, Brian Yanny, and Andrew R. Zentner. Probing the Fundamental Nature of Dark Matter with the Large Synoptic Survey Telescope. *arXiv e-prints*, page arXiv:1902.01055, February 2019.

[152] Richard D’Souza and Eric F. Bell. The infall of dwarf satellite galaxies are influenced by their host’s massive accretions. *MNRAS*, May 2021.

[153] Xiao Long Du. *Structure Formation with Ultralight Axion Dark Matter*. PhD thesis, Gottingen U., 2018.

[154] Xiaolong Du, Bodo Schwabe, Jens C. Niemeyer, and David Bürger. Tidal disruption of fuzzy dark matter subhalo cores. *Phys. Rev. D*, 97(6):063507, March 2018.

[155] Aaron A. Dutton and Andrea V. Macciò. Cold dark matter haloes in the Planck era: evolution of structural parameters for Einasto and NFW profiles. *MNRAS*, 441(4):3359–3374, July 2014.

[156] Cora Dvorkin, Kfir Blum, and Marc Kamionkowski. Constraining dark matter-baryon scattering with linear cosmology. *Phys. Rev. D*, 89:023519, Jan 2014.

- [157] Oliver D. Elbert, James S. Bullock, Shea Garrison-Kimmel, Miguel Rocha, Jose Oñorbe, and Annika H. G. Peter. Core formation in dwarf haloes with self-interacting dark matter: no fine-tuning necessary. *MNRAS*, 453(1):29–37, Oct 2015.
- [158] Oliver D. Elbert, James S. Bullock, Manoj Kaplinghat, Shea Garrison-Kimmel, Andrew S. Graus, and Miguel Rocha. A Testable Conspiracy: Simulating Baryonic Effects on Self-interacting Dark Matter Halos. *ApJ*, 853(2):109, Feb 2018.
- [159] Timon Emken and Chris Kouvaris. How blind are underground and surface detectors to strongly interacting dark matter? *Phys. Rev. D*, 97:115047, Jun 2018.
- [160] Christoph Engler, Annalisa Pillepich, Anna Pasquali, Dylan Nelson, Vicente Rodriguez-Gomez, Kun Ting Eddie Chua, Eva K. Grebel, Volker Springel, Federico Marinacci, Rainer Weinberger, Mark Vogelsberger, and Lars Hernquist. The abundance of satellites around Milky Way- and M31-like galaxies with the TNG50 simulation: a matter of diversity. *arXiv e-prints*, page arXiv:2101.12215, January 2021.
- [161] Wolfgang Enzi, Riccardo Murgia, Oliver Newton, Simona Vegetti, Carlos Frenk, Matteo Viel, Marius Cautun, Christopher D. Fassnacht, Matt Auger, Giulia Despali, John McKean, Léon V. E. Koopmans, and Mark Lovell. Joint constraints on thermal relic dark matter from a selection of astrophysical probes. *arXiv e-prints*, page arXiv:2010.13802, October 2020.
- [162] Adrienne L. Erickcek, Paul J. Steinhardt, Dan McCammon, and Patrick C. McGuire. Constraints on the interactions between dark matter and baryons from the x-ray quantum calorimetry experiment. *Phys. Rev. D*, 76:042007, Aug 2007.
- [163] D. Erkal, V. Belokurov, C. F. P. Laporte, S. E. Koposov, T. S. Li, C. J. Grillmair, N. Kallivayalil, A. M. Price-Whelan, N. W. Evans, K. Hawkins, D. Hendel, C. Mateu, J. F. Navarro, A. del Pino, C. T. Slater, S. T. Sohn, and Orphan Aspen Treasury Collaboration. The total mass of the Large Magellanic Cloud from its perturbation on the Orphan stream. *MNRAS*, 487(2):2685–2700, Aug 2019.
- [164] Denis Erkal and Vasily A. Belokurov. Limit on the LMC mass from a census of its satellites. *MNRAS*, 495(3):2554–2563, July 2020.
- [165] R. Errani, J. Penarrubia, and G. Tormen. Constraining the distribution of dark matter in dwarf spheroidal galaxies with stellar tidal streams. *MNRAS*, 449:L46–L50, Apr 2015.
- [166] Raphaël Errani and Julio F. Navarro. The asymptotic tidal remnants of cold dark matter subhalos. *MNRAS*, May 2021.
- [167] Miguel Escudero, Laura Lopez-Honorez, Olga Mena, Sergio Palomares-Ruiz, and Pablo Villanueva-Domingo. A fresh look into the interacting dark matter scenario. *JCAP*, 2018(6):007, Jun 2018.

- [168] Tilly A. Evans, Azadeh Fattahi, Alis J. Deason, and Carlos S. Frenk. How unusual is the Milky Way's assembly history? *MNRAS*, 497(4):4311–4321, July 2020.
- [169] Azadeh Fattahi, Julio F. Navarro, Carlos S. Frenk, Kyle A. Oman, Till Sawala, and Matthieu Schaller. Tidal stripping and the structure of dwarf galaxies in the Local Group. *MNRAS*, 476(3):3816–3836, May 2018.
- [170] Michael A. Fedderke, Peter W. Graham, and Surjeet Rajendran. Axion dark matter detection with CMB polarization. *Phys. Rev. D*, 100(1):015040, July 2019.
- [171] Jonathan L. Feng, Manoj Kaplinghat, Huitzu Tu, and Hai-Bo Yu. Hidden charged dark matter. *JCAP*, 2009(7):004, Jul 2009.
- [172] Davide Fiacconi, Piero Madau, Doug Potter, and Joachim Stadel. Cold Dark Matter Substructures in Early-type Galaxy Halos. *ApJ*, 824(2):144, June 2016.
- [173] Catherine E. Fielder, Yao-Yuan Mao, Jeffrey A. Newman, Andrew R. Zentner, and Timothy C. Licquia. Predictably missing satellites: subhalo abundances in Milky Way-like haloes. *MNRAS*, 486(4):4545–4568, Jul 2019.
- [174] C. Firmani, E. D’Onghia, V. Avila-Reese, G. Chincarini, and X. Hernández. Evidence of self-interacting cold dark matter from galactic to galaxy cluster scales. *MNRAS*, 315(3):L29–L32, Jul 2000.
- [175] A. Fitts, M. Boylan-Kolchin, J. S. Bullock, D. R. Weisz, K. El-Badry, C. Wheeler, C.-A. Faucher-Giguère, E. Quataert, P. F. Hopkins, D. Kereš, A. Wetzel, and C. C. Hayward. No assembly required: mergers are mostly irrelevant for the growth of low-mass dwarf galaxies. *MNRAS*, 479:319–331, September 2018.
- [176] A. Fitts, M. Boylan-Kolchin, O. D. Elbert, J. S. Bullock, P. F. Hopkins, J. Oñorbe, A. Wetzel, C. Wheeler, C.-A. Faucher-Giguère, D. Kereš, E. D. Skillman, and D. R. Weisz. fire in the field: simulating the threshold of galaxy formation. *MNRAS*, 471:3547–3562, nov 2017.
- [177] Alex Fitts, Michael Boylan-Kolchin, Brandon Bozek, James S. Bullock, Andrew Graus, Victor Robles, Philip F. Hopkins, Kareem El-Badry, Shea Garrison-Kimmel, Claude-André Faucher-Giguère, Andrew Wetzel, and Dušan Kereš. Dwarf galaxies in CDM, WDM, and SIDM: disentangling baryons and dark matter physics. *MNRAS*, 490(1):962–977, Nov 2019.
- [178] Seth Flaxman, Andrew Wilson, Daniel Neill, Hannes Nickisch, and Alex Smola. Fast kronecker inference in gaussian processes with non-gaussian likelihoods. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 607–616, Lille, France, 07–09 Jul 2015. PMLR.

- [179] Andreea S. Font, Ian G. McCarthy, and Vasily Belokurov. Can cosmological simulations capture the diverse satellite populations of observed Milky Way analogues? *arXiv e-prints*, page arXiv:2011.12974, November 2020.
- [180] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman. emcee: The MCMC Hammer. *PASP*, 125:306, March 2013.
- [181] T. K. Fritz, G. Battaglia, M. S. Pawlowski, N. Kallivayalil, R. van der Marel, S. T. Sohn, C. Brook, and G. Besla. Gaia DR2 proper motions of dwarf galaxies within 420 kpc. Orbits, Milky Way mass, tidal influences, planar alignments, and group infall. *A&A*, 619:A103, Nov 2018.
- [182] A. Bastidas Fry, F. Governato, A. Pontzen, T. Quinn, M. Tremmel, L. Anderson, H. Menon, A. M. Brooks, and J. Wadsley. All about baryons: revisiting SIDM predictions at small halo masses. *MNRAS*, 452(2):1468–1479, Sep 2015.
- [183] Gaia Collaboration, A. Helmi, F. van Leeuwen, P. J. McMillan, D. Massari, T. Antoja, A. C. Robin, L. Lindegren, U. Bastian, F. Arenou, and et al. Gaia Data Release 2. Kinematics of globular clusters and dwarf galaxies around the Milky Way. *A&A*, 616:A12, Aug 2018.
- [184] J. Gan, X. Kang, F. C. van den Bosch, and J. Hou. An improved model for the dynamical evolution of dark matter subhaloes. *MNRAS*, 408:2201–2212, November 2010.
- [185] S. Garrison-Kimmel, M. Boylan-Kolchin, J. S. Bullock, and K. Lee. ELVIS: Exploring the Local Volume in Simulations. *MNRAS*, 438:2578–2596, March 2014.
- [186] Shea Garrison-Kimmel, Michael Boylan-Kolchin, James S. Bullock, and Evan N. Kirby. Too big to fail in the Local Group. *MNRAS*, 444(1):222–236, October 2014.
- [187] Shea Garrison-Kimmel, James S. Bullock, Michael Boylan-Kolchin, and Emma Bardwell. Organized chaos: scatter in the relation between stellar mass and halo mass in small galaxies. *MNRAS*, 464(3):3108–3120, Jan 2017.
- [188] Shea Garrison-Kimmel, Philip F. Hopkins, Andrew Wetzel, James S. Bullock, Michael Boylan-Kolchin, Dušan Kereš, Claude-André Faucher-Giguère, Kareem El-Badry, Astrid Lamberts, Eliot Quataert, and Robyn Sanderson. The Local Group on FIRE: dwarf galaxy populations across a suite of hydrodynamic simulations. *MNRAS*, 487(1):1380–1399, Jul 2019.
- [189] Shea Garrison-Kimmel, Shunsaku Horiuchi, Kevork N. Abazajian, James S. Bullock, and Manoj Kaplinghat. Running with BICEP2: implications for small-scale problems in CDM. *MNRAS*, 444(1):961–970, October 2014.

- [190] Shea Garrison-Kimmel, Andrew Wetzel, James S. Bullock, Philip F. Hopkins, Michael Boylan-Kolchin, Claude-André Faucher-Giguère, Dušan Kereš, Eliot Quataert, Robyn E. Sanderson, Andrew S. Graus, and Tyler Kelley. Not so lumpy after all: modelling the depletion of dark matter subhaloes by Milky Way-like galaxies. *MNRAS*, 471:1709–1727, Oct 2017.
- [191] Raphaël Gavazzi, Tommaso Treu, Jason D. Rhodes, Léon V. E. Koopmans, Adam S. Bolton, Scott Burles, Richard J. Massey, and Leonidas A. Moustakas. The Sloan Lens ACS Survey. IV. The Mass Density Profile of Early-Type Galaxies out to 100 Effective Radii. *ApJ*, 667(1):176–190, September 2007.
- [192] M. Geha, B. Willman, J. D. Simon, L. E. Strigari, E. N. Kirby, D. R. Law, and J. Strader. The Least-Luminous Galaxy: Spectroscopy of the Milky Way Satellite Segue 1. *ApJ*, 692:1464–1475, February 2009.
- [193] Marla Geha, Risa H. Wechsler, Yao-Yuan Mao, Erik J. Tollerud, Benjamin Weiner, Rebecca Bernstein, Ben Hoyle, Sebastian Marchi, Phil J. Marshall, Ricardo Muñoz, and Yu Lu. The SAGA Survey. I. Satellite Galaxy Populations around Eight Milky Way Analogs. *ApJ*, 847(1):4, Sep 2017.
- [194] Anna Genina, Justin I. Read, Carlos S. Frenk, Shaun Cole, Alejandro Benitez-Llambay, Aaron D. Ludlow, Julio F. Navarro, Kyle A. Oman, and Andrew Robertson. To beta or not to beta: can higher-order Jeans analysis break the mass-anisotropy degeneracy in simulated dwarfs? *arXiv e-prints*, page arXiv:1911.09124, Nov 2019.
- [195] Daniel Gilman, Simon Birrer, Anna Nierenberg, Tommaso Treu, Xiaolong Du, and Andrew Benson. Warm dark matter chills out: constraints on the halo mass function and the free-streaming length of dark matter with eight quadruple-image strong gravitational lenses. *MNRAS*, 491(4):6077–6101, Feb 2020.
- [196] Daniel Gilman, Simon Birrer, Tommaso Treu, Anna Nierenberg, and Andrew Benson. Probing dark matter structure down to 10^7 solar masses: flux ratio statistics in gravitational lenses with line-of-sight haloes. *MNRAS*, 487(4):5721–5738, August 2019.
- [197] Daniel Gilman, Xiaolong Du, Andrew Benson, Simon Birrer, Anna Nierenberg, and Tommaso Treu. Constraints on the mass-concentration relation of cold dark matter halos with 11 strong gravitational lenses. *MNRAS*, 492(1):L12–L16, February 2020.
- [198] Vera Gluscevic and Kimberly K. Boddy. Constraints on scattering of kev–tev dark matter with protons in the early universe. *Phys. Rev. Lett.*, 121:081301, Aug 2018.
- [199] Antareep Gogoi, Prolay Chanda, and Subinoy Das. Degenerate dark matter micro-nuggets from light free-streaming sterile neutrino and the Hubble anomaly. *arXiv e-prints*, page arXiv:2005.11889, May 2020.

- [200] Martin Götz and Jesper Sommer-Larsen. Galaxy formation: Warm dark matter, missing satellites, and the angular momentum problem. *Ap&SS*, 284(2):341–344, April 2003.
- [201] F. Governato, A. Zolotov, A. Pontzen, C. Christensen, S. H. Oh, A. M. Brooks, T. Quinn, S. Shen, and J. Wadsley. Cuspy no more: how outflows affect the central dark matter and baryon distribution in Λ cold dark matter galaxies. *MNRAS*, 422:1231–1240, may 2012.
- [202] A. S. Graus, J. S. Bullock, T. Kelley, M. Boylan-Kolchin, S. Garrison-Kimmel, and Y. Qi. How low does it go? Too few Galactic satellites with standard reionization quenching. *MNRAS*, 488:4585–4595, October 2019.
- [203] Andrew S. Graus, James S. Bullock, Michael Boylan-Kolchin, and Anna M. Nierenberg. Through a Smoother Lens: An expected absence of LCDM substructure detections from hydrodynamic and dark matter only simulations. *MNRAS*, 480(1):1322–1332, October 2018.
- [204] A. M. Green, S. Hofmann, and D. J. Schwarz. The power spectrum of SUSY-CDM on subgalactic scales. *MNRAS*, 353:L23–L27, September 2004.
- [205] Sheridan B. Green, Frank C. van den Bosch, and Fangzhou Jiang. The tidal evolution of dark matter substructure – II. The impact of artificial disruption on subhalo mass functions and radial profiles. *arXiv e-prints*, page arXiv:2103.01227, March 2021.
- [206] B. F. Griffen, A. P. Ji, G. A. Dooley, F. A. Gómez, M. Vogelsberger, B. W. O’Shea, and A. Frebel. The Caterpillar Project: A Large Suite of Milky Way Sized Halos. *ApJ*, 818:10, February 2016.
- [207] C. J. Grillmair. Detection of a 60 deg-long Dwarf Galaxy Debris Stream. *ApJl*, 645:L37–L40, July 2006.
- [208] C. J. Grillmair. Four New Stellar Debris Streams in the Galactic Halo. *ApJ*, 693:1118–1127, March 2009.
- [209] Q. Guo, A. P. Cooper, C. Frenk, J. Helly, and W. A. Hellwing. The Milky Way system in Λ cold dark matter cosmological simulations. *MNRAS*, 454:550–559, November 2015.
- [210] Q. Guo and S. White. Numerical resolution limits on subhalo abundance matching. *MNRAS*, 437:3228–3235, February 2014.
- [211] Q. Guo, S. White, M. Boylan-Kolchin, G. De Lucia, G. Kauffmann, G. Lemson, C. Li, V. Springel, and S. Weinmann. From dwarf spheroidals to cD galaxies: simulating the galaxy population in a Λ CDM cosmology. *MNRAS*, 413:101–131, May 2011.

- [212] Steffen Hagstotz, Pablo F. de Salas, Stefano Gariazzo, Martina Gerbino, Massimiliano Lattanzi, Sunny Vagnozzi, Katherine Freese, and Sergio Pastor. Bounds on light sterile neutrino mass and mixing from cosmology and laboratory searches. *arXiv e-prints*, page arXiv:2003.02289, March 2020.
- [213] J. Han, S. Cole, C. S. Frenk, and Y. Jing. A unified model for the spatial and mass distribution of subhaloes. *MNRAS*, 457:1208–1223, April 2016.
- [214] J. R. Hargis, B. Willman, and A. H. G. Peter. Too Many, Too Few, or Just Right? The Predicted Number and Distribution of Milky Way Dwarf Galaxies. *ApJl*, 795:L13, November 2014.
- [215] David Harvey, Yves Revaz, Andrew Robertson, and Loic Hausammann. The impact of cored density profiles on the observable quantities of dwarf spheroidal galaxies. *MNRAS*, 481(1):L89–L93, Nov 2018.
- [216] Yashar D. Hezaveh, Neal Dalal, Gilbert Holder, Theodore Kisner, Michael Kuhlen, and Laurence Perreault Levasseur. Measuring the power spectrum of dark matter substructure using strong gravitational lensing. *JCAP*, 11(11):048–048, nov 2016.
- [217] Yashar D. Hezaveh, Neal Dalal, Daniel P. Marrone, Yao-Yuan Mao, Warren Morningstar, Di Wen, Roger D. Bland ford, John E. Carlstrom, Christopher D. Fassnacht, Gilbert P. Holder, Athol Kemball, Philip J. Marshall, Norman Murray, Laurence Perreault Levasseur, Joaquin D. Vieira, and Risa H. Wechsler. Detection of Lensing Substructure Using ALMA Observations of the Dusty Galaxy SDP.81. *ApJ*, 823:37, May 2016.
- [218] G. Hinshaw, D. Larson, E. Komatsu, D. N. Spergel, C. L. Bennett, J. Dunkley, M. R. Nolta, M. Halpern, R. S. Hill, N. Odegard, et al. Nine-year wilkinson microwave anisotropy probe (wmap) observations: Cosmological parameter results. *ApJ Suppl.*, 208(2):19, sep 2013.
- [219] S. R. Hinton. ChainConsumer. *J. Open Source Software*, 1:00045, August 2016.
- [220] D. Homma, M. Chiba, S. Okamoto, Y. Komiyama, M. Tanaka, M. Tanaka, M. N. Ishigaki, K. Hayashi, N. Arimoto, J. A. Garmilla, R. H. Lupton, M. A. Strauss, S. Miyazaki, S.-Y. Wang, and H. Murayama. Searches for new Milky Way satellites from the first two years of data of the Subaru/Hyper Suprime-Cam survey: Discovery of Cetus III. *PASJ*, 70:S18, January 2018.
- [221] Daisuke Homma, Masashi Chiba, Yutaka Komiyama, Masayuki Tanaka, Sakurako Okamoto, Mikito Tanaka, Miho N. Ishigaki, Kohei Hayashi, Nobuo Arimoto, Scott G. Carlsten, Robert H. Lupton, Michael A. Strauss, Satoshi Miyazaki, Gabriel Torrealba, Shiang-Yu Wang, and Hitoshi Murayama. Boötes. IV. A new Milky Way satellite discovered in the Subaru Hyper

Suprime-Cam Survey and implications for the missing satellite problem. *PASJ*, 71(5):94, October 2019.

[222] Daisuke Homma, Masashi Chiba, Sakurako Okamoto, Yutaka Komiyama, Masayuki Tanaka, Mikito Tanaka, Miho N. Ishigaki, Masayuki Akiyama, Nobuo Arimoto, José A. Garmilla, Robert H. Lupton, Michael A. Strauss, Hisanori Furusawa, Satoshi Miyazaki, Hitoshi Murayama, Atsushi J. Nishizawa, Masahiro Takada, Tomonori Usuda, and Shiang-Yu Wang. A New Milky Way Satellite Discovered in the Subaru/Hyper Suprime-Cam Survey. *ApJ*, 832(1):21, Nov 2016.

[223] P. F. Hopkins, D. Kereš, J. Oñorbe, C.-A. Faucher-Giguère, E. Quataert, N. Murray, and J. S. Bullock. Galaxies on fire (feedback in realistic environments): stellar feedback explains cosmologically inefficient star formation. *MNRAS*, 445:581–603, nov 2014.

[224] Philip F. Hopkins, Andrew Wetzel, Dušan Kereš, Claude-André Faucher-Giguère, Eliot Quataert, Michael Boylan-Kolchin, Norman Murray, Christopher C. Hayward, Shea Garrison-Kimmel, Cameron Hummels, Robert Feldmann, Paul Torrey, Xiangcheng Ma, Daniel Anglés-Alcázar, Kung-Yi Su, Matthew Orr, Denise Schmitz, Ivanna Escala, Robyn Sanderson, Michael Y. Grudić, Zachary Hafen, Ji-Hoon Kim, Alex Fitts, James S. Bullock, Coral Wheeler, T. K. Chan, Oliver D. Elbert, and Desika Narayanan. FIRE-2 simulations: physics versus numerics in galaxy formation. *MNRAS*, 480(1):800–863, October 2018.

[225] Shunsaku Horiuchi, Philip J. Humphrey, Jose Oñorbe, Kevork N. Abazajian, Manoj Kaplinghat, and Shea Garrison-Kimmel. Sterile neutrino dark matter bounds from galaxies of the Local Group. *Phys. Rev. D*, 89(2):025017, January 2014.

[226] J. W. Hsueh, W. Enzi, S. Vegetti, M. W. Auger, C. D. Fassnacht, G. Despali, L. V. E. Koopmans, and J. P. McKean. SHARP - VII. New constraints on the dark matter free-streaming properties and substructure abundance from gravitationally lensed quasars. *MNRAS*, 492(2):3047–3059, February 2020.

[227] Wayne Hu, Rennan Barkana, and Andrei Gruzinov. Fuzzy cold dark matter: The wave properties of ultralight particles. *Phys. Rev. Lett.*, 85:1158–1161, Aug 2000.

[228] Lam Hui, Jeremiah P. Ostriker, Scott Tremaine, and Edward Witten. Ultralight scalars as cosmological dark matter. *Phys. Rev. D*, 95:043541, Feb 2017.

[229] Ran Huo, Manoj Kaplinghat, Zhen Pan, and Hai-Bo Yu. Signatures of self-interacting dark matter in the matter power spectrum and the CMB. *Phys. Lett. B*, 783:76–81, August 2018.

[230] Alejandro Ibarra, Bradley J. Kavanagh, and Andreas Rappelt. Impact of substructure on local dark matter searches. *JCAP*, 2019(12):013, Dec 2019.

- [231] Masahiro Ibe and Hai-Bo Yu. Distinguishing dark matter annihilation enhancement scenarios via halo shapes. *Physics Letters B*, 692(2):70–73, Aug 2010.
- [232] Vid Iršič, Matteo Viel, Martin G. Haehnelt, James S. Bolton, Stefano Cristiani, George D. Becker, Valentina D’Odorico, Guido Cupani, Tae-Sun Kim, Trystyn A. M. Berg, Sebastian López, Sara Ellison, Lise Christensen, Kelly D. Denney, and Gábor Worseck. New constraints on the free-streaming of warm dark matter from intermediate and small scale lyman- α forest data. *Phys. Rev. D*, 96:023522, Jul 2017.
- [233] Vid Iršič, Matteo Viel, Martin G. Haehnelt, James S. Bolton, and George D. Becker. First Constraints on Fuzzy Dark Matter from Lyman- α Forest Data and Hydrodynamical Simulations. *Phys. Rev. Lett.*, 119(3):031302, July 2017.
- [234] M. J. Irwin, V. Belokurov, N. W. Evans, E. V. Ryan-Weber, J. T. A. de Jong, S. Koposov, D. B. Zucker, S. T. Hodgkin, G. Gilmore, P. Prema, L. Hebb, A. Begum, M. Fellhauer, P. C. Hewett, R. C. Kennicutt, Jr., M. I. Wilkinson, D. M. Bramich, S. Vidrih, H.-W. Rix, T. C. Beers, J. C. Barentine, H. Brewington, M. Harvanek, J. Krzesinski, D. Long, A. Nitta, and S. A. Snedden. Discovery of an Unusual Dwarf Galaxy in the Outskirts of the Milky Way. *ApJl*, 656:L13–L16, February 2007.
- [235] Tomoaki Ishiyama and Shin’ichiro Ando. The abundance and structure of subhaloes near the free streaming scale and their impact on indirect dark matter searches. *MNRAS*, 492(3):3662–3671, March 2020.
- [236] Tomoaki Ishiyama, Toshiyuki Fukushige, and Junichiro Makino. Variation of the Subhalo Abundance in Dark Matter Halos. *ApJ*, 696(2):2115–2125, May 2009.
- [237] Ž. Ivezić, S. M. Kahn, J. A. Tyson, B. Abel, E. Acosta, R. Allsman, D. Alonso, Y. AlSayyad, S. F. Anderson, J. Andrew, and et al. LSST: from Science Drivers to Reference Design and Anticipated Data Products. *ArXiv e-prints*, May 2008.
- [238] Ethan D. Jahn, Laura V. Sales, Andrew Wetzel, Michael Boylan-Kolchin, T. K. Chan, Kareem El-Badry, Alexandres Lazar, and James S. Bullock. Dark and luminous satellites of LMC-mass galaxies in the FIRE simulations. *MNRAS*, 489(4):5348–5364, November 2019.
- [239] H. Jeffreys. *Theory of probability*. Oxford University Press, Oxford, England, third edition, 1961.
- [240] P. Jethwa, D. Erkal, and V. Belokurov. A Magellanic origin of the DES dwarfs. *MNRAS*, 461(2):2212–2233, Sep 2016.
- [241] P. Jethwa, D. Erkal, and V. Belokurov. The upper bound on the lowest mass halo. *MNRAS*, 473:2060–2083, January 2018.

- [242] F. Jiang and F. C. van den Bosch. Comprehensive assessment of the too big to fail problem. *MNRAS*, 453:3575–3592, November 2015.
- [243] F. Jiang and F. C. van den Bosch. Statistics of dark matter substructure - I. Model and universal fitting functions. *MNRAS*, 458:2848–2869, May 2016.
- [244] Fangzhou Jiang, Avishai Dekel, Omer Kneller, Sharon Lapiner, Daniel Ceverino, Joel R. Primack, Sandra M. Faber, Andrea V. Macciò, Aaron A. Dutton, Shy Genel, and Rachel S. Somerville. Is the dark-matter halo spin a predictor of galaxy spin and size? *MNRAS*, 488(4):4801–4815, Oct 2019.
- [245] S. T. John and James Hensman. Large-Scale Cox Process Inference using Variational Fourier Features. *arXiv e-prints*, page arXiv:1804.01016, Apr 2018.
- [246] G. Jungman, M. Kamionkowski, and K. Griest. Supersymmetric dark matter. *Phys. Rep.*, 267:195–373, March 1996.
- [247] Felix Kahlhoefer, Manoj Kaplinghat, Tracy R. Slatyer, and Chih-Liang Wu. Diversity in density profiles of self-interacting dark matter satellite halos. *JCAP*, 2019(12):010, Dec 2019.
- [248] Felix Kahlhoefer, Kai Schmidt-Hoberg, Mads T. Frandsen, and Subir Sarkar. Colliding clusters and dark matter self-interactions. *MNRAS*, 437(3):2865–2881, Jan 2014.
- [249] Nitya Kallivayalil, Laura V. Sales, Paul Zivick, Tobias K. Fritz, Andrés Del Pino, Sangmo Tony Sohn, Gurtina Besla, Roeland P. van der Marel, Julio F. Navarro, and Elena Sacchi. The missing satellites of the magellanic clouds? gaia proper motions of the recently discovered ultra-faint galaxies. *ApJ*, 867(1):19, oct 2018.
- [250] Nitya Kallivayalil, Roeland P. van der Marel, Gurtina Besla, Jay Anderson, and Charles Alcock. Third-epoch Magellanic Cloud Proper Motions. I. Hubble Space Telescope/WFC3 Data and Orbit Implications. *ApJ*, 764(2):161, Feb 2013.
- [251] Ayuki Kamada, Manoj Kaplinghat, Andrew B. Pace, and Hai-Bo Yu. Self-Interacting Dark Matter Can Explain Diverse Galactic Rotation Curves. *Phys. Rev. Lett.*, 119(11):111102, Sep 2017.
- [252] Marc Kamionkowski and Andrew R. Liddle. The Dearth of Halo Dwarf Galaxies: Is There Power on Short Scales? *Phys. Rev. Lett.*, 84(20):4525–4528, May 2000.
- [253] Manoj Kaplinghat, Ryan E. Keeley, Tim Linden, and Hai-Bo Yu. Tying Dark Matter to Baryons with Self-Interactions. *Phys. Rev. Lett.*, 113(2):021302, Jul 2014.
- [254] Manoj Kaplinghat, Tao Ren, and Hai-Bo Yu. Dark Matter Cores and Cusps in Spiral Galaxies and their Explanations. *arXiv e-prints*, page arXiv:1911.00544, Nov 2019.

- [255] Manoj Kaplinghat, Sean Tulin, and Hai-Bo Yu. Dark Matter Halos as Particle Colliders: Unified Solution to Small-Scale Structure Puzzles from Dwarfs to Clusters. *Phys. Rev. Lett.*, 116(4):041302, Jan 2016.
- [256] Manoj Kaplinghat, Mauro Valli, and Hai-Bo Yu. Too big to fail in light of Gaia. *MNRAS*, 490(1):231–242, Nov 2019.
- [257] Harley Katz, Marius Ramsøy, Joakim Rosdahl, Taysun Kimm, Jérémie Blaizot, Martin G. Haehnelt, Léo Michel-Dansac, Thibault Garel, Clotilde Laigle, Julien Devriendt, and Adrienne Slyz. How to Quench a Dwarf Galaxy: The Impact of Inhomogeneous Reionization on Dwarf Galaxies and Cosmic Filaments. *MNRAS*, March 2020.
- [258] N. Katz and S. D. M. White. Hierarchical galaxy formation - overmerging and the formation of an x-ray cluster. *ApJ*, 412:455–478, aug 1993.
- [259] T. Kelley, J. S. Bullock, S. Garrison-Kimmel, M. Boylan-Kolchin, M. S. Pawlowski, and A. S. Graus. Phat ELVIS: The inevitable effect of the Milky Way’s disc on its dark matter subhaloes. *MNRAS*, 487:4409–4423, August 2019.
- [260] Rachel Kennedy, Carlos Frenk, Shaun Cole, and Andrew Benson. Constraining the warm dark matter particle mass with Milky Way satellites. *MNRAS*, 442(3):2487–2495, 06 2014.
- [261] D. Kim and H. Jerjen. Horologium II: A Second Ultra-faint Milky Way Satellite in the Horologium Constellation. *ApJl*, 808:L39, August 2015.
- [262] J.-h. Kim, T. Abel, O. Agertz, G. L. Bryan, D. Ceverino, C. Christensen, C. Conroy, A. Dekel, N. Y. Gnedin, N. J. Goldbaum, J. Guedes, O. Hahn, A. Hobbs, P. F. Hopkins, C. B. Hummels, F. Iannuzzi, D. Keres, A. Klypin, A. V. Kravtsov, M. R. Krumholz, M. Kuhlen, S. N. Leitner, P. Madau, L. Mayer, C. E. Moody, K. Nagamine, M. L. Norman, J. Onorbe, B. W. O’Shea, A. Pillepich, J. R. Primack, T. Quinn, J. I. Read, B. E. Robertson, M. Rocha, D. H. Rudd, S. Shen, B. D. Smith, A. S. Szalay, R. Teyssier, R. Thompson, K. Todoroki, M. J. Turk, J. W. Wadsley, J. H. Wise, A. Zolotov, and t. AGORA Collaboration29. The agora high-resolution galaxy simulations comparison project. *ApJs*, 210:14, jan 2014.
- [263] Jihn E. Kim. Weak-interaction singlet and strong CP invariance. *Phys. Rev. Lett.*, 43(2):103–107, July 1979.
- [264] Stacy Y. Kim, Annika H. G. Peter, and Jonathan R. Hargis. Missing satellites problem: Completeness corrections to the number of satellite galaxies in the milky way are consistent with cold dark matter predictions. *Phys. Rev. Lett.*, 121:211302, Nov 2018.
- [265] A. Klypin, I. Karachentsev, D. Makarov, and O. Nasonova. Abundance of field galaxies. *MNRAS*, 454:1798–1810, December 2015.

- [266] A. Klypin, A. V. Kravtsov, O. Valenzuela, and F. Prada. Where are the missing galactic satellites? *ApJ*, 522:82–92, sep 1999.
- [267] A. Knebe, S. R. Knollmann, S. I. Muldrew, F. R. Pearce, M. A. Aragon-Calvo, Y. Ascasibar, P. S. Behroozi, D. Ceverino, S. Colombi, J. Diemand, K. Dolag, B. L. Falck, P. Fasel, J. Gardner, S. Gottlöber, C.-H. Hsu, F. Iannuzzi, A. Klypin, Z. Lukić, M. Maciejewski, C. McBride, M. C. Neyrinck, S. Planelles, D. Potter, V. Quilis, Y. Rasera, J. I. Read, P. M. Ricker, F. Roy, V. Springel, J. Stadel, G. Stinson, P. M. Sutter, V. Turchaninov, D. Tweed, G. Yepes, and M. Zemp. Haloes gone MAD: The Halo-Finder Comparison Project. *MNRAS*, 415:2293–2318, August 2011.
- [268] S. R. Knollmann and A. Knebe. Ahf: Amiga’s halo finder. *ApJs*, 182:608–624, jun 2009.
- [269] Jun Koda and Paul R. Shapiro. Gravothermal collapse of isolated self-interacting dark matter haloes: N-body simulation versus the fluid model. *MNRAS*, 415(2):1125–1137, Aug 2011.
- [270] Rohit Kondapally, George A. Russell, Christopher J. Conselice, and Samantha J. Penny. Faint satellite population of the NGC-3175 Group - a Local Group analogue. *MNRAS*, 481(2):1759–1773, Dec 2018.
- [271] S. Koposov, V. Belokurov, N. W. Evans, P. C. Hewett, M. J. Irwin, G. Gilmore, D. B. Zucker, H.-W. Rix, M. Fellhauer, E. F. Bell, and E. V. Glushkova. The Luminosity Function of the Milky Way Satellites. *ApJ*, 686:279–291, October 2008.
- [272] S. E. Koposov, V. Belokurov, G. Torrealba, and N. W. Evans. Beasts of the Southern Wild: Discovery of Nine Ultra Faint Satellites in the Vicinity of the Magellanic Clouds. *ApJ*, 805:130, June 2015.
- [273] S. E. Koposov, J. Yoo, H.-W. Rix, D. H. Weinberg, A. V. Macciò, and J. M. Escudé. A Quantitative Explanation of the Observed Population of Milky Way Satellite Galaxies. *ApJ*, 696:2179–2194, May 2009.
- [274] Sergey E. Koposov, Matthew G. Walker, Vasily Belokurov, Andrew R. Casey, Alex Geringer-Sameth, Dougal Mackey, Gary Da Costa, Denis Erkal, Prashin Jethwa, Mario Mateo, Edward W. Olszewski, and John I. Bailey. Snake in the Clouds: a new nearby dwarf galaxy in the Magellanic bridge*. *MNRAS*, 479:5343–5361, Oct 2018.
- [275] A. V. Kravtsov. The Size-Virial Radius Relation of Galaxies. *ApJL*, 764:L31, February 2013.
- [276] A. V. Kravtsov, O. Y. Gnedin, and A. A. Klypin. The Tumultuous Lives of Galactic Dwarfs and the Missing Satellites Problem. *ApJ*, 609:482–497, July 2004.
- [277] Andrey Kravtsov. Dark Matter Substructure and Dwarf Galactic Satellites. *Advances in Astronomy*, 2010:281913, January 2010.

- [278] Christina D. Kreisch, Francis-Yan Cyr-Racine, and Olivier Doré. The Neutrino Puzzle: Anomalies, Interactions, and Cosmological Tensions. *Phys. Rev. D*, 101(12):123505, 2020.
- [279] Gordan Krnjaic and Samuel D. McDermott. Implications of BBN bounds for cosmic ray upscattered dark matter. *Phys. Rev. D*, 101(12):123022, June 2020.
- [280] M. Kuhlen, P. Madau, and M. R. Krumholz. Dwarf Galaxy Formation with H₂-regulated Star Formation. II. Gas-rich Dark Galaxies at Redshift 2.5. *ApJ*, 776:34, October 2013.
- [281] Janis Kummer, Felix Kahlhoefer, and Kai Schmidt-Hoberg. Effective description of dark matter self-interactions in small dark matter haloes. *MNRAS*, 474(1):388–399, Feb 2018.
- [282] B. P. M. Laevens, N. F. Martin, R. A. Ibata, H.-W. Rix, E. J. Bernard, E. F. Bell, B. Sesar, A. M. N. Ferguson, E. F. Schlafly, C. T. Slater, W. S. Burgett, K. C. Chambers, H. Flewelling, K. A. Hodapp, N. Kaiser, R.-P. Kudritzki, R. H. Lupton, E. A. Magnier, N. Metcalfe, J. S. Morgan, P. A. Price, J. L. Tonry, R. J. Wainscoat, and C. Waters. A New Faint Milky Way Satellite Discovered in the Pan-STARRS1 3 π Survey. *ApJl*, 802:L18, April 2015.
- [283] Benjamin P. M. Laevens, Nicolas F. Martin, Edouard J. Bernard, Edward F. Schlafly, Branimir Sesar, Hans-Walter Rix, Eric F. Bell, Annette M. N. Ferguson, Colin T. Slater, William E. Sweeney, Rosemary F. G. Wyse, Avon P. Huxor, William S. Burgett, Kenneth C. Chambers, Peter W. Draper, Klaus A. Hodapp, Nicholas Kaiser, Eugene A. Magnier, Nigel Metcalfe, John L. Tonry, Richard J. Wainscoat, and Christopher Waters. Sagittarius II, Draco II and Laevens 3: Three New Milky Way Satellites Discovered in the Pan-STARRS 1 3 π Survey. *ApJ*, 813:44, November 2015.
- [284] Lachlan Lancaster, Cara Giovanetti, Philip Mocz, Yonatan Kahn, Mariangela Lisanti, and David N. Spergel. Dynamical friction in a Fuzzy Dark Matter universe. *JCAP*, 01(1):001, January 2020.
- [285] R. Laureijs, J. Amiaux, S. Arduini, J. L. Auguères, J. Brinchmann, R. Cole, M. Cropper, C. Dabin, L. Duvet, A. Ealet, B. Garilli, P. Gondoin, L. Guzzo, J. Hoar, H. Hoekstra, R. Holmes, T. Kitching, T. Maciaszek, Y. Mellier, F. Pasian, W. Percival, J. Rhodes, G. Saavedra Criado, M. Sauvage, R. Scaramella, L. Valenziano, S. Warren, R. Bender, F. Castander, A. Cimatti, O. Le Fèvre, H. Kurki-Suonio, M. Levi, P. Lilje, G. Meylan, R. Nichol, K. Pedersen, V. Popa, R. Rebolo Lopez, H. W. Rix, H. Rottgering, W. Zeilinger, F. Grupp, P. Hudelot, R. Massey, M. Meneghetti, L. Miller, S. Paltani, S. Paulin-Henriksson, S. Pires, C. Saxton, T. Schrabback, G. Seidel, J. Walsh, N. Aghanim, L. Amendola, J. Bartlett, C. Baccigalupi, J. P. Beaulieu, K. Benabed, J. G. Cuby, D. Elbaz, P. Fosalba, G. Gavazzi, A. Helmi, I. Hook, M. Irwin, J. P. Kneib, M. Kunz, F. Mannucci, L. Moscardini, C. Tao, R. Teyssier, J. Weller,

G. Zamorani, M. R. Zapatero Osorio, O. Boulade, J. J. Foumond, A. Di Giorgio, P. Guttridge, A. James, M. Kemp, J. Martignac, A. Spencer, D. Walton, T. Blümchen, C. Bonoli, F. Bortoletto, C. Cerna, L. Corcione, C. Fabron, K. Jahnke, S. Ligori, F. Madrid, L. Martin, G. Morgante, T. Pamplona, E. Prieto, M. Riva, R. Toledo, M. Trifoglio, F. Zerbi, F. Abdalla, M. Douspis, C. Grenet, S. Borgani, R. Bouwens, F. Courbin, J. M. Delouis, P. Dubath, A. Fontana, M. Frailis, A. Grazian, J. Koppenhöfer, O. Mansutti, M. Melchior, M. Mignoli, J. Mohr, C. Neissner, K. Noddle, M. Poncet, M. Scodeggio, S. Serrano, N. Shane, J. L. Starck, C. Surace, A. Taylor, G. Verdoes-Kleijn, C. Vuerli, O. R. Williams, A. Zacchei, B. Altieri, I. Escudero Sanz, R. Kohley, T. Oosterbroek, P. Astier, D. Bacon, S. Bardelli, C. Baugh, F. Bellagamba, C. Benoist, D. Bianchi, A. Biviano, E. Branchini, C. Carbone, V. Cardone, D. Clements, S. Colombi, C. Conselice, G. Cresci, N. Deacon, J. Dunlop, C. Fedeli, F. Fontanot, P. Franzetti, C. Giocoli, J. Garcia-Bellido, J. Gow, A. Heavens, P. Hewett, C. Heymans, A. Holland, Z. Huang, O. Ilbert, B. Joachimi, E. Jennings, E. Kerins, A. Kiessling, D. Kirk, R. Kotak, O. Krause, O. Lahav, F. van Leeuwen, J. Lesgourgues, M. Lombardi, M. Magliocchetti, K. Maguire, E. Majerotto, R. Maoli, F. Marulli, S. Maurogordato, H. McCracken, R. McLure, A. Melchiorri, A. Merson, M. Moresco, M. Nonino, P. Norberg, J. Peacock, R. Pello, M. Penny, V. Pettorino, C. Di Porto, L. Pozzetti, C. Quercellini, M. Radovich, A. Rassat, N. Roche, S. Ronayette, E. Rossetti, B. Sartoris, P. Schneider, E. Semboloni, S. Serjeant, F. Simpson, C. Skordis, G. Smadja, S. Smartt, P. Spano, S. Spiro, M. Sullivan, A. Tilquin, R. Trotta, L. Verde, Y. Wang, G. Williger, G. Zhao, J. Zoubian, and E. Zucca. Euclid Definition Study Report. *arXiv e-prints*, page arXiv:1110.3193, October 2011.

[286] Alexandres Lazar and James S. Bullock. Accurate mass estimates from the proper motions of dispersion-supported galaxies. *MNRAS*, 493(4):5825–5837, March 2020.

[287] Alexandres Lazar, James S. Bullock, Michael Boylan-Kolchin, Robert Feldmann, Onur Çatmabacak, and Leonidas Moustakas. Out of sight, out of mind? The impact of correlated clustering in substructure lensing. *MNRAS*, 502(4):6064–6079, April 2021.

[288] B. V. Lehmann, Y.-Y. Mao, M. R. Becker, S. W. Skillman, and R. H. Wechsler. The Concentration Dependence of the Galaxy-Halo Connection: Modeling Assembly Bias with Abundance Matching. *ApJ*, 834:37, January 2017.

[289] Ming Li, Liang Gao, and Jie Wang. The abundance of satellite galaxies in the inner region of Λ CDM Milky Way sized haloes. *MNRAS*, 483:2000–2006, Feb 2019.

[290] Y.-S. Li, G. De Lucia, and A. Helmi. On the nature of the Milky Way satellites. *MNRAS*, 401:2036–2052, January 2010.

- [291] Zhao-Zhou Li, Yong-Zhong Qian, Jiaxin Han, Ting S. Li, Wenting Wang, and Y. P. Jing. Constraining the Milky Way Mass Profile with Phase-Space Distribution of Satellite Galaxies. *arXiv e-prints*, page arXiv:1912.02086, Dec 2019.
- [292] Zhao-Zhou Li, Yong-Zhong Qian, Jiaxin Han, Wenting Wang, and Y. P. Jing. A Versatile and Accurate Method for Halo Mass Determination from Phase-space Distribution of Satellite Galaxies. *ApJ*, 886(1):69, Nov 2019.
- [293] T. C. Licquia, J. A. Newman, and J. Brinchmann. Unveiling the Milky Way: A New Technique for Determining the Optical Color and Luminosity of Our Galaxy. *ApJ*, 809:96, August 2015.
- [294] Abraham Loeb and Neal Weiner. Cores in Dwarf Galaxies from Dark Matter with a Yukawa Potential. *Phys. Rev. Lett.*, 106(17):171302, Apr 2011.
- [295] J. Loveday, P. Norberg, I. K. Baldry, J. Bland-Hawthorn, S. Brough, M. J. I. Brown, S. P. Driver, L. S. Kelvin, and S. Phillipps. Galaxy and mass assembly (gama): maximum-likelihood determination of the luminosity function and its evolution. *MNRAS*, 451:1540–1552, aug 2015.
- [296] M. R. Lovell. The halo mass function in alternative dark matter models. *Mon. Not. R. Astron. Soc.*, 493(1):L11–L15, March 2020.
- [297] Mark R. Lovell. Toward a General Parameterization of the Warm Dark Matter Halo Mass Function. *ApJ*, 897(2):147, July 2020.
- [298] Mark R. Lovell, Marius Cautun, Carlos S. Frenk, Wojciech A. Hellwing, and Oliver Newton. The spatial distribution of Milky Way satellites, gaps in streams and the nature of dark matter. *arXiv e-prints*, page arXiv:2104.03322, April 2021.
- [299] Mark R. Lovell, Carlos S. Frenk, Vincent R. Eke, Adrian Jenkins, Liang Gao, and Tom Theuns. The properties of warm dark matter haloes. *MNRAS*, 439(1):300–317, Mar 2014.
- [300] Mark R. Lovell, Violeta Gonzalez-Perez, Sownak Bose, Alexey Boyarsky, Shaun Cole, Carlos S. Frenk, and Oleg Ruchayskiy. Addressing the too big to fail problem with baryon physics and sterile neutrino dark matter. *MNRAS*, 468(3):2836–2849, Jul 2017.
- [301] LSST Science Collaboration, P. A. Abell, J. Allison, S. F. Anderson, J. R. Andrew, J. R. P. Angel, L. Armus, D. Arnett, S. J. Asztalos, T. S. Axelrod, and et al. Lsst science book, version 2.0. *ArXiv e-prints*, dec 2009.
- [302] Y. Lu, A. Benson, Y.-Y. Mao, S. Tonnesen, A. H. G. Peter, A. R. Wetzel, M. Boylan-Kolchin, and R. H. Wechsler. The connection between the host halo and the satellite galaxies of the milky way. *ApJ*, 830:59, oct 2016.

- [303] Y. Lu, A. Benson, A. Wetzel, Y.-Y. Mao, S. Tonnesen, A. H. G. Peter, M. Boylan-Kolchin, and R. H. Wechsler. The importance of preventive feedback: Inference from observations of the stellar masses and metallicities of milky way dwarf galaxies. *ApJ*, 846:66, sep 2017.
- [304] Ragnhild Lunnan, Mark Vogelsberger, Anna Frebel, Lars Hernquist, Adam Lidz, and Michael Boylan-Kolchin. The Effects of Patchy Reionization on Satellite Galaxies of the Milky Way. *ApJ*, 746(1):109, Feb 2012.
- [305] E. Luque, A. Queiroz, B. Santiago, A. Pieres, E. Balbinot, K. Bechtol, A. Drlica-Wagner, A. Fausti Neto, L. N. da Costa, M. A. G. Maia, B. Yanny, T. Abbott, S. Allam, A. Benoit-Lévy, E. Bertin, D. Brooks, E. Buckley-Geer, D. L. Burke, A. Carnero Rosell, M. Carrasco Kind, J. Carretero, C. E. Cunha, S. Desai, H. T. Diehl, J. P. Dietrich, T. F. Eifler, D. A. Finley, B. Flaugher, P. Fosalba, J. Frieman, D. W. Gerdes, D. Gruen, G. Gutierrez, K. Honscheid, D. J. James, K. Kuehn, N. Kuropatkin, O. Lahav, T. S. Li, M. March, J. L. Marshall, P. Martini, R. Miquel, E. Neilsen, R. C. Nichol, B. Nord, R. Ogando, A. A. Plazas, A. K. Romer, A. Roodman, E. Sanchez, V. Scarpine, M. Schubnell, I. Sevilla-Noarbe, R. C. Smith, M. Soares-Santos, F. Sobreira, E. Suchyta, M. E. C. Swanson, G. Tarle, J. Thaler, D. Tucker, A. R. Walker, and Y. Zhang. Digging deeper into the Southern skies: a compact Milky Way companion discovered in first-year Dark Energy Survey data. *MNRAS*, 458(1):603–612, May 2016.
- [306] D. Lynden-Bell. Dwarf galaxies and globular clusters in high velocity hydrogen streams. *MNRAS*, 174:695–710, Mar 1976.
- [307] Chung-Pei Ma and Edmund Bertschinger. Cosmological Perturbation Theory in the Synchronous and Conformal Newtonian Gauges. *ApJ*, 455:7, December 1995.
- [308] Karime Maamari, Vera Gluscevic, Kimberly K. Boddy, Ethan O. Nadler, and Risa H. Wechsler. Bounds on Velocity-dependent Dark Matter-Proton Scattering from Milky Way Satellite Abundance. *ApJL*, 907(2):L46, February 2021.
- [309] A. V. Macciò and F. Fontanot. How cold is dark matter? Constraints from Milky Way satellites. *MNRAS*, 404:L16–L20, May 2010.
- [310] A. V. Macciò, X. Kang, F. Fontanot, R. S. Somerville, S. Koposov, and P. Monaco. Luminosity function and radial distribution of Milky Way satellites in a Λ CDM Universe. *MNRAS*, 402:1995–2008, March 2010.
- [311] A. V. Macciò, S. M. Udrescu, A. A. Dutton, A. Obreja, L. Wang, G. R. Stinson, and X. Kang. NIHAO X: Reconciling the local galaxy velocity function with Cold Dark Matter via mock HI observations. *ArXiv e-prints*, July 2016.

- [312] Andrea V. Macciò, Sinziana Paduroiu, Donnino Anderhalden, Aurel Schneider, and Ben Moore. Cores in warm dark matter haloes: a Catch 22 problem. *MNRAS*, 424(2):1105–1112, August 2012.
- [313] Prasanta Chandra Mahalanobis. On the generalized distance in statistics. *Proceedings of the National Institute of Sciences (Calcutta)*, 2:49–55, 1936.
- [314] M. Shafi Mahdawi and Glennys R. Farrar. Constraints on Dark Matter with a moderately large and velocity-dependent DM-nucleon cross-section. *JCAP*, 10(10):007, October 2018.
- [315] Y.-Y. Mao, M. Williamson, and R. H. Wechsler. The dependence of subhalo abundance on halo concentration. *ApJ*, 810:21, sep 2015.
- [316] Yao-Yuan Mao, Marla Geha, Risa H. Wechsler, Benjamin Weiner, Erik J. Tollerud, Ethan O. Nadler, and Nitya Kallivayalil. The SAGA Survey. II. Building a Statistical Sample of Satellite Systems around Milky Way-like Galaxies. *ApJ*, 907(2):85, February 2021.
- [317] David J. E. Marsh. Axion cosmology. *Phys. Rep.*, 643:1–79, July 2016.
- [318] N. F. Martin, D. L. Nidever, G. Besla, K. Olsen, A. R. Walker, A. K. Vivas, R. A. Gruendl, C. C. Kaleida, R. R. Muñoz, R. D. Blum, A. Saha, B. C. Conn, E. F. Bell, Y.-H. Chu, M.-R. L. Cioni, T. J. L. de Boer, C. Gallart, S. Jin, A. Kunder, S. R. Majewski, D. Martinez-Delgado, A. Monachesi, M. Monelli, L. Monteagudo, N. E. D. Noël, E. W. Olszewski, G. S. Stringfellow, R. P. van der Marel, and D. Zaritsky. Hydra II: A Faint and Compact Milky Way Dwarf Galaxy Found in the Survey of the Magellanic Stellar History. *ApJl*, 804:L5, May 2015.
- [319] Nicolas F. Martin, Rodrigo A. Ibata, Geraint F. Lewis, Alan McConnachie, Arif Babul, Nicholas F. Bate, Edouard Bernard, Scott C. Chapman, Michelle M. L. Collins, Anthony R. Conn, Denija Crnojević, Mark A. Fardal, Annette M. N. Ferguson, Michael Irwin, A. Dougal Mackey, Brendan McMonigal, Julio F. Navarro, and R. Michael Rich. The PAndAS View of the Andromeda Satellite System. II. Detailed Properties of 23 M31 Dwarf Spheroidal Galaxies. *ApJ*, 833(2):167, December 2016.
- [320] S. Mau, W. Cerny, A. B. Pace, Y. Choi, A. Drlica-Wagner, L. Santana-Silva, A. H. Riley, D. Erkal, G. S. Stringfellow, M. Adamów, J. L. Carlin, R. A. Gruendl, D. Hernandez-Lang, N. Kuropatkin, T. S. Li, C. E. Martínez-Vázquez, E. Morganson, B. Mutlu-Pakdil, E. H. Neilsen, D. L. Nidever, K. A. G. Olsen, D. J. Sand, E. J. Tollerud, D. L. Tucker, B. Yanny, A. Zenteno, S. Allam, W. A. Barkhouse, K. Bechtol, E. F. Bell, P. Balaji, D. Crnojević, J. Esteves, P. S. Ferguson, C. Gallart, A. K. Hughes, D. J. James, P. Jethwa, L. C. Johnson, K. Kuehn, S. Majewski, Y. Y. Mao, P. Massana, M. McNanna, A. Monachesi, E. O. Nadler, N. E. D. Noël, A. Palmese, F. Paz-Chinchon, A. Pieres, J. Sanchez, N. Shipp, J. D. Simon, M. Soares-Santos, K. Tavangar, R. P. van der Marel, A. K. Vivas, A. R. Walker, and R. H.

Wechsler. Two Ultra-faint Milky Way Stellar Systems Discovered in Early Data from the DECam Local Volume Exploration Survey. *ApJ*, 890(2):136, February 2020.

[321] A. W. McConnachie. The Observed Properties of Dwarf Galaxies in and around the Local Group. *AJ*, 144:4, July 2012.

[322] Carisa Miller, Adrienne Erickcek, and Riccardo Murgia. Constraining Nonthermal Dark Matter's Impact on the Matter Power Spectrum. *arXiv e-prints*, page arXiv:1908.10369, Aug 2019.

[323] Quinn E. Minor, Manoj Kaplinghat, Tony H. Chan, and Emily Simon. Inferring the concentration of dark matter subhalos perturbing strongly lensed images. *arXiv e-prints*, page arXiv:2011.10629, November 2020.

[324] G. Monari, B. Famaey, I. Carrillo, T. Piffl, M. Steinmetz, R. F. G. Wyse, F. Anders, C. Chiappini, and K. Janßen. The escape speed curve of the Galaxy obtained from Gaia DR2 implies a heavy Milky Way. *A&A*, 616:L9, Aug 2018.

[325] B. Moore, S. Ghigna, F. Governato, G. Lake, T. Quinn, J. Stadel, and P. Tozzi. Dark matter substructure within galactic halos. *ApJl*, 524:L19–L22, oct 1999.

[326] Surhud More, Benedikt Diemer, and Andrey V. Kravtsov. The Splashback Radius as a Physical Halo Boundary and the Growth of Halo Mass. *ApJ*, 810(1):36, Sep 2015.

[327] J. A. Muñoz, P. Madau, A. Loeb, and J. Diemand. Probing the epoch of reionization with Milky Way satellites. *MNRAS*, 400:1593–1602, December 2009.

[328] F. Munshi, A. M. Brooks, E. Applebaum, D. R. Weisz, F. Governato, and T. R. Quinn. Going, going, gone dark: Quantifying the scatter in the faintest dwarf galaxies. *ArXiv e-prints*, May 2017.

[329] F. Munshi, F. Governato, A. M. Brooks, C. Christensen, S. Shen, S. Loebman, B. Moster, T. Quinn, and J. Wadsley. Reproducing the Stellar Mass/Halo Mass Relation in Simulated Λ CDM Galaxies: Theory versus Observational Estimates. *ApJ*, 766:56, March 2013.

[330] Ferah Munshi, Alyson M. Brooks, Charlotte Christensen, Elaad Applebaum, Kelly Holley-Bockelmann, Thomas R. Quinn, and James Wadsley. Dancing in the Dark: Uncertainty in Ultrafaint Dwarf Galaxy Predictions from Cosmological Simulations. *ApJ*, 874(1):40, Mar 2019.

[331] Iain Murray. *Advances in Markov chain Monte Carlo methods*. PhD thesis, Gatsby Computational Neuroscience Unit, University College London, 2007.

[332] E. O. Nadler, A. Drlica-Wagner, K. Bechtol, S. Mau, R. H. Wechsler, V. Gluscevic, K. Boddy, A. B. Pace, T. S. Li, M. McNanna, A. H. Riley, J. García-Bellido, Y. Y. Mao, G. Green, D. L. Burke, A. Peter, B. Jain, T. M. C. Abbott, M. Aguena, S. Allam, J. Annis, S. Avila, D. Brooks, M. Carrasco Kind, J. Carretero, M. Costanzi, L. N. da Costa, J. De Vicente, S. Desai, H. T. Diehl, P. Doel, S. Everett, A. E. Evrard, B. Flaugher, J. Frieman, D. W. Gerdes, D. Gruen, R. A. Gruendl, J. Gschwend, G. Gutierrez, S. R. Hinton, K. Honscheid, D. Huterer, D. J. James, E. Krause, K. Kuehn, N. Kuropatkin, O. Lahav, M. A. G. Maia, J. L. Marshall, F. Menanteau, R. Miquel, A. Palmese, F. Paz-Chinchón, A. A. Plazas, A. K. Romer, E. Sanchez, V. Scarpine, S. Serrano, I. Sevilla-Noarbe, M. Smith, M. Soares-Santos, E. Suchyta, M. E. C. Swanson, G. Tarle, D. L. Tucker, A. R. Walker, W. Wester, and DES Collaboration. Constraints on Dark Matter Properties from Observations of Milky Way Satellite Galaxies. *Phys. Rev. Lett.*, 126(9):091101, March 2021.

[333] E. O. Nadler, R. H. Wechsler, K. Bechtol, Y. Y. Mao, G. Green, A. Drlica-Wagner, M. McNanna, S. Mau, A. B. Pace, J. D. Simon, A. Kravtsov, S. Dodelson, T. S. Li, A. H. Riley, M. Y. Wang, T. M. C. Abbott, M. Aguena, S. Allam, J. Annis, S. Avila, G. M. Bernstein, E. Bertin, D. Brooks, D. L. Burke, A. Carnero Rosell, M. Carrasco Kind, J. Carretero, M. Costanzi, L. N. da Costa, J. De Vicente, S. Desai, A. E. Evrard, B. Flaugher, P. Fosalba, J. Frieman, J. García-Bellido, E. Gaztanaga, D. W. Gerdes, D. Gruen, J. Gschwend, G. Gutierrez, W. G. Hartley, S. R. Hinton, K. Honscheid, E. Krause, K. Kuehn, N. Kuropatkin, O. Lahav, M. A. G. Maia, J. L. Marshall, F. Menanteau, R. Miquel, A. Palmese, F. Paz-Chinchón, A. A. Plazas, A. K. Romer, E. Sanchez, B. Santiago, V. Scarpine, S. Serrano, M. Smith, M. Soares-Santos, E. Suchyta, G. Tarle, D. Thomas, T. N. Varga, A. R. Walker, and DES Collaboration. Milky Way Satellite Census. II. Galaxy-Halo Connection Constraints Including the Impact of the Large Magellanic Cloud. *ApJ*, 893(1):48, April 2020.

[334] Ethan O. Nadler, Arka Banerjee, Susmita Adhikari, Yao-Yuan Mao, and Risa H. Wechsler. Signatures of Velocity-dependent Dark Matter Self-interactions in Milky Way-mass Halos. *ApJ*, 896(2):112, June 2020.

[335] Ethan O. Nadler, Vera Gluscevic, Kimberly K. Boddy, and Risa H. Wechsler. Constraints on Dark Matter Microphysics from the Milky Way Satellite Population. *ApJL*, 878(2):L32, Jun 2019.

[336] Ethan O. Nadler, Yao-Yuan Mao, Gregory M. Green, and Risa H. Wechsler. Modeling the Connection between Subhalos and Satellites in Milky Way-like Systems. *ApJ*, 873(1):34, Mar 2019.

- [337] Ethan O. Nadler, Yao-Yuan Mao, Risa H. Wechsler, Shea Garrison-Kimmel, and Andrew Wetzel. Modeling the impact of baryons on subhalo populations with machine learning. *ApJ*, 859(2):129, 2018.
- [338] J. F. Navarro, C. S. Frenk, and S. D. M. White. A Universal Density Profile from Hierarchical Clustering. *ApJ*, 490:493–508, December 1997.
- [339] Maria K. Neuzil, Philip Mansfield, and Andrey V. Kravtsov. The Sheet of Giants: Unusual properties of the Milky Way’s immediate neighbourhood. *MNRAS*, 494(2):2600–2617, April 2020.
- [340] Oliver Newton, Marius Cautun, Adrian Jenkins, Carlos S. Frenk, and John C. Helly. The total satellite population of the Milky Way. *MNRAS*, 479:2853–2870, Sep 2018.
- [341] Oliver Newton, Matteo Leo, Marius Cautun, Adrian Jenkins, Carlos S. Frenk, Mark R. Lovell, John C. Helly, and Andrew J. Benson. Constraints on the properties of warm dark matter using the satellite galaxies of the Milky Way. *arXiv e-prints*, page arXiv:2011.08865, November 2020.
- [342] A. M. Nierenberg, D. Gilman, T. Treu, G. Brammer, S. Birrer, L. Moustakas, A. Agnello, T. Anguita, C. D. Fassnacht, V. Motta, A. H. G. Peter, and D. Sluse. Double dark matter vision: twice the number of compact-source lenses with narrow-line lensing and the WFC3 grism. *MNRAS*, 492(4):5314–5335, March 2020.
- [343] A. M. Nierenberg, T. Treu, G. Brammer, A. H. G. Peter, C. D. Fassnacht, C. R. Keeton, C. S. Kochanek, K. B. Schmidt, D. Sluse, and S. A. Wright. Probing dark matter substructure in the gravitational lens HE 0435-1223 with the WFC3 grism. *MNRAS*, 471(2):2224–2236, October 2017.
- [344] A. M. Nierenberg, T. Treu, S. A. Wright, C. D. Fassnacht, and M. W. Auger. Detection of substructure with adaptive optics integral field spectroscopy of the gravitational lens B1422+231. *MNRAS*, 442(3):2434–2445, August 2014.
- [345] Hiroya Nishikawa, Kimberly K. Boddy, and Manoj Kaplinghat. Accelerated core collapse in tidally stripped self-interacting dark matter halos. *Phys. Rev. D*, 101(6):063009, March 2020.
- [346] Kenneth M. Nollett and Gary Steigman. BBN and the CMB constrain neutrino coupled light WIMPs. *Phys. Rev. D*, 91(8):083505, April 2015.
- [347] J. Oñorbe, M. Boylan-Kolchin, J. S. Bullock, P. F. Hopkins, D. Kereš, C.-A. Faucher-Giguère, E. Quataert, and N. Murray. Forged in fire: cusps, cores and baryons in low-mass dwarf galaxies. *MNRAS*, 454:2092–2106, dec 2015.

- [348] Masamune Oguri and Ryuichi Takahashi. Probing Dark Low-mass Halos and Primordial Black Holes with Frequency-dependent Gravitational Lensing Dispersions of Gravitational Waves. *ApJ*, 901(1):58, September 2020.
- [349] T. Okamoto, L. Gao, and T. Theuns. Mass loss of galaxies due to an ultraviolet background. *MNRAS*, 390:920–928, November 2008.
- [350] K. A. Oman, J. F. Navarro, A. Fattah, C. S. Frenk, T. Sawala, S. D. M. White, R. Bower, R. A. Crain, M. Furlong, M. Schaller, J. Schaye, and T. Theuns. The unexpected diversity of dwarf galaxy rotation curves. *MNRAS*, 452:3650–3665, oct 2015.
- [351] Kyle A. Oman. The apostle simulations: Rotation curves derived from synthetic 21-cm observations. In *Proceedings of the International Astronomical Union*, volume 13, page 213–218, Cambridge, United Kingdom, 2017. Cambridge University Press.
- [352] Junpei Ooba, Horoyuki Tashiro, and Kenji Kadota. Cosmological constraints on the velocity-dependent baryon-dark matter coupling. *arXiv e-prints*, page arXiv:1902.00826, February 2019.
- [353] S. A. Pardy, E. D’Onghia, J. F. Navarro, R. Grand, F. A. Gomez, F. Marinacci, R. Pakmor, C. Simpson, and V. Springel. Satellites of satellites: the case for Carina and Fornax. *MNRAS*, 492:1543–1549, January 2020.
- [354] E. Patel, G. Besla, S. T. Sohn, and K. Mandel. Estimating the Mass of the Milky Way Using the Ensemble of Classical Satellite Galaxies. In *American Astronomical Society Meeting Abstracts #232*, volume 232 of *American Astronomical Society Meeting Abstracts*, page 402.03, June 2018.
- [355] Ekta Patel, Gurtina Besla, and Kaisey Mandel. Orbits of massive satellite galaxies - II. Bayesian estimates of the Milky Way and Andromeda masses using high-precision astrometry and cosmological simulations. *MNRAS*, 468(3):3428–3449, Jul 2017.
- [356] Ekta Patel, Nitya Kallivayalil, Nicolas Garavito-Camargo, Gurtina Besla, Daniel R. Weisz, Roeland P. van der Marel, Michael Boylan-Kolchin, Marcel S. Pawlowski, and Facundo A. Gómez. The Orbital Histories of Magellanic Satellites Using Gaia DR2 Proper Motions. *ApJ*, 893(2):121, April 2020.
- [357] Alexandre Payez, Carmelo Evoli, Tobias Fischer, Maurizio Giannotti, Alessandro Mirizzi, and Andreas Ringwald. Revisiting the SN1987A gamma-ray limit on ultralight axion-like particles. *JCAP*, 02(2):006, February 2015.
- [358] J. Peñarrubia, J. F. Navarro, and A. W. McConnachie. The Tidal Evolution of Local Group Dwarf Spheroidals. *ApJ*, 673:226–240, January 2008.

- [359] J. Peñarrubia, J. F. Navarro, A. W. McConnachie, and N. F. Martin. The Signature of Galactic Tides in Local Group Dwarf Spheroidals. *ApJ*, 698:222–232, June 2009.
- [360] Jorge Peñarrubia, Facundo A. Gómez, Gurtina Besla, Denis Erkal, and Yin-Zhe Ma. A timing constraint on the (total) mass of the Large Magellanic Cloud. *MNRAS*, 456(1):L54–L58, Feb 2016.
- [361] R. D. Peccei and Helen R. Quinn. CP Conservation in the Presence of Instantons. *Phys. Rev. Lett.*, 38:1440–1443, 1977.
- [362] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, October 2011.
- [363] Kerstin Perez, Kenny C. Y. Ng, John F. Beacom, Cora Hersh, Shunsaku Horiuchi, and Roman Krivonos. Almost closing the ν MSM sterile neutrino dark matter window with nustar. *Phys. Rev. D*, 95:123002, Jun 2017.
- [364] Annika H. G. Peter, Christopher E. Moody, and Marc Kamionkowski. Dark-matter decays and self-gravitating halos. *Phys. Rev. D*, 81(10):103501, May 2010.
- [365] A. Pillepich, P. Madau, and L. Mayer. Building Late-type Spiral Galaxies by In-situ and Ex-situ Star Formation. *ApJ*, 799:184, February 2015.
- [366] Juan C. B. Pineda, Christopher C. Hayward, Volker Springel, and Claudia Mendes de Oliveira. Rotation curve fitting and its fatal attraction to cores in realistically simulated galaxy observations. *MNRAS*, 466(1):63–87, Apr 2017.
- [367] Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, and et al. Planck 2015 results. XIII. Cosmological parameters. *A&A*, 594:A13, September 2016.
- [368] Emil Polisensky and Massimo Ricotti. Constraints on the dark matter particle mass from the number of Milky Way satellites. *Phys. Rev. D*, 83(4):043506, February 2011.
- [369] A. Pontzen and F. Governato. How supernova feedback turns dark matter cusps into cores. *MNRAS*, 421:3464–3471, apr 2012.
- [370] Lorenzo Posti and Amina Helmi. Mass and shape of the Milky Way’s dark matter halo with globular clusters from Gaia and Hubble. *A&A*, 621:A56, Jan 2019.

- [371] William H. Press and Paul Schechter. Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation. *ApJ*, 187:425–438, February 1974.
- [372] A. Pujol, R. A. Skibba, E. Gaztañaga, A. Benson, J. Blaizot, R. Bower, J. Carretero, F. J. Castander, A. Cattaneo, S. A. Cora, D. J. Croton, W. Cui, D. Cunnamma, G. De Lucia, J. E. Devriendt, P. J. Elahi, A. Font, F. Fontanot, J. Garcia-Bellido, I. D. Gargiulo, V. Gonzalez-Perez, J. Helly, B. M. B. Henriques, M. Hirschmann, A. Knebe, J. Lee, G. A. Mamon, P. Monaco, J. Onions, N. D. Padilla, F. R. Pearce, C. Power, R. S. Somerville, C. Srisawat, P. A. Thomas, E. Tollet, C. A. Vega-Martínez, and S. K. Yi. nifty cosmology: the clustering consistency of galaxy formation models. *MNRAS*, 469:749–762, jul 2017.
- [373] Anthony R. Pullen, Andrew J. Benson, and Leonidas A. Moustakas. Nonlinear Evolution of Dark Matter Subhalos and Applications to Warm Dark Matter. *ApJ*, 792(1):24, September 2014.
- [374] Chris W. Purcell and Andrew R. Zentner. Bailing out the Milky Way: variation in the properties of massive dwarfs among galaxy-sized systems. *JCAP*, 2012(12):007, December 2012.
- [375] Harikrishnan Ramani, Tanner Trickler, and Kathryn M. Zurek. Observability of dark matter substructure with pulsar timing correlations. *JCAP*, 2020(12):033, December 2020.
- [376] J. I. Read and D. Erkal. Abundance matching with the mean star formation rate: there is no missing satellites problem in the Milky Way. *ArXiv e-prints*, July 2018.
- [377] J. I. Read, M. G. Walker, and P. Steger. The case for a cold dark matter cusp in Draco. *MNRAS*, 481(1):860–877, Nov 2018.
- [378] J. I. Read, M. G. Walker, and P. Steger. Dark matter heats up in dwarf galaxies. *MNRAS*, 484(1):1401–1420, Mar 2019.
- [379] R. M. Reddick, R. H. Wechsler, J. L. Tinker, and P. S. Behroozi. The Connection between Galaxies and Dark Matter Structures in the Local Universe. *ApJ*, 771:30, July 2013.
- [380] Tao Ren, Anna Kwa, Manoj Kaplinghat, and Hai-Bo Yu. Reconciling the Diversity and Uniformity of Galactic Rotation Curves with Self-Interacting Dark Matter. *Physical Review X*, 9(3):031020, Jul 2019.
- [381] Martin P. Rey, Andrew Pontzen, Oscar Agertz, Matthew D. A. Orkney, Justin I. Read, Amélie Saintonge, and Christian Pedersen. EDGE: The Origin of Scatter in Ultra-faint Dwarf Stellar Masses and Surface Brightnesses. *ApJl*, 886(1):L3, Nov 2019.

- [382] Christopher S. Reynolds, M. C. David Marsh, Helen R. Russell, Andrew C. Fabian, Robyn Smith, Francesco Tombesi, and Sylvain Veilleux. Astrophysical Limits on Very Light Axion-like Particles from Chandra Grating Spectroscopy of NGC 1275. *ApJ*, 890(1):59, February 2020.
- [383] Jack Richings, Carlos Frenk, Adrian Jenkins, Andrew Robertson, Azadeh Fattahi, Robert J. J. Grand, Julio Navarro, Rüdiger Pakmor, Facundo A. Gomez, Federico Marinacci, and Kyle A. Oman. Subhalo destruction in the APOSTLE and AURIGA simulations. *MNRAS*, 492(4):5780–5793, March 2020.
- [384] Jack Richings, Carlos Frenk, Adrian Jenkins, Andrew Robertson, and Matthieu Schaller. A high-resolution cosmological simulation of a strong gravitational lens. *MNRAS*, 501(3):4657–4668, March 2021.
- [385] Alexander H. Riley, Azadeh Fattahi, Andrew B. Pace, Louis E. Strigari, Carlos S. Frenk, Facundo A. Gómez, Robert J. J. Grand, Federico Marinacci, Julio F. Navarro, Rüdiger Pakmor, Christine M. Simpson, and Simon D. M. White. The velocity anisotropy of the Milky Way satellite system. *MNRAS*, 486(2):2679–2694, Jun 2019.
- [386] A. Ringwald. Ultralight Particle Dark Matter. In *25th Rencontres de Blois on Particle Physics and Cosmology*, 10 2013.
- [387] E. Ritondale, S. Vegetti, G. Despali, M. W. Auger, L. V. E. Koopmans, and J. P. McKean. Low-mass halo perturbations in strong gravitational lenses at redshift $z \sim 0.5$ are consistent with CDM. *MNRAS*, 485(2):2179–2193, May 2019.
- [388] C. P. Robert and D. Wraith. Computational methods for Bayesian model choice. In P. M. Goggans and C.-Y. Chan, editors, *American Institute of Physics Conference Series*, volume 1193 of *American Institute of Physics Conference Series*, pages 251–262, December 2009.
- [389] Andrew Robertson, Richard Massey, and Vincent Eke. Cosmic particle colliders: simulations of self-interacting dark matter with anisotropic scattering. *MNRAS*, 467(4):4719–4730, Jun 2017.
- [390] Victor H. Robles, James S. Bullock, Oliver D. Elbert, Alex Fitts, Alejandro González-Samaniego, Michael Boylan-Kolchin, Philip F. Hopkins, Claude-André Faucher-Giguère, Dušan Kereš, and Christopher C. Hayward. SIDM on FIRE: hydrodynamical self-interacting dark matter simulations of low-mass dwarf galaxies. *MNRAS*, 472(3):2945–2954, Dec 2017.
- [391] Victor H. Robles, Tyler Kelley, James S. Bullock, and Manoj Kaplinghat. The Milky Way’s halo and subhaloes in self-interacting dark matter. *MNRAS*, 490(2):2117–2123, Dec 2019.

- [392] Miguel Rocha, Annika H. G. Peter, James S. Bullock, Manoj Kaplinghat, Shea Garrison-Kimmel, Jose Oñorbe, and Leonidas A. Moustakas. Cosmological simulations with self-interacting dark matter - I. Constant-density cores and substructure. *MNRAS*, 430(1):81–104, Mar 2013.
- [393] A. Rodríguez-Puebla, V. Avila-Reese, and N. Drory. The Massive Satellite Population of Milky-Way-sized Galaxies. *ApJ*, 773:172, August 2013.
- [394] Keir K. Rogers and Hiranya V. Peiris. General framework for cosmological dark matter bounds using N -body simulations. *Phys. Rev. D*, 103(4):043526, February 2021.
- [395] Keir K. Rogers and Hiranya V. Peiris. Strong Bound on Canonical Ultralight Axion Dark Matter from the Lyman-Alpha Forest. *Phys. Rev. Lett.*, 126(7):071302, February 2021.
- [396] Nina Roth, Andrew Pontzen, and Hiranya V. Peiris. Genetically modified haloes: towards controlled experiments in Λ CDM galaxy formation. *MNRAS*, 455(1):974–986, January 2016.
- [397] Vera C. Rubin and Jr. Ford, W.Kent. Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions. *ApJ*, 159:379–403, 1970.
- [398] Mohammadtaher Safarzadeh and David N. Spergel. Ultra-light Dark Matter Is Incompatible with the Milky Way’s Dwarf Satellites. *ApJ*, 893(1):21, April 2020.
- [399] T. Sakamoto and T. Hasegawa. Discovery of a Faint Old Stellar System at 150 kpc. *ApJl*, 653:L29–L32, December 2006.
- [400] Laura V. Sales, Julio F. Navarro, Nitya Kallivayalil, and Carlos S. Frenk. Identifying true satellites of the Magellanic Clouds. *MNRAS*, 465(2):1879–1888, Feb 2017.
- [401] Laura V. Sales, Wenting Wang, Simon D. M. White, and Julio F. Navarro. Satellites and haloes of dwarf galaxies. *MNRAS*, 428(1):573–578, January 2013.
- [402] Omid Sameie, Peter Creasey, Hai-Bo Yu, Laura V. Sales, Mark Vogelsberger, and Jesús Zavala. The impact of baryonic discs on the shapes and profiles of self-interacting dark matter haloes. *MNRAS*, 479(1):359–367, Sep 2018.
- [403] Omid Sameie, Hai-Bo Yu, Laura V. Sales, Mark Vogelsberger, and Jesús Zavala. Self-Interacting Dark Matter Subhalos in the Milky Way’s Tides. *Phys. Rev. Lett.*, 124(14):141102, April 2020.
- [404] Jenna Samuel, Andrew Wetzel, Erik Tollerud, Shea Garrison-Kimmel, Sarah Loebman, Karreem El-Badry, Philip F. Hopkins, Michael Boylan-Kolchin, Claude-André Faucher-Giguère, James S. Bullock, Samantha Benincasa, and Jeremy Bailin. A profile in FIRE: resolving

the radial distributions of satellite galaxies in the Local Group with simulations. *MNRAS*, 491(1):1471–1490, Jan 2020.

[405] Isabel M. Santos-Santos, Arianna Di Cintio, Chris B. Brook, Andrea Macciò, Aaron Dutton, and Rosa Domínguez-Tenreiro. NIHAO - XIV. Reproducing the observed diversity of dwarf galaxy rotation curve shapes in Λ CDM. *MNRAS*, 473(4):4392–4403, Feb 2018.

[406] Isabel M. E. Santos-Santos, Julio F. Navarro, Andrew Robertson, Alejandro Benítez-Llambay, Kyle A. Oman, Mark R. Lovell, Carlos S. Frenk, Aaron D. Ludlow, Azadeh Fattahi, and Adam Ritz. Baryonic clues to the puzzling diversity of dwarf galaxy rotation curves. *MNRAS*, 495(1):58–77, April 2020.

[407] Abir Sarkar, Subinoy Das, and Shiv K. Sethi. How late can the dark matter form in our universe? *JCAP*, 2015(3):004, March 2015.

[408] T. Sawala, C. S. Frenk, A. Fattahi, J. F. Navarro, R. G. Bower, R. A. Crain, C. Dalla Vecchia, M. Furlong, J. C. Helly, A. Jenkins, K. A. Oman, M. Schaller, J. Schaye, T. Theuns, J. Trayford, and S. D. M. White. The APOSTLE simulations: solutions to the Local Group’s cosmic puzzles. *MNRAS*, 457:1931–1943, April 2016.

[409] T. Sawala, C. S. Frenk, A. Fattahi, J. F. Navarro, R. G. Bower, R. A. Crain, C. Dalla Vecchia, M. Furlong, A. Jenkins, I. G. McCarthy, Y. Qu, M. Schaller, J. Schaye, and T. Theuns. Bent by baryons: the low-mass galaxy-halo relation. *MNRAS*, 448:2941–2947, April 2015.

[410] T. Sawala, C. S. Frenk, A. Fattahi, J. F. Navarro, T. Theuns, R. G. Bower, R. A. Crain, M. Furlong, A. Jenkins, M. Schaller, and J. Schaye. The chosen few: the low-mass haloes that host faint galaxies. *MNRAS*, 456:85–97, February 2016.

[411] T. Sawala, P. Pihajoki, P. H. Johansson, C. S. Frenk, J. F. Navarro, K. A. Oman, and S. D. M. White. Shaken and stirred: the milky way’s dark substructures. *MNRAS*, 467:4383–4400, jun 2017.

[412] J. A. Schewtschenko, C. M. Baugh, R. J. Wilkinson, C. Boehm, S. Pascoli, and T. Sawala. Dark matter-radiation interactions: the structure of Milky Way satellite galaxies. *Mon. Not. R. Astron. Soc.*, 461(3):2282–2287, September 2016.

[413] Hsi-Yu Schive, Tzihong Chiueh, Tom Broadhurst, and Kuan-Wei Huang. Contrasting Galaxy Formation from Quantum Wave Dark Matter, ψ DM, with Λ CDM, using Planck and Hubble Data. *ApJ*, 818(1):89, February 2016.

[414] Aurel Schneider. Astrophysical constraints on resonantly produced sterile neutrino dark matter. *JCAP*, 04(4):059, April 2016.

- [415] Aurel Schneider, Robert E. Smith, Andrea V. Macciò, and Ben Moore. Non-linear evolution of cosmological structures in warm dark matter models. *MNRAS*, 424:684–698, July 2012.
- [416] Aurel Schneider, Robert E. Smith, and Darren Reed. Halo mass function and the free streaming scale. *MNRAS*, 433(2):1573–1587, August 2013.
- [417] Nils Schöneberg, Julien Lesgourges, and Deanna C. Hooper. The BAO+BBN take on the Hubble tension. *JCAP*, 10:029, 2019.
- [418] Katelin Schutz. Subhalo mass function and ultralight bosonic dark matter. *Phys. Rev. D*, 101(12):123026, June 2020.
- [419] Shi Shao, Marius Cautun, Alis J. Deason, Carlos S. Frenk, and Tom Theuns. Evolution of LMC/M33-mass dwarf galaxies in the EAGLE simulation. *MNRAS*, 479(1):284–296, Sep 2018.
- [420] Harlow Shapley. Galactic and extragalactic studies. *Proceedings of the National Academy of Sciences*, 25(11):565–569, 1939.
- [421] Ravi K. Sheth, H. J. Mo, and Giuseppe Tormen. Ellipsoidal collapse and an improved model for the number and spatial distribution of dark matter haloes. *MNRAS*, 323(1):1–12, May 2001.
- [422] Xiangdong Shi and George M. Fuller. New dark matter candidate: Nonthermal sterile neutrinos. *Phys. Rev. Lett.*, 82:2832–2835, Apr 1999.
- [423] Joshua Simon, Keith Bechtol, Alex Drlica-Wagner, Marla Geha, Vera Gluscevic, Alex Ji, Evan Kirby, Ting S. Li, Ethan O. Nadler, Andrew B. Pace, Annika Peter, and Risa Wechsler. Dynamical masses for a complete census of local dwarf galaxies. *BAAS*, 51(3), 5 2019.
- [424] Joshua D. Simon. Gaia Proper Motions and Orbits of the Ultra-faint Milky Way Satellites. *ApJ*, 863:89, Aug 2018.
- [425] Joshua D. Simon. The Faintest Dwarf Galaxies. *Annu. Rev. Astron. Astrophys.*, 57:375–415, August 2019.
- [426] Joshua D. Simon, Marla Geha, Quinn E. Minor, Gregory D. Martinez, Evan N. Kirby, James S. Bullock, Manoj Kaplinghat, Louis E. Strigari, Beth Willman, Philip I. Choi, Erik J. Tollerud, and Joe Wolf. A Complete Spectroscopic Survey of the Milky Way Satellite Segue 1: The Darkest Galaxy. *ApJ*, 733:46, May 2011.
- [427] Christine M. Simpson, Robert J. J. Grand, Facundo A. Gómez, Federico Marinacci, Rüdiger Pakmor, Volker Springel, David J. R. Campbell, and Carlos S. Frenk. Quenching and ram pressure stripping of simulated Milky Way satellite galaxies. *MNRAS*, 478(1):548–567, Jul 2018.

- [428] S. W. Skillman, M. S. Warren, M. J. Turk, R. H. Wechsler, D. E. Holz, and P. M. Sutter. Dark Sky Simulations: Early Data Release. *ArXiv e-prints*, July 2014.
- [429] R. S. Somerville. Can photoionization squelching resolve the substructure crisis? *ApJl*, 572:L23–L26, jun 2002.
- [430] Alessandro Sonnenfeld, Tommaso Treu, Philip J. Marshall, Sherry H. Suyu, Raphaël Gavazzi, Matthew W. Auger, and Carlo Nipoti. The SL2S Galaxy-scale Lens Sample. V. Dark Matter Halos and Stellar IMF of Massive Early-type Galaxies Out to Redshift 0.8. *ApJ*, 800(2):94, February 2015.
- [431] D. Spergel, N. Gehrels, C. Baltay, D. Bennett, J. Breckinridge, M. Donahue, A. Dressler, B. S. Gaudi, T. Greene, O. Guyon, C. Hirata, J. Kalirai, N. J. Kasdin, B. Macintosh, W. Moos, S. Perlmutter, M. Postman, B. Rauscher, J. Rhodes, Y. Wang, D. Weinberg, D. Benford, M. Hudson, W. S. Jeong, Y. Mellier, W. Traub, T. Yamada, P. Capak, J. Colbert, D. Masters, M. Penny, D. Savransky, D. Stern, N. Zimmerman, R. Barry, L. Bartusek, K. Carpenter, E. Cheng, D. Content, F. Dekens, R. Demers, K. Grady, C. Jackson, G. Kuan, J. Kruk, M. Melton, B. Nemati, B. Parvin, I. Poberezhskiy, C. Peddie, J. Ruffa, J. K. Wallace, A. Whipple, E. Wollack, and F. Zhao. Wide-Field InfrarRed Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report. *arXiv e-prints*, page arXiv:1503.03757, March 2015.
- [432] David N. Spergel and Paul J. Steinhardt. Observational Evidence for Self-Interacting Cold Dark Matter. *Phys. Rev. Lett.*, 84(17):3760–3763, Apr 2000.
- [433] V. Springel, J. Wang, M. Vogelsberger, A. Ludlow, A. Jenkins, A. Helmi, J. F. Navarro, C. S. Frenk, and S. D. M. White. The aquarius project: the subhaloes of galactic haloes. *MNRAS*, 391:1685–1711, dec 2008.
- [434] Chaichalit Srisawat, Alexander Knebe, Frazer R. Pearce, Aurel Schneider, Peter A. Thomas, Peter Behroozi, Klaus Dolag, Pascal J. Elahi, Jiaxin Han, John Helly, Yipeng Jing, Intae Jung, Jaehyun Lee, Yao-Yuan Mao, Julian Onions, Vicente Rodriguez-Gomez, Dylan Tweed, and Sukyoung K. Yi. Sussing Merger Trees: The Merger Trees Comparison Project. *MNRAS*, 436(1):150–162, Nov 2013.
- [435] Sam G. Stafford, Shaun T. Brown, Ian G. McCarthy, Andreea S. Font, Andrew Robertson, and Robert Poole-McKenzie. Exploring extensions to the standard cosmological model and the impact of baryons on small scales. *MNRAS*, 497(3):3809–3829, August 2020.
- [436] E. Starkenburg, A. Helmi, G. De Lucia, Y.-S. Li, J. F. Navarro, A. S. Font, C. S. Frenk, V. Springel, C. A. Vera-Ciro, and S. D. M. White. The satellites of the Milky Way - insights from semi-analytic modelling in a Λ CDM cosmology. *MNRAS*, 429:725–743, February 2013.

[437] Stephen Stopyra, Andrew Pontzen, Hiranya Peiris, Nina Roth, and Martin P. Rey. GenetIC—A New Initial Conditions Generator to Support Genetically Modified Zoom Simulations. *ApJS*, 252(2):28, February 2021.

[438] The MSE Science Team, Carine Babusiaux, Maria Bergemann, Adam Burgasser, Sara Ellison, Daryl Haggard, Daniel Huber, Manoj Kaplinghat, Ting Li, Jennifer Marshall, Sarah Martell, Alan McConnachie, Will Percival, Aaron Robotham, Yue Shen, Sivarani Thirupathi, Kim-Vy Tran, Christophe Yèche, David Yong, Vardan Adibekyan, Victor Silva Aguirre, George Angelou, Martin Asplund, Michael Balogh, Projwal Banerjee, Michele Bannister, Daniela Barría, Giuseppina Battaglia, Amelia Bayo, Keith Bechtol, Paul G. Beck, Timothy C. Beers, Earl P. Bellinger, Trystyn Berg, Joachim M. Bestenlehner, Maciej Bilicki, Bertram Bitsch, Joss Bland-Hawthorn, Adam S. Bolton, Alessandro Boselli, Jo Bovy, Angela Bragaglia, Derek Buzasi, Elisabetta Caffau, Jan Cami, Timothy Carleton, Luca Casagrande, Santi Cassisi, Márcio Catelan, Chihway Chang, Luca Cortese, Ivana Damjanov, Luke J. M. Davies, Richard de Grijs, Gisella de Rosa, Alis Deason, Paola di Matteo, Alex Drlica-Wagner, Denis Erkal, Ana Escorza, Laura Ferrarese, Scott W. Fleming, Andreu Font-Ribera, Ken Freeman, Boris T. Gänsicke, Maksim Gabdeev, Sarah Gallagher, Davide Gandolfi, Rafael A. García, Patrick Gaulme, Marla Geha, Mario Gennaro, Mark Gieles, Karoline Gilbert, Yjan Gordon, Aruna Goswami, Johnny P. Greco, Carl Grillmair, Guillaume Guiglion, Vincent Hénault-Brunet, Patrick Hall, Gerald Handler, Terese Hansen, Nimish Hathi, Despina Hatzidimitriou, Misha Haywood, Juan V. Hernández Santisteban, Lynne Hillenbrand, Andrew M. Hopkins, Culian Howlett, Michael J. Hudson, Rodrigo Ibata, Dragana Ilić, Pascale Jablonka, Alexander Ji, Linhua Jiang, Stephanie Juneau, Amanda Karakas, Drisya Karinkuzhi, Stacy Y. Kim, Xu Kong, Iraklis Konstantopoulos, Jens-Kristian Krogager, Claudia Lagos, Rosine Lallement, Chervin Laporte, Yveline Lebreton, Khee-Gan Lee, Geraint F. Lewis, Sophia Lianou, Xin Liu, Nicolas Lodieu, Jon Loveday, Szabolcs Mészáros, Martin Makler, Yao-Yuan Mao, Danilo Marchesini, Nicolas Martin, Mario Mateo, Carl Melis, Thibault Merle, Andrea Miglio, Faizan Gohar Mohammad, Karan Molaverdikhani, Richard Monier, Thierry Morel, Benoit Mosser, David Nataf, Lina Necib, Hilding R. Neilson, Jeffrey A. Newman, A. M. Nierenberg, Brian Nord, Pasquier Noterdaeme, Chris O’Dea, Mahmoudreza Oshagh, Andrew B. Pace, Nathalie Palanque-Delabrouille, Gajendra Pandey, Laura C. Parker, Marcel S. Pawłowski, Annika H. G. Peter, Patrick Petitjean, Andreea Petric, Vinicius Placco, Luka Č. Popović, Adrian M. Price-Whelan, Andrej Prsa, Swara Ravindranath, R. Michael Rich, John Ruan, Jan Rybicki, Charli Sakari, Robyn E. Sanderson, Ricardo Schiavon, Carlo Schimd, Aldo Serenelli, Arnaud Siebert, Małgorzata Siudek, Rodolfo Smiljanic, Daniel Smith, Jennifer Sobeck, Else Starkenburg, Dennis Stello, Gyula M. Szabó, Robert Szabo, Matthew A. Taylor, Karun Thanjavur, Guillaume Thomas, Erik Tollerud, Silvia Toonen, Pier-Emmanuel Tremblay, Laurence Tresse, Maria Tsantaki, Marica Valentini, Sophie Van Eck, Andrei Variu, Kim Venn, Eva Villaver,

Matthew G. Walker, Yiping Wang, Yuting Wang, Michael J. Wilson, Nicolas Wright, Siyi Xu, Mutlu Yildiz, Huawei Zhang, Konstanze Zwintz, Borja Anguiano, Megan Bedell, William Chaplin, Remo Collet, Jean-Charles Cuillandre, Pierre-Alain Duc, Nicolas Flagey, JJ Hermes, Alexis Hill, Devika Kamath, Mary Beth Laychak, Katarzyna Małek, Mark Marley, Andy Sheinis, Doug Simons, Sérgio G. Sousa, Kei Szeto, Yuan-Sen Ting, Simona Vegetti, Lisa Wells, Ferdinand Babas, Steve Bauman, Alessandro Bosselli, Pat Côté, Matthew Colless, Johan Comparat, Helene Courtois, David Crampton, Scott Croom, Luke Davies, Richard de Grijs, Kelly Denny, Daniel Devost, Paola di Matteo, Simon Driver, Mirian Fernandez-Lorenzo, Raja Guhathakurta, Zhanwen Han, Clare Higgs, Vanessa Hill, Kevin Ho, Andrew Hopkins, Mike Hudson, Rodrigo Ibata, Sidik Isani, Matt Jarvis, Andrew Johnson, Eric Jullo, Nick Kaiser, Jean-Paul Kneib, Jun Koda, George Koshy, Shan Mignot, Rick Murowinski, Jeff Newman, Adi Nusser, Anna Pancoast, Eric Peng, Celine Peroux, Christophe Pichon, Bianca Poggianti, Johan Richard, Derrick Salmon, Arnaud Seibert, Prajval Shastri, Dan Smith, Firoza Sutaria, Charling Tao, Edwar Taylor, Brent Tully, Ludovic van Waerbeke, Tom Vermeulen, Matthew Walker, Jon Willis, Chris Willot, and Kanoa Withington. The Detailed Science Case for the Maunakea Spectroscopic Explorer, 2019 edition. *arXiv e-prints*, page arXiv:1904.04907, April 2019.

[439] A. A. Thoul and D. H. Weinberg. Hydrodynamic Simulations of Galaxy Formation. II. Photoionization and the Formation of Low-Mass Galaxies. *ApJ*, 465:608, July 1996.

[440] Erik J. Tollerud, James S. Bullock, Louis E. Strigari, and Beth Willman. Hundreds of Milky Way Satellites? Luminosity Bias in the Satellite Luminosity Function. *ApJ*, 688(1):277–289, Nov 2008.

[441] Erik J. Tollerud and J. E. G. Peek. Where Are All of the Gas-bearing Local Dwarf Galaxies? Quantifying Possible Impacts of Reionization. *ApJ*, 857(1):45, Apr 2018.

[442] E. Tollet, A. V. Macciò, A. A. Dutton, G. S. Stinson, L. Wang, C. Penzo, T. A. Gutcke, T. Buck, X. Kang, C. Brook, A. Di Cintio, B. W. Keller, and J. Wadsley. Nihao - iv: core creation and destruction in dark matter density profiles across cosmic time. *MNRAS*, 456:3542–3552, mar 2016.

[443] G. Torrealba, V. Belokurov, S. E. Koposov, K. Bechtol, A. Drlica-Wagner, K. A. G. Olsen, A. K. Vivas, B. Yanny, P. Jethwa, and A. R. Walker. Discovery of two neighbouring satellites in the Carina constellation with MagLiteS. *MNRAS*, 475(4):5085–5097, Apr 2018.

[444] G. Torrealba, V. Belokurov, S. E. Koposov, T. S. Li, M. G. Walker, J. L. Sanders, A. Geringer-Sameth, D. B. Zucker, K. Kuehn, N. W. Evans, and W. Dehnen. The hidden giant: discovery of an enormous Galactic dwarf satellite in Gaia DR2. *MNRAS*, 488(2):2743–2766, September 2019.

- [445] G. Torrealba, S. E. Koposov, V. Belokurov, and M. Irwin. The feeble giant. Discovery of a large and diffuse Milky Way dwarf galaxy in the constellation of Crater. *MNRAS*, 459:2370–2378, July 2016.
- [446] G. Torrealba, S. E. Koposov, V. Belokurov, M. Irwin, M. Collins, M. Spencer, R. Ibata, M. Mateo, A. Bonaca, and P. Jethwa. At the survey limits: discovery of the Aquarius 2 dwarf galaxy in the VST ATLAS and the SDSS data. *MNRAS*, 463(1):712–722, Nov 2016.
- [447] Scott Tremaine and James E. Gunn. Dynamical role of light neutral leptons in cosmology. *Phys. Rev. Lett.*, 42(6):407–410, February 1979.
- [448] Sean Tulin and Hai-Bo Yu. Dark matter self-interactions and small scale structure. *Phys. Rep.*, 730:1–57, Feb 2018.
- [449] Sean Tulin, Hai-Bo Yu, and Kathryn M. Zurek. Beyond collisionless dark matter: Particle physics dynamics for dark matter halo structure. *Phys. Rev. D*, 87(11):115007, Jun 2013.
- [450] Mauro Valli and Hai-Bo Yu. Dark matter self-interactions from the internal dynamics of dwarf spheroidals. *NatAs*, 2:907–912, Aug 2018.
- [451] F. C. van den Bosch. Dissecting the evolution of dark matter subhaloes in the bolshoi simulation. *MNRAS*, 468:885–909, jun 2017.
- [452] F. C. van den Bosch, F. Jiang, D. Campbell, and P. Behroozi. On the segregation of dark matter substructure. *MNRAS*, 455:158–177, jan 2016.
- [453] F. C. van den Bosch and G. Ogiya. Dark matter substructure in numerical simulations: a tale of discreteness noise, runaway instabilities, and artificial disruption. *MNRAS*, 475:4066–4087, April 2018.
- [454] F. C. van den Bosch, G. Ogiya, O. Hahn, and A. Burkert. Disruption of dark matter substructure: fact or fiction? *MNRAS*, 474:3043–3066, March 2018.
- [455] S. Vegetti, G. Despali, M. R. Lovell, and W. Enzi. Constraining sterile neutrino cosmologies with strong gravitational lensing observations at redshift $z \sim 0.2$. *MNRAS*, 481(3):3661–3669, December 2018.
- [456] S. Vegetti, D. J. Lagattuta, J. P. McKean, M. W. Auger, C. D. Fassnacht, and L. V. E. Koopmans. Gravitational detection of a low-mass dark satellite galaxy at cosmological distance. *Nature*, 481:341–343, Jan 2012.
- [457] R. Verbeke, E. Papastergis, A. A. Ponomareva, S. Rathi, and S. De Rijcke. A new astrophysical solution to the Too Big To Fail problem. Insights from the moria simulations. *A&A*, 607:A13, October 2017.

- [458] Licia Verde, Tommaso Treu, and Adam G. Riess. Tensions between the early and late Universe. *Nature Astronomy*, 3:891–895, September 2019.
- [459] Matteo Viel, George D. Becker, James S. Bolton, and Martin G. Haehnelt. Warm dark matter as a solution to the small scale crisis: New constraints from high redshift lyman- α forest data. *Phys. Rev. D*, 88:043502, Aug 2013.
- [460] Matteo Viel, Julien Lesgourgues, Martin G. Haehnelt, Sabino Matarrese, and Antonio Riotto. Constraining warm dark matter candidates including sterile neutrinos and light gravitinos with WMAP and the Lyman- α forest. *Phys. Rev. D*, 71(6):063534, March 2005.
- [461] Mark Vogelsberger, Jesus Zavala, and Abraham Loeb. Subhaloes in self-interacting galactic dark matter haloes. *MNRAS*, 423(4):3740–3752, Jul 2012.
- [462] Mark Vogelsberger, Jesús Zavala, Katelin Schutz, and Tracy R. Slatyer. Evaporating the Milky Way halo and its satellites with inelastic self-interacting dark matter. *MNRAS*, 484(4):5437–5452, Apr 2019.
- [463] Digvijay Wadekar and Glennys R. Farrar. First direct astrophysical constraints on dark matter interactions with ordinary matter at very low velocities. *arXiv e-prints*, page arXiv:1903.12190, Mar 2019.
- [464] S. M. Walsh, H. Jerjen, and B. Willman. A Pair of Boötes: A New Milky Way Satellite. *ApJl*, 662:L83–L86, June 2007.
- [465] S. M. Walsh, B. Willman, and H. Jerjen. The Invisibles: A Detection Algorithm to Trace the Faintest Milky Way Satellites. *AJ*, 137:450–469, January 2009.
- [466] J. Wang, S. Bose, C. S. Frenk, L. Gao, A. Jenkins, V. Springel, and S. D. M. White. Universal structure of dark matter haloes over a mass range of 20 orders of magnitude. *Nature*, 585(7823):39–42, September 2020.
- [467] Jie Wang and Simon D. M. White. Discreteness effects in simulations of hot/warm dark matter. *Mon. Not. R. Astron. Soc.*, 380(1):93–103, September 2007.
- [468] L. Wang, C. Li, G. Kauffmann, and G. De Lucia. Modelling galaxy clustering in a high-resolution simulation of structure formation. *MNRAS*, 371:537–547, September 2006.
- [469] Mei-Yu Wang, Annika H. G. Peter, Louis E. Strigari, Andrew R. Zentner, Bryan Arant, Shea Garrison-Kimmel, and Miguel Rocha. Cosmological simulations of decaying dark matter: implications for small-scale structure of dark matter haloes. *Mon. Not. R. Astron. Soc.*, 445(1):614–629, November 2014.

- [470] Asher Wasserman, Pieter van Dokkum, Aaron J. Romanowsky, Jean Brodie, Shany Danieli, Duncan A. Forbes, Roberto Abraham, Christopher Martin, Matt Matuszewski, Alexa Vil-laume, et al. Spatially Resolved Stellar Kinematics of the Ultra-diffuse Galaxy Dragonfly 44. II. Constraints on Fuzzy Dark Matter. *ApJ*, 885(2):155, November 2019.
- [471] Laura L. Watkins, Roeland P. van der Marel, Sangmo Tony Sohn, and N. Wyn Evans. Evidence for an Intermediate-mass Milky Way from Gaia DR2 Halo Globular Cluster Motions. *ApJ*, 873(2):118, Mar 2019.
- [472] Jeremy J. Webb and Jo Bovy. High-resolution simulations of dark matter subhalo disruption in a Milky-Way-like tidal field. *MNRAS*, 499(1):116–128, September 2020.
- [473] Risa H. Wechsler, James S. Bullock, Joel R. Primack, Andrey V. Kravtsov, and Avishai Dekel. Concentrations of dark halos from their assembly histories. *ApJ*, 568(1):52–70, mar 2002.
- [474] Risa H. Wechsler and Jeremy L. Tinker. The Connection Between Galaxies and Their Dark Matter Halos. *Annual Review of Astronomy and Astrophysics*, 56:435–487, Sep 2018.
- [475] Daniel R. Weisz and Michael Boylan-Kolchin. Local Group ultra-faint dwarf galaxies in the reionization era. *MNRAS*, 469(1):L83–L88, Jul 2017.
- [476] Daniel R. Weisz, Nicolas F. Martin, Andrew E. Dolphin, Saundra M. Albers, Michelle L. M. Collins, Annette M. N. Ferguson, Geraint F. Lewis, Dougal Mackey, Alan McConnachie, R. Michael Rich, and Evan D. Skillman. Comparing the Quenching Times of Faint M31 and Milky Way Satellite Galaxies. *ApJL*, 885(1):L8, November 2019.
- [477] Daniel R. Weisz, Evan D. Skillman, Sebastian L. Hidalgo, Matteo Monelli, Andrew E. Dolphin, Alan McConnachie, Edouard J. Bernard, Carme Gallart, Antonio Aparicio, Michael Boylan-Kolchin, Santi Cassisi, Andrew A. Cole, Henry C. Ferguson, Mike Irwin, Nicolas F. Martin, Lucio Mayer, Kristen B. W. McQuinn, Julio F. Navarro, and Peter B. Stetson. Comparing M31 and Milky Way Satellites: The Extended Star Formation Histories of Andromeda II and Andromeda XVI. *ApJ*, 789(1):24, July 2014.
- [478] A. R. Wetzel, P. F. Hopkins, J.-h. Kim, C.-A. Faucher-Giguère, D. Kereš, and E. Quataert. Reconciling dwarf galaxies with Λ CDM cosmology: Simulating a realistic population of satellites around a milky way-mass galaxy. *ApJL*, 827:L23, aug 2016.
- [479] A. R. Wetzel, J. L. Tinker, C. Conroy, and F. C. van den Bosch. Galaxy evolution near groups and clusters: ejected satellites and the spatial extent of environmental quenching. *MNRAS*, 439:2687–2700, April 2014.
- [480] Andrew R. Wetzel, Alis J. Deason, and Shea Garrison-Kimmel. Satellite Dwarf Galaxies in a Hierarchical Universe: Infall Histories, Group Preprocessing, and Reionization. *ApJ*, 807(1):49, July 2015.

- [481] C. Wheeler, J. Oñorbe, J. S. Bullock, M. Boylan-Kolchin, O. D. Elbert, S. Garrison-Kimmel, P. F. Hopkins, and D. Kereš. Sweating the small stuff: simulating dwarf galaxies, ultra-faint dwarf galaxies, and their own tiny satellites. *MNRAS*, 453:1305–1316, October 2015.
- [482] Coral Wheeler, Philip F. Hopkins, Andrew B. Pace, Shea Garrison-Kimmel, Michael Boylan-Kolchin, Andrew Wetzel, James S. Bullock, Dušan Kereš, Claude-André Faucher-Giguère, and Eliot Quataert. Be it therefore resolved: cosmological simulations of dwarf galaxies with 30 solar mass resolution. *MNRAS*, 490(3):4447–4463, December 2019.
- [483] Martin White and Rupert A. C. Croft. Suppressing Linear Power on Dwarf Galaxy Halo Scales. *ApJ*, 539(2):497–504, August 2000.
- [484] B. Willman. In Pursuit of the Least Luminous Galaxies. *Advances in Astronomy*, 2010:285454, 2010.
- [485] B. Willman, M. R. Blanton, A. A. West, J. J. Dalcanton, D. W. Hogg, D. P. Schneider, N. Wherry, B. Yanny, and J. Brinkmann. A New Milky Way Companion: Unusual Globular Cluster or Extreme Dwarf Satellite? *AJ*, 129:2692–2700, June 2005.
- [486] B. Willman, J. J. Dalcanton, D. Martinez-Delgado, A. A. West, M. R. Blanton, D. W. Hogg, J. C. Barentine, H. J. Brewington, M. Harvanek, S. J. Kleinman, J. Krzesinski, D. Long, E. H. Neilsen, Jr., A. Nitta, and S. A. Snedden. A New Milky Way Dwarf Galaxy in Ursa Major. *ApJl*, 626:L85–L88, June 2005.
- [487] J. Wolf, G. D. Martinez, J. S. Bullock, M. Kaplinghat, M. Geha, R. R. Muñoz, J. D. Simon, and F. F. Avedo. Accurate masses for dispersion-supported galaxies. *MNRAS*, 406:1220–1237, August 2010.
- [488] A. H. Wright, A. S. G. Robotham, S. P. Driver, M. Alpaslan, S. K. Andrews, I. K. Baldry, J. Bland -Hawthorn, S. Brough, M. J. I. Brown, M. Colless, E. da Cunha, L. J. M. Davies, Alister W. Graham, B. W. Holwerda, A. M. Hopkins, P. R. Kafle, L. S. Kelvin, J. Loveday, S. J. Maddox, M. J. Meyer, A. J. Moffett, P. Norberg, S. Phillipps, K. Rowlands, E. N. Taylor, L. Wang, and S. M. Wilkins. Galaxy And Mass Assembly (GAMA): the galaxy stellar mass function to $z = 0.1$ from the r-band selected equatorial regions. *MNRAS*, 470(1):283–302, Sep 2017.
- [489] Weishuang Linda Xu, Cora Dvorkin, and Andrew Chael. Probing sub-gev dark matter-baryon scattering with cosmological observables. *Phys. Rev. D*, 97:103530, May 2018.
- [490] Shengqi Yang, Xiaolong Du, Andrew J. Benson, Anthony R. Pullen, and Annika H. G. Peter. A new calibration method of sub-halo orbital evolution for semi-analytic models. *MNRAS*, 498(3):3902–3913, September 2020.

- [491] X. Yang, H. J. Mo, F. C. van den Bosch, Y. Zhang, and J. Han. Evolution of the Galaxy-Dark Matter Connection and the Assembly of Galaxies in Dark Matter Halos. *ApJ*, 752:41, June 2012.
- [492] B. Yniguez, S. Garrison-Kimmel, M. Boylan-Kolchin, and J. S. Bullock. On the stark difference in satellite distributions around the Milky Way and Andromeda. *MNRAS*, 439:73–82, March 2014.
- [493] J. Zavala, M. Vogelsberger, and M. G. Walker. Constraining self-interacting dark matter with the Milky way’s dwarf spheroidals. *MNRAS*, 431:L20–L24, Apr 2013.
- [494] Jesús Zavala and Carlos S. Frenk. Dark Matter Haloes and Subhaloes. *Galaxies*, 7(4):81, September 2019.
- [495] Jesús Zavala, Mark R. Lovell, Mark Vogelsberger, and Jan D. Burger. Diverse dark matter density at sub-kiloparsec scales in Milky Way satellites: Implications for the nature of dark matter. *Phys. Rev. D*, 100(6):063007, Sep 2019.
- [496] A. R. Zentner, A. A. Berlind, J. S. Bullock, A. V. Kravtsov, and R. H. Wechsler. The Physics of Galaxy Clustering. I. A Model for Subhalo Populations. *ApJ*, 624:505–525, May 2005.
- [497] A. R. Zentner and J. S. Bullock. Halo Substructure and the Power Spectrum. *ApJ*, 598:49–72, November 2003.
- [498] Dali Zhang, Yu Luo, and Xi Kang. The effect of the Large Magellanic Cloud on the satellite galaxy population in Milky Way analogous galaxies. *MNRAS*, 486(2):2440–2448, Jun 2019.
- [499] Guangtun Zhu, Zheng Zheng, W. P. Lin, Y. P. Jing, Xi Kang, and Liang Gao. The Dependence of the Occupation of Galaxies on the Halo Formation Time. *ApJL*, 639(1):L5–L8, March 2006.
- [500] Q. Zhu, L. Hernquist, F. Marinacci, V. Springel, and Y. Li. Baryonic impact on the dark matter orbital properties of milky way-sized haloes. *MNRAS*, 466:3876–3886, apr 2017.
- [501] Q. Zhu, F. Marinacci, M. Maji, Y. Li, V. Springel, and L. Hernquist. Baryonic impact on the dark matter distribution in milky way-sized galaxies and their satellites. *MNRAS*, 458:1559–1580, may 2016.
- [502] A. Zolotov, A. M. Brooks, B. Willman, F. Governato, A. Pontzen, C. Christensen, A. Dekel, T. Quinn, S. Shen, and J. Wadsley. Baryons matter: Why luminous satellite galaxies have reduced central masses. *ApJ*, 761:71, dec 2012.
- [503] D. B. Zucker, V. Belokurov, N. W. Evans, J. T. Kleyna, M. J. Irwin, M. I. Wilkinson, M. Felli-hauer, D. M. Bramich, G. Gilmore, H. J. Newberg, B. Yanny, J. A. Smith, P. C. Hewett, E. F. Bell, H.-W. Rix, O. Y. Gnedin, S. Vidrih, R. F. G. Wyse, B. Willman, E. K. Grebel,

D. P. Schneider, T. C. Beers, A. Y. Kniazev, J. C. Barentine, H. Brewington, J. Brinkmann, M. Harvanek, S. J. Kleinman, J. Krzesinski, D. Long, A. Nitta, and S. A. Snedden. A Curious Milky Way Satellite in Ursa Major. *ApJl*, 650:L41–L44, October 2006.

[504] F. Zwicky. Die Rotverschiebung von extragalaktischen Nebeln. *Helv. Phys. Acta*, 6:110–127, 1933.