

Forecasts for Galaxy Formation and Dark Matter Constraints from Dwarf Galaxy Surveys

Ethan O. Nadler^{1,2} , Vera Gluscevic² , Trey Driskell² , Risa H. Wechsler^{3,4,5} , Leonidas A. Moustakas⁶ ,

Andrew Benson¹ , and Yao-Yuan Mao⁷ 

¹ Carnegie Observatories, 813 Santa Barbara Street, Pasadena, CA 91101, USA; enadler@carnegiescience.edu

² Department of Physics & Astronomy, University of Southern California, Los Angeles, CA 90007, USA

³ Kavli Institute for Particle Astrophysics & Cosmology, P. O. Box 2450, Stanford University, Stanford, CA 94305, USA

⁴ Department of Physics, Stanford University, 382 Via Pueblo Mall, Stanford, CA 94305, USA

⁵ SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

⁶ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA

⁷ Department of Physics and Astronomy, University of Utah, Salt Lake City, UT 84112, USA

Received 2024 January 18; revised 2024 March 15; accepted 2024 April 5; published 2024 May 17

Abstract

The abundance of faint dwarf galaxies is determined by the underlying population of low-mass dark matter (DM) halos and the efficiency of galaxy formation in these systems. Here, we quantify potential galaxy formation and DM constraints from future dwarf satellite galaxy surveys. We generate satellite populations using a suite of Milky Way (MW)-mass cosmological zoom-in simulations and an empirical galaxy–halo connection model, and assess sensitivity to galaxy formation and DM signals when marginalizing over galaxy–halo connection uncertainties. We find that a survey of all satellites around one MW-mass host can constrain a galaxy formation cutoff at peak virial masses of $\mathcal{M}_{50} = 10^8 M_\odot$ at the 1σ level; however, a tail toward low \mathcal{M}_{50} prevents a 2σ measurement. In this scenario, combining hosts with differing bright satellite abundances significantly reduces uncertainties on \mathcal{M}_{50} at the 1σ level, but the 2σ tail toward low \mathcal{M}_{50} persists. We project that observations of one (two) complete satellite populations can constrain warm DM models with $m_{\text{WDM}} \approx 10$ keV (20 keV). Subhalo mass function (SHMF) suppression can be constrained to $\approx 70\%$, 60% , and 50% that in cold dark matter (CDM) at peak virial masses of 10^8 , 10^9 , and $10^{10} M_\odot$, respectively; SHMF enhancement constraints are weaker (≈ 20 , 4, and 2 times that in CDM, respectively) due to galaxy–halo connection degeneracies. These results motivate searches for faint dwarf galaxies beyond the MW and indicate that ongoing missions like Euclid and upcoming facilities including the Vera C. Rubin Observatory and Nancy Grace Roman Space Telescope will probe new galaxy formation and DM physics.

Unified Astronomy Thesaurus concepts: Dark matter (353); Dwarf galaxies (416); Galaxy abundances (574); Galaxy formation (595)

1. Introduction

The faintest dwarf galaxies inhabit the smallest dark matter halos that can form galaxies. These systems are sensitive probes of galaxy formation models and simultaneously test dark matter (DM) particle physics (e.g., Bullock & Boylan-Kolchin 2017; Gluscevic et al. 2019; Simon 2019; Nadler 2021; Sales et al. 2022), including DM free-streaming, interactions, and decays, wave interference of ultralight DM, and non-standard DM production mechanisms (e.g., see Bechtol et al. 2022 for an overview).

To disentangle the information dwarf galaxy surveys can offer, it is crucial to understand the interplay between galaxy formation and DM physics in low-mass halos. For example, reionization inhibits star formation in halos with peak masses below $\approx 10^8$ – $10^9 M_\odot$ because of their low virial temperatures and limited gas reservoirs (e.g., Benítez-Llambay & Frenk 2020; Munshi et al. 2021; Kravtsov & Manwadkar 2022; Ahvazi et al. 2024), suppressing the population of dwarf galaxies that occupy these systems. Observations of the faintest dwarf galaxies can shed light on these processes and can also reveal deviations in underlying halo populations from cold dark

matter (CDM), which would signal new DM physics on small scales. Moreover, a measurement of the galaxy formation cutoff is necessary to infer the existence of “dark” (galaxy-free) halos—which remain a key, untested predicted of many DM models—in combination with gravitational probes of even smaller halos (e.g., Gilman et al. 2020a; Banik et al. 2021b; Nadler et al. 2021b).

Current constraints on the galaxy formation cutoff are limited by the small number of known ultrafaint dwarf galaxies, which have almost exclusively been detected as satellites of the Milky Way (MW; e.g., Drlica-Wagner et al. 2020). Additionally, this measurement is plagued by large theoretical uncertainties on the faint end of the galaxy–halo connection (e.g., see Wechsler & Tinker 2018 for a review). Recently, Nadler et al. (2020b) analyzed the full-sky MW satellite population detected by the Dark Energy Survey (DES; Abbott et al. 2018) and Pan-STARRS1 (PS1; Chambers et al. 2016) to place a 95% confidence upper limit of $8.5 \times 10^7 M_\odot$ on the peak virial mass at which 50% of halos host galaxies, hereafter \mathcal{M}_{50} .⁸ This constraint accounts for observational selection effects, marginalizes over the relevant theoretical uncertainties, and implies stringent limits on a variety of DM scenarios that suppress the abundance of halos with peak



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

⁸ We define “galaxies” as dark matter-dominated systems with absolute V -band magnitudes $M_V < 0$ mag, which are composed of $\gtrsim 100$ stars in our fiducial models, on average, following Nadler et al. (2020b).

masses of $\approx 10^8 M_\odot$ (e.g., Nadler et al. 2021c; Newton et al. 2021; Dekker et al. 2022). However, current data do not permit a measurement of the galaxy formation cutoff because only tens of ultrafaints have been detected; this number is too small to rule out statistical fluctuations that mimic a cutoff or to overcome theoretical uncertainties on the predicted abundance of MW satellite galaxies.

Upcoming observational facilities will substantially improve this situation by expanding our census of faint dwarf galaxies. For example, the Vera C. Rubin Observatory (Rubin; Ivezić et al. 2008) is expected to detect much of the remaining MW satellite population, as well as many resolved dwarf galaxies out to ≈ 5 Mpc (Drlica-Wagner et al. 2019; Mutlu-Pakdil et al. 2021); comparable sensitivity has been projected for the Chinese Space Station Telescope (Qu et al. 2023). The Nancy Grace Roman Space Telescope (Roman; Spergel et al. 2015) will also be a powerful dwarf galaxy discovery machine (Gezari et al. 2022), with the ability to detect resolved stars out to even larger distances than Rubin. The Euclid Dark Energy Mission (Euclid Collaboration et al. 2022; launched 2023 July 1) is expected to achieve similar sensitivity (Racca et al. 2016). Meanwhile, despite its small field of view, JWST can resolve stars out to much larger distances than Roman or Euclid and thus may play an important role in detecting faint dwarf galaxies throughout the Local Volume and beyond (e.g., Conroy et al. 2023). Looking further ahead, next-generation optical/infrared space telescopes such as the Habitable Worlds Observatory, which is the top priority of the Astro2020 Decadal Survey, will revolutionize the search for dwarf galaxies at cosmic distances (e.g., see The LUVOIR Team 2019 for estimates of future large space telescopes' detection sensitivities).

To prepare for these advances, we study measurements of galaxy formation and DM physics in the context of future dwarf galaxy surveys. In particular, we use a suite of MW-mass cosmological zoom-in simulations (Nadler et al. 2023) to assess the sensitivity of satellite galaxy populations to a reionization-like galaxy formation cutoff, a suppression of low-mass halo abundances characteristic of warm dark matter (WDM), and the amplitude of the subhalo mass function (SHMF). We focus on subhalos and satellite galaxies (rather than isolated halos and field galaxies) because, as discussed above, future observational facilities capable of resolving ultrafaint dwarfs beyond the Local Group will have relatively small fields of view; thus, massive central galaxies will provide natural targets around which to search for dwarf satellite galaxies in deep pointings. We focus on MW-mass systems because lower-mass centrals host fewer satellites and because MW-mass systems are prevalent in the Local Volume (e.g., Mutlu-Pakdil et al. 2021; Carlsten et al. 2022) and nearby Universe (e.g., Geha et al. 2017; Mao et al. 2021).

By combining our simulations with an expanded version of the galaxy–halo connection model and probabilistic framework from Nadler et al. (2019b, 2020b, 2021c), we explore degeneracies between galaxy formation and DM effects on dwarf satellite populations and quantify signals of new physics that future galaxy surveys can uncover. We forecast that observations of all satellite galaxies around two MW-mass hosts (e.g., the MW and M31 or another nearby galaxy) can constrain a galaxy formation cutoff at peak virial halo masses of $10^8 M_\odot$ and place stringent upper limits on cutoffs at lower halo masses. We show that the same data can constrain WDM

models with masses of ≈ 10 –20 keV, significantly improving upon current constraints (Bullock & Boylan-Kolchin 2017; Drlica-Wagner et al. 2019). Furthermore, we find that the SHMF can be measured in a model-independent fashion at peak virial masses between $10^8 M_\odot$ and $10^{10} M_\odot$, paving the way for tests of any DM or early-Universe physics that suppresses or enhances low-mass halo abundances relative to CDM.

This paper is organized as follows. In Section 2, we describe our simulations and galaxy–halo connection modeling framework; in Section 3, we describe our procedure for forecasting constraints from future dwarf measurements. We then present the results of our galaxy formation (Section 4), WDM (Section 5), and SHMF (Section 6) forecasts. We discuss our results in the context of upcoming and future observations, along with areas for theoretical development, in Section 7. We discuss the implications of our forecasts for the detection of “dark” halos in Section 8, and we conclude in Section 9.

We adopt the same cosmological parameters used for our MW zoom-in simulations: $h = 0.7$, $\Omega_m = 0.286$, $\Omega_\Lambda = 0.714$, $\sigma_8 = 0.82$, and $n_s = 0.96$ (Hinshaw et al. 2013). We quote (sub) halo virial masses using the Bryan & Norman (1998) virial overdensity, which corresponds to $\Delta_{\text{vir}} \approx 99$ times the critical density of the Universe at $z = 0$ in this cosmology. Throughout, we refer to DM halos within a host’s virial radius as “subhalos,” and we denote galaxies that occupy a host’s subhalos as “satellites.” Furthermore, “log” always refers to the base-10 logarithm.

2. Modeling Dwarf Galaxy Populations

We begin with an overview of our forecasting framework, including the cosmological zoom-in simulations we base our forecasts on (Section 2.1), our galaxy–halo connection model (Section 2.2), and our treatment of dwarf population statistics in beyond-CDM scenarios (Section 2.3); we then illustrate the predictions of our model for subhalo and dwarf satellite galaxy populations (Section 2.4).

2.1. Cosmological Zoom-in Simulations

Our forecasts are based on subhalo populations from cosmological dark matter-only zoom-in simulations of MW-mass systems. Specifically, we use the “Milky Way-mass” suite from the Symphony zoom-in compilation (Mao et al. 2015; Nadler et al. 2023). This suite contains 45 host halos in a virial mass range of $10^{12.09 \pm 0.02} M_\odot$ and accurately resolves the abundance of subhalos with present-day virial masses of $M_{\text{vir}} > 1.2 \times 10^8 M_\odot$ (Nadler et al. 2023). Subhalos below this convergence threshold contribute negligibly to the observable satellite populations in our main forecasts, which are restricted to dwarf galaxies with $M_V < 0$ mag that occupy subhalos with $M_{\text{vir}} > 10^8 M_\odot$, on average (Nadler et al. 2020b). Thus, poorly resolved subhalos do not impact our results. We use the entire subhalo population within the virial radius of each MW-mass system—with no additional constraints imposed on e.g., subhalo distance from the host center, infall time, or size—to forward-model the entire observable satellite population of each host.

Our simulated MW-mass hosts have a range of secondary halo properties and inhabit a variety of cosmic environments (Fielder et al. 2019; Nadler et al. 2023). Thus, host-to-host scatter in subhalo populations, which results in

subhalo abundance variations up to a factor of ≈ 2 at fixed host halo mass (Mao et al. 2015), is naturally included in our predictions. However, we do not attempt to model the specific region of the nearby Universe inhabited by the closest 40 MW-mass systems, which provide natural observational targets for our forecasts. In particular, our simulations span a narrow range of host halo mass, whereas there is substantial variation in stellar mass and inferred host halo mass for nearby central galaxies. For example, the volume-limited sample of 31 massive hosts within 12 Mpc from Carlsten et al. (2022) spans central stellar masses of $9.9 \lesssim \log(M_*/M_\odot) \lesssim 11.1$, and Danieli et al. (2023) models 27 hosts from this population (excluding galaxy groups) with semianalytic realizations of host halos with $10.5 \lesssim \log(M_{\text{host}}/M_\odot) \lesssim 13.3$. Our forecasts should therefore be viewed as illustrative estimates of the constraining power from future dwarf satellite galaxy surveys rather than specific predictions for upcoming observations. We discuss aspects of our simulations and galaxy–halo connection model that may be generalized to capture a broader range of host halo masses, secondary properties, and environments in Section 7.3.

2.2. Galaxy–Halo Connection Model

We use the CDM galaxy–halo connection model from Nadler et al. (2019b, 2020b) with a few minor modifications described below. This model associates satellite galaxies with subhalos using extrapolated abundance-matching relations to predict satellite luminosity and size. Each satellite’s absolute V -band magnitude, M_V , is predicted from its subhalo’s peak maximum circular velocity, V_{peak} , using an extrapolated abundance-matching relation anchored to the GAMA luminosity function at $M_r = -13$ mag (Loveday et al. 2015), assuming a value for the faint-end luminosity function slope α , and for the scatter σ_M in absolute magnitude at fixed V_{peak} .⁹ Both σ_M and α are free parameters in our forecasts. Stellar mass–halo mass relation constraints derived from independent galaxy samples (e.g., Danieli et al. 2023) are consistent with the relation derived from GAMA, lending confidence to our approach.

Each satellite’s mean azimuthally averaged projected half-light radii, $r_{1/2}$ is predicted following Kravtsov (2013) via

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

where $R_{\text{vir,acc}}$ denotes a subhalo’s virial radius, measured when it accretes into its host’s virial radius; $R_0 = 10$ kpc is a fixed normalization constant. The amplitude parameter \mathcal{A} and power-law slope n are free parameters in our analyses. After mean sizes are predicted for a given set of model parameters, each satellite’s half-light radius is drawn from a lognormal distribution with scatter $\sigma_{\log R}$, which is also a free parameter. A power-law galaxy–halo size relation as in Equation (1) has been shown to fit a variety of hydrodynamic simulation results reasonably well (e.g., Jiang et al. 2019), although some semianalytic models (SAMs) predict a break in the size relation for ultrafaint dwarf galaxies, which we do not attempt to model (e.g., Kravtsov & Manwadkar 2022).

⁹ After performing abundance matching, we convert r -band to V -band magnitude via $M_V = M_r + 0.2$ mag (Geha et al. 2017; Nadler et al. 2019b).

We model the fraction f_{gal} of subhalos that host galaxies as a function of peak subhalo virial mass, M_{peak} , using a modified version of the model in Graus et al. (2019) via

$$f_{\text{gal}}(M_{\text{peak}}) = \frac{1}{2} \left[1 + \text{erf} \left(\frac{\log(M_{\text{peak}}/M_\odot) - \log(M_{50}/M_\odot)}{\sqrt{2} \mathcal{S}_{\text{gal}}} \right) \right], \quad (2)$$

where M_{50} is the peak virial mass at which 50% of halos host galaxies. \mathcal{S}_{gal} parameterizes the occupation fraction shape; smaller \mathcal{S}_{gal} implies a steeper cutoff, and f_{gal} approaches a step function at M_{50} as $\mathcal{S}_{\text{gal}} \rightarrow 0$.¹⁰ Both M_{50} and \mathcal{S}_{gal} are free parameters in our forecasts. Note that occupation fraction shapes predicted by hydrodynamic simulations and SAMs are similar to Equation (2) (e.g., Munshi et al. 2021; Kravtsov & Manwadkar 2022; Ahvazi et al. 2024).

Finally, we model subhalo disruption due to the central galaxy using the model from Nadler et al. (2018), which was calibrated to the hydrodynamic zoom-in simulations from Garrison-Kimmel et al. (2017b), which used the Feedback In Realistic Environments (FIRE) code. We apply this model to the subhalos of each host in the Symphony MW suite while allowing for variations in the strength of subhalo disruption relative to FIRE by setting

$$p_{\text{disrupt}} = (p_{\text{disrupt,0}})^{1/\mathcal{B}}, \quad (3)$$

where $p_{\text{disrupt,0}}$ is the fiducial subhalo disruption probability assigned by the Nadler et al. (2018) model, which depends on each subhalo’s mass and maximum circular velocity at infall, accretion time, and orbital properties. Note that we do not modify $p_{\text{disrupt,0}}$ in our beyond-CDM scenarios because it mainly depends on a given subhalo’s accretion time and orbital properties (Garrison-Kimmel et al. 2017b; Nadler et al. 2018), which do not significantly differ between CDM and models like WDM (e.g., Knebe et al. 2008).

To implement the disruption model in our analysis, each satellite is weighted by $1 - p_{\text{disrupt}}$ when number counts are computed from realizations of our model. In Equation (3), \mathcal{B} is a free parameter that controls the strength of disruption. In particular, smaller (larger) values of \mathcal{B} correspond to weaker (stronger) disruption due to baryons, and $\mathcal{B} = 1$ corresponds to the fiducial amount of disruption measured in FIRE simulations. We adopt a wide prior on \mathcal{B} because the efficiency of subhalo disruption due to central galaxies (e.g., Webb & Bovy 2020; Green et al. 2022) and the potential numerical origins of CDM subhalo disruption in general (van den Bosch et al. 2018; Errani & Navarro 2021; Green et al. 2021; Benson & Du 2022; Mansfield et al. 2023) are highly uncertain; our forecasts conservatively bracket these uncertainties.

For simplicity, we do not include orphan subhalos in our forecasts. Nadler et al. (2020b) showed, using the orphan model from Nadler et al. (2019b), that orphans do not affect galaxy–halo connection constraints derived from current MW satellite observations in our modeling framework (however, see Bose et al. 2020). Because our main analyses do not consider fainter satellites (and thus lower-mass subhalos) than the Nadler et al. (2020b) analysis, we do not expect orphans to impact our results. Moreover, the convergence analyses in

¹⁰ This parameter is dimensionless, so we label it as \mathcal{S}_{gal} rather than σ_{gal} as in Nadler et al. (2020b).

Nadler et al. (2023) indicate that our simulations accurately predict subhalo abundances above the fiducial resolution cuts described above. Modeling dwarf galaxies that occupy halos with $M_{\text{vir}} \lesssim 10^8 M_{\odot}$ would require modeling highly stripped subhalos, which is difficult for standard halo finders (Mansfield et al. 2023) but will be an important area for future work.

In summary, we model satellite galaxy populations in CDM with eight unknown parameters: the galaxy–halo luminosity parameters α and σ_M , the galaxy–halo size parameters \mathcal{A} , n , and $\sigma_{\log R}$, the occupation fraction parameters \mathcal{M}_{50} and \mathcal{S}_{gal} , and the disruption parameter \mathcal{B} .

2.3. Modeling Subhalo Populations Beyond CDM

The galaxy–halo connection model described in Section 2.2 can easily be applied to beyond-CDM scenarios. Here, we rely on subhalo populations from our CDM simulations (Section 2.1) to analyze WDM and models with a generalized SHMF. Thus, our analyses of these scenarios rely on the assumption that the following input subhalo properties do not significantly differ relative to CDM: V_{peak} (for the luminosity model), $R_{\text{vir,acc}}$ (for the size model), M_{peak} (for the occupation fraction model), and (mainly) accretion time and orbital properties (for the disruption model). In addition, the radial and angular distribution of subhalos is taken directly from our CDM simulations. We justify these assumptions for the WDM and generalized SHMF cases below.

2.3.1. Thermal-relic WDM

We consider thermal-relic WDM as a benchmark model with suppressed small-scale power and low-mass (sub)halo abundances (Bond & Szalay 1983; Bode et al. 2001). The WDM particle mass, m_{WDM} , sets the free-streaming scale below which the linear matter power spectrum is suppressed. We parameterize this cutoff via (Nadler et al. 2021c)

$$M_{\text{hm}}(m_{\text{WDM}}) = 5 \times 10^8 \left(\frac{m_{\text{WDM}}}{3 \text{ keV}} \right)^{-10/3} M_{\odot}, \quad (4)$$

where the half-mode mass, M_{hm} , is related to the unknown WDM particle mass, m_{WDM} . In particular, M_{hm} is defined by the wavenumber at which the linear matter power spectrum is suppressed by 75% relative to CDM (Nadler et al. 2019a). Thus, theoretical uncertainties in Equation (4) are small, and we hold this relation fixed in our forecasts.¹¹

We model the suppression of the SHMF using the fit to WDM zoom-in simulations from Lovell et al. (2014),

$$f_{\text{WDM}}(M_{\text{peak}}, m_{\text{WDM}}) = \left[1 + \left(\frac{\alpha M_{\text{hm}}(m_{\text{WDM}})}{M_{\text{peak}}} \right)^{\beta} \right]^{\gamma}, \quad (5)$$

where $f_{\text{WDM}} \equiv (dn_{\text{WDM}}/dM_{\text{peak}})/(dn_{\text{CDM}}/dM_{\text{peak}})$ with $\alpha = 2.7$, $\beta = 1.0$, and $\gamma = -0.99$. For the WDM analyses we present below, M_{hm} is a free parameter and each satellite is weighted by

¹¹ Vogel & Abazajian (2023) found that Equation (4) is slightly inaccurate, at the $\approx 10\%$ level, when compared to transfer functions output by linear Boltzmann solvers for $m_{\text{WDM}} \gtrsim 3 \text{ keV}$ (also see Appendix C of Decant et al. 2022). We do not correct for this discrepancy here because the (sub)halo mass function suppression we adopt is based on Equation (4); this will be addressed in a forthcoming study (E. O. Nadler et al. 2024, in preparation).

$f_{\text{WDM}} \times (1 - p_{\text{disrupt}})$ when number counts are computed from realizations of our model.

We focus on models with $m_{\text{WDM}} \gtrsim 5 \text{ keV}$ ($M_{\text{hm}} \gtrsim 10^8 M_{\odot}$). As a result, we primarily study halos above the half-mode mass (and well above the free-streaming mass) in our forecasts. In this regime, WDM simulations indicate that M_{peak} and V_{peak} are only reduced by $\mathcal{O}(10\%)$ relative to matched CDM halos (e.g., Bozek et al. 2019; Fitts et al. 2019). We expect $R_{\text{vir,acc}}$ to be reduced by a similar amount because it is mainly determined by M_{peak} and the amount of pre-infall stripping a subhalo experiences, which does not differ significantly between the models (e.g., Elahi et al. 2014). Meanwhile, as described in Section 2.2, orbital properties of given WDM subhalos are not significantly altered relative to CDM.

Finally, we assume that the WDM subhalo radial distribution is unchanged relative to CDM. WDM simulations of MW-mass hosts indicate that this is a good assumption for subhalos more massive than M_{hm} (Lovell et al. 2021), which we primarily study. Modeling lower-mass subhalos would benefit from the direct use of WDM simulations. Thus, our method for applying the model to WDM is well motivated, and we leave an extension of our framework to beyond-CDM simulations for future work.

2.3.2. Model-independent SHMF

We also consider a purely empirical modification to the SHMF, in which subhalo abundances vary above or below those in CDM, as a function of peak subhalo mass. Specifically, we bin in decades of subhalo mass and define

$$\xi_i \equiv \log [f_{\text{NCDM}}(10^{i-0.5} < M_{\text{peak}}/M_{\odot} < 10^{i+0.5})], \quad (6)$$

where $f_{\text{NCDM}} \equiv (dn_{\text{NCDM}}/dM_{\text{peak}})/(dn_{\text{CDM}}/dM_{\text{peak}})$. In this model, ξ_8 , ξ_9 , and ξ_{10} parameterize SHMF deviations from CDM centered on peak masses of 10^8 , 10^9 , and $10^{10} M_{\odot}$, respectively. For our SHMF forecasts, ξ_8 , ξ_9 , and ξ_{10} are free parameters, and each satellite is weighted by $f_{\text{NCDM}} \times (1 - p_{\text{disrupt}})$ when number counts are computed from realizations of our model.

By construction, our generalized SHMF scenarios do not alter any of the subhalo properties relevant for our galaxy–halo connection model or quantities like the radial distribution. Although these assumptions may be broken in specific beyond-CDM scenarios, this approach allows us to isolate the effects of varying the SHMF.

2.4. Model Illustration

We will generate “true” satellite populations for our forecasts by evaluating the galaxy–halo connection model at fixed values of σ_M , \mathcal{M}_{50} , \mathcal{S}_{gal} , and beyond-CDM parameters; the remaining parameters are always set to the best-fit values derived from the MW satellite population in Nadler et al. (2020b), i.e., $\alpha = -1.44$, $\mathcal{B} = 0.93$, $\sigma_{\log R} = 0.63 \text{ dex}$, and $\mathcal{A} = 37 \text{ pc}$. However, all eight galaxy–halo connection parameters and additional beyond-CDM parameters are then left free and inferred in our forecasting analyses.

Before performing these full parameter recovery tests, we illustrate the predictions of our model and the constraining power of satellite population measurements. Figure 1 shows mock observations of mean satellite luminosity functions from our 40 MW-mass zoom-in simulations. We compare mock observations of “true” satellite luminosity functions in

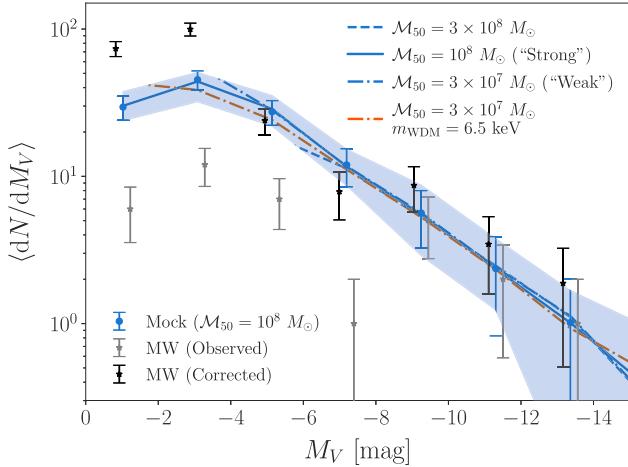


Figure 1. Mean satellite luminosity function and 1σ Poisson uncertainty for the MW satellite population observed by DES and PS1 (gray), and a completeness-corrected total luminosity function from (Drlica-Wagner et al. 2020; black). Blue error bars show the predicted luminosity function mean and 1σ Poisson uncertainty for 40 complete satellite populations of MW-mass hosts generated from our galaxy–halo connection model with a “strong” galaxy formation cutoff ($M_{50} = 10^8 M_\odot$, $\sigma_M = 0.2$ dex, and $S_{\text{gal}} = 0.2$; all remaining galaxy–halo connection parameters are fixed to the input values listed in Section 2.4). Blue lines show predictions for the same σ_M and S_{gal} , with $M_{50} = 3 \times 10^8$ (dashed), 10^8 (solid), and $3 \times 10^7 M_\odot$ (dotted-dashed, which we refer to as a “weak” cutoff throughout). The blue band shows 16th–84th percentile host-to-host scatter for $M_{50} = 10^8 M_\odot$. All predictions assume a satellite detectability threshold of $M_V < 0$ mag and $\mu < 32$ mag arcsec $^{-2}$, and error bars are offset horizontally for clarity.

a model with $M_{50} = 10^8 M_\odot$ (shown as error bars) to predictions with $M_{50} = 3 \times 10^7$, 10^8 , and $3 \times 10^8 M_\odot$ (shown as lines). Below, we refer to $M_{50} = 10^8 M_\odot$ ($3 \times 10^7 M_\odot$) as a “strong” (“weak”) cutoff scenario, and we set $\sigma_M = 0.2$ dex, and $S_{\text{gal}} = 0.2$ in both scenarios; see Section 4 for details. Note that all remaining galaxy–halo connection parameters are fixed to the input values listed above.

For this illustration, we optimistically assume that all satellites of each host with $M_V < 0$ mag and $\mu_V < 32$ mag arcsec $^{-2}$ are detectable; we expand on our observational assumptions in Section 3.1 and discuss them in the context of near and long-term observational capabilities in Sections 7.1 and 7.2. Given this assumption, the model with $M_{50} = 3 \times 10^8 M_\odot$ is detectable using the complete satellite population of a single MW-mass host. Meanwhile, models with $M_{50} \lesssim 10^8 M_\odot$ are more difficult to detect because they only affect the abundances of the faintest satellites we consider. For comparison, the gray points in Figure 1 show the luminosity function of kinematically confirmed and candidate MW satellites detected by DES and PS1 (Drlica-Wagner et al. 2020). Because MW satellite observations are incomplete, we also show the completeness-corrected estimate from Drlica-Wagner et al. (2020) for comparison. Note that Nadler et al. (2020b) inferred $M_{50} < 8.5 \times 10^7 M_\odot$ at 95% confidence using these satellites, including the observational selection function derived in Drlica-Wagner et al. (2020).

Figure 1 also shows a “weak” galaxy formation cutoff combined with an $m_{\text{WDM}} = 6.5$ keV WDM model, corresponding to the 95% confidence WDM limit from the MW satellite population derived in Nadler et al. (2021c). These astrophysical and DM models combine to yield a luminosity

function similar to a “strong” galaxy formation cutoff in CDM, illustrating how degeneracies between galaxy formation and DM physics can manifest in dwarf galaxy populations.

To illustrate our beyond-CDM predictions, the left panel of Figure 2 compares the suppression of dwarf galaxy abundances in our “strong” and “weak” galaxy formation cutoff scenarios to the suppression of the SHMF in various WDM models. We show WDM masses of $m_{\text{WDM}} = 4.9$ keV (the fiducial value in our WDM forecasts), 5.6 keV (the limit from one complete satellite population in our results below, assuming a “strong” galaxy formation cutoff), 7.9 keV (constrained at 68% confidence by the MW satellite population; Nadler et al. 2021c), and 17.0 keV (constrained by two complete satellite in our results below, again assuming a “strong” cutoff). The shapes of the galaxy occupation fraction and SHMF cutoffs differ in detail; thus, sufficiently precise measurements of dwarf galaxy abundances can potentially disentangle these effects. However, note that the occupation fraction shape is fixed to $S_{\text{gal}} = 0.2$ in this illustration, while our forecasts treat S_{gal} as an unknown parameter.

Finally, the right panel of Figure 2 illustrates our model-independent SHMF parameterization. We show uniformly suppressed (dotted red) and enhanced (dashed blue) models constrained at 68% confidence by one complete satellite population in our forecasts; we also show a model that mixes SHMF suppression and enhancement in different peak virial mass decades (dotted-dashed purple) at a level that is consistent with our projected dwarf galaxy sensitivity (see Section 6.1 for details). For context, the size of these SHMF deviations relative to CDM is comparable to the SHMF sensitivity inferred from recent stellar stream and strong lensing measurements (Gilman et al. 2020a; Banik et al. 2021a). Both the shape and normalization of the SHMF can be altered in our parameterization, and only some of this parameter space can be mimicked by galaxy formation or WDM cutoffs.

3. Forecasting Framework

Our forecasting framework starts from realizations of the simulated satellite population following Section 2. We then analyze these mock data using probabilistic inference. Our observational assumptions are outlined in Section 3.1, our likelihood formalism is described in Section 3.2, and our evaluation of the likelihood is described in Section 3.3.

3.1. Observational Assumptions

We make the following observational assumptions for all forecasts presented below:

1. All satellite galaxies within the virial radius of each MW-mass host that have $M_V < 0$ mag and $\mu_V < 32$ mag arcsec $^{-2}$ are detectable, where $\mu_V \equiv M_V + 36.57 + 2.5 \log[2\pi(r_{1/2}/\text{kpc})^2]$ is the effective surface brightness averaged within the half-light radius.
2. Host halo properties (and all latent variables that determine a host’s subhalo population) are perfectly known.

The second assumption is implicit in our procedure and follows because we compare predicted satellite populations, generated using the subhalo population around a given simulated host, to a “true” satellite population generated from the same host. Both assumptions are optimistic; we make these choices in order to

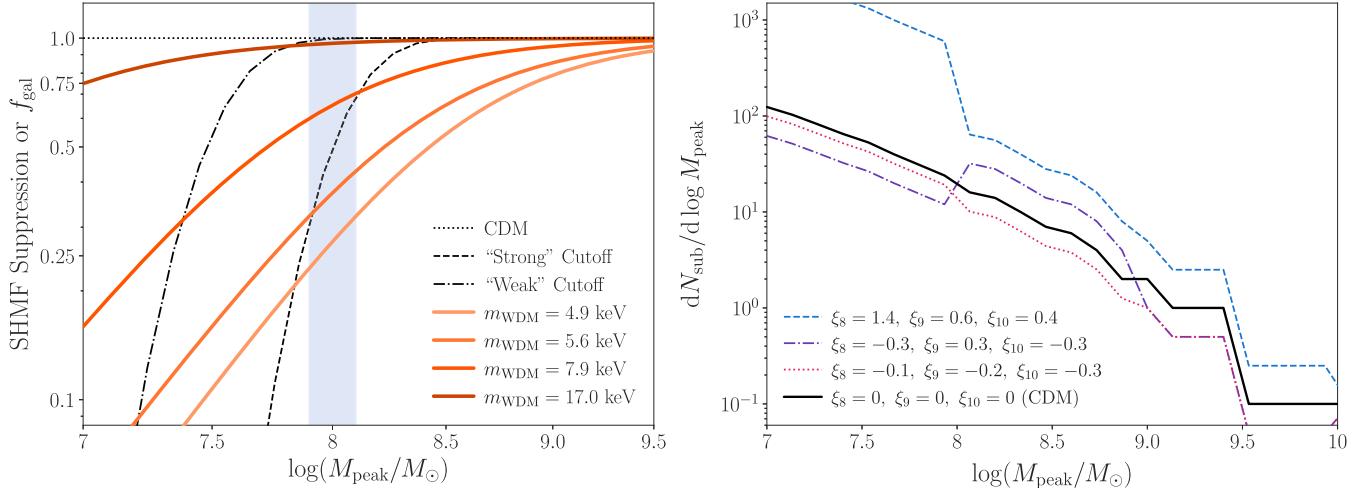


Figure 2. Left: suppression of the subhalo mass function (defined as the ratio of the WDM to CDM SHMF) for WDM models with $m_{\text{WDM}} = 4.9, 5.6, 7.9$, and 17.0 keV from lightest to darkest red. Our “strong” and “weak” galaxy formation cutoff scenarios are shown on the same axes. The blue band shows the range of \mathcal{M}_{50} probed at 68% confidence by our “strong” galaxy formation cutoff forecast, and thus illustrates the range of halo masses we typically probe. Right: mean SHMF (black) from our 40 CDM MW-mass zoom-in simulations, measured using peak virial mass. Dashed blue, dotted-dashed purple, and dotted red lines illustrate the effects of our model-independent SHMF parameterization, which can enhance or suppress the SHMF relative to CDM as a function of peak halo mass, depending on the values of ξ_8 , ξ_9 , and ξ_{10} . The uniformly suppressed (dashed blue) and enhanced (dotted red) models correspond to our projections for 68% confidence SHMF constraints from one complete host.

assess the potential of dwarf satellite populations for constraining galaxy formation and DM physics. We place these assumptions in the context of current and future observational sensitivity and discuss related areas for modeling work in Section 7.

3.2. Likelihood

To fit the mock data and assess parameter uncertainties, we generalize the likelihood framework from Nadler et al. (2019b, 2020b), following Nadler (2021). In particular, we model realizations of satellite luminosities and sizes as a Poisson process, drawn from a distribution uniquely predicted by a set of parameter values $\boldsymbol{\theta} = \{\alpha, \sigma_M, \mathcal{A}, n, \sigma_{\log R}, \mathcal{M}_{50}, \mathcal{S}_{\text{gal}}, \mathcal{B}\}$, and additionally M_{hm} or ξ_8, ξ_9 , and ξ_{10} in our beyond-CDM forecasts. The probability of detecting $n_{\text{obs},ij}$ satellites in luminosity and size bin j , in host i , is given by

$$P(n_{\text{obs},ij}|\boldsymbol{\theta}) = \left(\frac{N_{\text{real.}} + 1}{N_{\text{real.}}}\right)^{-(N_{ij}+1)} \times (N_{\text{real.}} + 1)^{-n_{\text{obs},ij}} \frac{(N_{ij} + n_{\text{obs},ij})!}{n_{\text{obs},ij}! N_{ij}!}, \quad (7)$$

where $N_{\text{real.}}$ is the total number of realizations per host at a given $\boldsymbol{\theta}$, $N_{ij} \equiv \sum_{k=1}^{N_{\text{real.}}} n_{\text{pred},ijk}$, and $n_{\text{pred},ijk}$ is the number of predicted satellites in host i , bin j , and realization k . In practice, we replace the factorials in Equation (7) with appropriate Gamma functions because our model produces noninteger satellite counts.

Before moving on, we provide context for this likelihood. In particular, Equation (7) is a modified Poisson likelihood derived by marginalizing over the (unknown) underlying Poisson rate in each bin of observable parameter space via

(Nadler et al. 2019b)

$$\begin{aligned} & P(n_{\text{obs},ij}|n_{\text{pred},ij1}, \dots, n_{\text{pred},ijN_{\text{real.}}}) \\ &= \int P(n_{\text{obs},ij}|\lambda_{ij}) P(\lambda_{ij}|n_{\text{pred},ij1}, \dots, n_{\text{pred},ijN_{\text{real.}}}) d\lambda_{ij} \\ &= \frac{1}{P(n_{\text{pred},ij1}, \dots, n_{\text{pred},ijN_{\text{real.}}})} \\ & \times \int P(n_{\text{obs},ij}|\lambda_{ij}) P(n_{\text{pred},ij1}|\lambda_{ij}) \dots P(n_{\text{pred},ijN_{\text{real.}}}|\lambda_{ij}) P(\lambda_{ij}) d\lambda_{ij} \\ &= \left(\frac{N_{\text{real.}} + 1}{N_{\text{real.}}}\right)^{-(n_{\text{pred},ij1} + \dots + n_{\text{pred},ijN_{\text{real.}}} + 1)} \\ & \times (N_{\text{real.}} + 1)^{-n_{\text{obs},ij}} \frac{(n_{\text{pred},ij1} + \dots + n_{\text{pred},ijN_{\text{real.}}} + n_{\text{obs},ij})!}{n_{\text{obs},ij}!(n_{\text{pred},ij1} + \dots + n_{\text{pred},ijN_{\text{real.}}})!}. \end{aligned} \quad (8)$$

Here, the final expression is obtained by substituting Poisson likelihoods for each conditional probability and a flat prior on positive values of λ_{ij} . The key assumption is that mock and observed satellites are drawn from the same (unknown) Poisson distribution; our likelihood converges to this distribution in the limit of many realizations (Nadler et al. 2019b).

When evaluating Equation (7), we include $N_{\text{real.}}$ realizations at a given $\boldsymbol{\theta}$ because our galaxy–halo connection model is stochastic. In this study, we use $N_{\text{real.}} = 10$ per host at a given $\boldsymbol{\theta}$. We have checked that varying $N_{\text{real.}}$ does not impact our results; in particular, our Markov Chain Monte Carlo (MCMC) evaluation of the posterior converges in fewer steps when $N_{\text{real.}}$ is larger, but the posterior distribution and total run time are unchanged. We follow Nadler et al. (2020b) by using 14 bins in M_V , evenly spaced from 0 to -20 mag, and two bins in surface brightness that split the sample at $\mu_V = 28$ mag arcsec $^{-2}$.

The joint likelihood of observing all galaxies in our mock data vector \mathbf{s}_{obs} , at each set of theoretical parameter values $\boldsymbol{\theta}$ for all of the hosts we consider in a given analysis is

$$\mathcal{L}(\mathbf{s}_{\text{obs}}|\boldsymbol{\theta}) = \prod_{\text{hosts } i} \prod_{\text{bins } j} P(n_{\text{obs},ij}|\boldsymbol{\theta}). \quad (9)$$

The mock data vector s_{obs} is generated from one realization of our model with the same fixed galaxy–halo connection and (if applicable) beyond-CDM parameters for all hosts, with specific values described for each forecast below.

We use Bayes’ theorem to compute the joint posterior distribution over θ for each forecast,

$$P(\theta|s_{\text{obs}}) = \frac{\mathcal{L}(s_{\text{obs}}|\theta)P(\theta)}{P(s_{\text{obs}})}, \quad (10)$$

where $P(\theta)$ is the joint prior distribution over θ , and $P(s_{\text{true}})$ is the Bayesian evidence. Appendix A provides all prior distributions used in our forecasts.

3.3. Posterior Sampling

We sample the posterior in Equation (10) using the MCMC sampler EMCEE (Foreman-Mackey et al. 2013), using 450 walkers and a combination of the StretchMove and KDEMove sampling algorithms. For our CDM and WDM forecasts, $\sim 10^6$ samples are sufficient to accurately characterize the posterior, while $\sim 5 \times 10^6$ samples are required for our SHMF forecasts due to higher dimensionality and more prominent degeneracies. We discard ~ 100 burn-in autocorrelation lengths for each chain, which leaves more than 100 independent samples for every scenario we consider.

It is difficult to sample our full galaxy–halo connection (plus beyond-CDM) parameter space using more than two hosts at a time. Because our model is evaluated on high-resolution cosmological simulations, likelihood evaluations are computationally expensive when using large numbers of hosts. In addition, degeneracies among galaxy–halo connection parameters (and, in our WDM and SHMF forecasts, between galaxy–halo connection and beyond-CDM parameters) slow down posterior sampling. The most challenging parameter in this context is σ_M , which is often degenerate with \mathcal{M}_{50} , \mathcal{S}_{gal} , and beyond-CDM parameters that mimic a cutoff; we discuss specific degeneracies in each forecast below. Furthermore, because the SHMF rises steeply with decreasing subhalo mass, abundant low-mass subhalos preferentially up-scatter to observable luminosities (a form of Eddington bias), resulting in non-Gaussian posteriors for individual hosts. The complexity of the fit increases with the number of hosts used to compute the joint likelihood since each host has a different underlying subhalo population and associated non-Gaussian posterior.

For each forecast, we therefore evaluate the (joint) likelihood using one and two hosts to derive our key results and characterize degeneracies in the posterior. We have checked that our posteriors for individual hosts are statistically consistent among all 45 of our Symphony MW zoom-ins, and we describe how constraints depend on combinations of two hosts below. To determine how the uncertainty of our forecasted constraints scales with the number of hosts used in the inference, N_{hosts} , we supplement our MCMC forecasts with projections for galaxy formation cutoff, WDM, and SHMF constraints that fix all remaining galaxy–halo connection parameters to their input values.

4. Galaxy Formation Forecasts

We first forecast sensitivity to measuring a cutoff in galaxy formation within CDM, for two fiducial values of \mathcal{M}_{50} :

1. “Weak” cutoff scenario: $\mathcal{M}_{50} = 3 \times 10^7 M_{\odot}$;

2. “Strong” cutoff scenario: $\mathcal{M}_{50} = 10^8 M_{\odot}$.

In both cases, we set $\sigma_M = 0.2$ dex, and $\mathcal{S}_{\text{gal}} = 0.2$. This value of σ_M is representative of the galaxy–halo connection scatter inferred at higher masses (Wechsler & Tinker 2018) and is near the upper limit derived from current MW satellite observations (Nadler et al. 2020b); thus, it is a helpful benchmark for our study. Meanwhile, our fiducial value of \mathcal{S}_{gal} yields a galaxy occupation fraction that agrees with the predictions of state-of-the-art galaxy formation models.

In particular, the “weak” cutoff scenario is broadly consistent with predictions from high-resolution simulations of galaxy formation at high redshifts (e.g., Côté et al. 2018) and from various SAMs (e.g., Kravtsov & Manwadkar 2022; Ahvazi et al. 2024). We note that the “weak” cutoff closely matches the mean galaxy occupation fraction predicted when both atomic and molecular hydrogen cooling are included in the Ahvazi et al. (2024) SAM. Meanwhile, the “strong” cutoff scenario is more consistent with predictions from hydrodynamic simulations run to lower redshifts (e.g., Sawala et al. 2016; Fitts et al. 2017; Munshi et al. 2021) and other SAMs (e.g., Benitez-Llambay & Frenk 2020). Furthermore, the “strong” cutoff is similar to (but slightly weaker than) the mean galaxy occupation fraction predicted when only atomic hydrogen cooling is included in the Ahvazi et al. (2024) SAM. We note that galaxy occupation in hydrodynamic simulations is resolution- and definition-dependent (e.g., Nadler et al. 2020b; Ahvazi et al. 2024), and that higher-resolution simulations (e.g., Wheeler et al. 2019) and convergence studies (e.g., Munshi et al. 2021) indicate that progressively lower-mass halos host “galaxies” (i.e., simulated systems with more than a critical number of star particles) as resolution increases.

The “strong” cutoff scenario is informative for our forecasts because it clearly reveals degeneracies between the galaxy formation cutoff, other galaxy–halo connection parameters, and beyond-CDM parameters. Note that current observations of the MW satellite population from DES and PS1 are beginning to probe the “strong” cutoff and its assumed scatter of $\sigma_M = 0.2$ dex (Nadler et al. 2020b). Meanwhile, the “weak” cutoff is not currently constrained and represents a theoretical target for the sensitivity of future dwarf galaxy surveys. Our “strong” and “weak” cutoff scenarios therefore bracket a physically and observationally motivated range of galaxy formation cutoffs.

4.1. Cutoff Mass Scale Posteriors

Figure 3 illustrates our results for the “strong” cutoff scenario: the left panel shows the marginalized posterior for \mathcal{M}_{50} given observations of all satellites around one (blue) and two (red) MW-mass hosts, and the right panel shows the corresponding galaxy occupation fraction posteriors. We find $\log(\mathcal{M}_{50}/M_{\odot}) = 7.9_{-1.6}^{+0.4}$ at 68% confidence for one host. Thus, \mathcal{M}_{50} can be probed at the $\approx 1\sigma$ level using a single complete satellite population; however, a long tail toward low \mathcal{M}_{50} prevents a measurement at the 2σ level. Note that our inferred distribution of \mathcal{M}_{50} is consistent with the input value of $\log(\mathcal{M}_{50}/M_{\odot}) = 8.0$, indicating that our galaxy–halo connection reconstruction is unbiased and that our inference framework is self-consistent.

The tail toward low \mathcal{M}_{50} arises because the “strong” cutoff primarily suppresses satellite abundances at the lowest

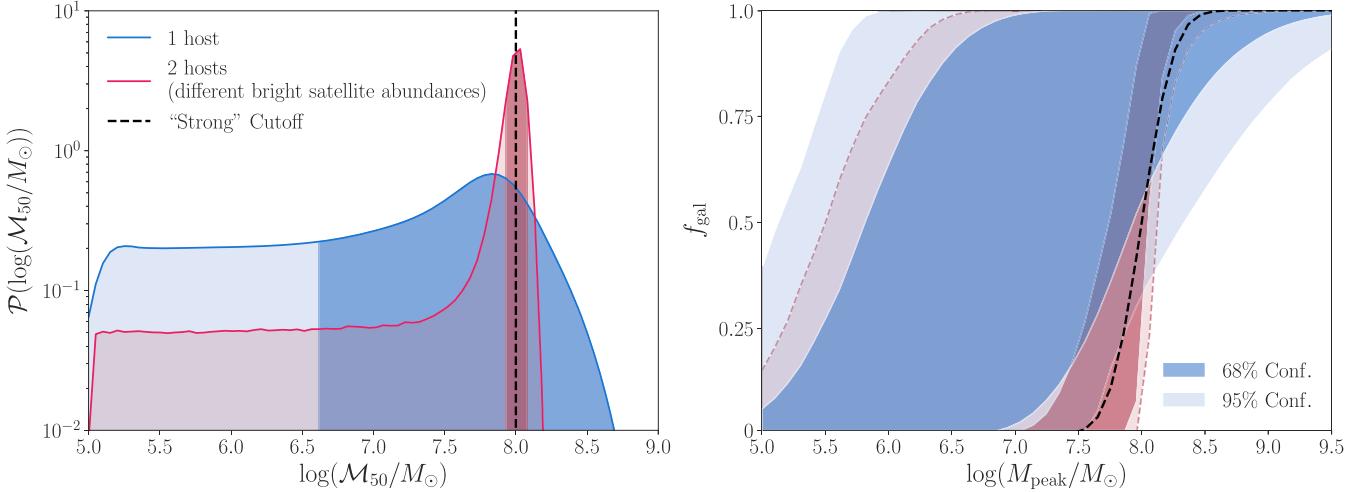


Figure 3. Left: marginalized posteriors for the galaxy formation cutoff mass scale \mathcal{M}_{50} (i.e., the peak virial mass at which 50% of halos host galaxies with $M_V < 0$ mag), from our forecasted measurements of one (blue) and two (red) complete satellite populations of MW-mass hosts. Dark (light) shaded regions show 68% (95%) confidence intervals. Right: marginalized posteriors for the galaxy occupation fraction, defined as the fraction of halos that host galaxies with $M_V < 0$ mag as a function of peak virial mass. Dark (light) bands show the 68% (95%) confidence intervals from our one- and two-host forecasts. In both panels, black dashed lines show the input “strong” galaxy formation cutoff.

luminosities we consider, and this effect is not significant enough to be distinguished from models that do not suppress the abundance of satellites with $M_V < 0$ mag and $\mu_V < 32$ mag arcsec $^{-2}$. As a result, models with $\mathcal{M}_{50} \lesssim 10^7 M_\odot$ are not distinguishable from each other and have roughly equal probability in the marginalized posterior. Our analysis of the full one-host posterior in Appendix C.1 shows that the low- \mathcal{M}_{50} tail is enhanced by a degeneracy with the luminosity scatter σ_M . For large σ_M , galaxies hosted by abundant, low-mass halos preferentially up-scatter in luminosity, resulting in higher predicted satellite abundances at all luminosities. In this regime, galaxy–halo connection parameters adjust so that lower-mass halos explain more of the mock data; in turn, large values of \mathcal{M}_{50} are disfavored because they prevent these low-mass halos from forming galaxies. In other words, higher-luminosity scatter increases the sensitivity of a fixed-depth survey to lower-mass subhalos, boosting the constraining power on large values of \mathcal{M}_{50} .

Combining the satellite populations of two hosts can improve this situation by setting a much tighter constraint on σ_M , leading to a less prominent tail toward low \mathcal{M}_{50} . We illustrate this effect in Figure 4 using two hosts with abundances of classical satellites ($M_V \lesssim -8$ mag) that differ at a level comparable to the 1σ host-to-host luminosity function scatter in our simulation suite, shown by the blue band in Figure 1. While σ_M is unconstrained in our “strong” cutoff forecast with one host, we recover $\sigma_M = (0.2 \pm 0.12)\text{dex}$ using these two hosts (recall that the input value is $\sigma_M = 0.2$ dex). In turn, this substantially improves our recovery of \mathcal{M}_{50} to $\log(\mathcal{M}_{50}/M_\odot) = 8.0 \pm 0.1$ at 68% confidence. However, \mathcal{M}_{50} remains difficult to measure at 95% confidence due to the low- \mathcal{M}_{50} tail, which persists even with two hosts. Using hosts with different bright satellite abundances effectively constrains σ_M because the host with more classical satellites can be fit by larger values of σ_M , while the host with fewer classical satellites cannot; thus, combining such hosts breaks the \mathcal{M}_{50} – σ_M degeneracy. Constraints on many other galaxy–halo connection parameters also significantly improve when such hosts are combined (see Appendix C.1).

In the “weak” cutoff scenario, our mock observations are less sensitive to \mathcal{M}_{50} because galaxy formation is suppressed in a smaller fraction of halos that host observable satellites. With one complete host and our fiducial detection thresholds, \mathcal{M}_{50} is only constrained at the $\approx 1\sigma$ level, with $\log(\mathcal{M}_{50}/M_\odot) = 7.1^{+0.3}_{-1.6}$; the tail toward low \mathcal{M}_{50} is even more prominent than in the “strong” cutoff scenario, which again prevents a 2σ measurement. These results do not significantly improve when combining two hosts, which only yield an upper limit on \mathcal{M}_{50} , with $\log(\mathcal{M}_{50}/M_\odot) = 5.3^{+1.4}_{-0.2}$.

Our two-host results in the “weak” cutoff scenario are less sensitive to the specific combination of hosts because the \mathcal{M}_{50} – σ_M degeneracy is not the limiting factor for \mathcal{M}_{50} constraints. Instead, for a “weak” cutoff, combining hosts with as many satellites as possible—rather than hosts with classical satellite abundances that differ substantially relative to the underlying host-to-host scatter—yields the strongest upper limits on \mathcal{M}_{50} . Note that current MW satellite observations mildly disfavor the “strong” cutoff (Nadler et al. 2020b). If future data strengthens this conclusion, our results imply that combining the Milky Way satellite population with hosts that have abundant bright satellite populations will maximize the galaxy formation constraints delivered by future surveys. Many such hosts have already been identified throughout the Local Volume (Carlsten et al. 2022) and nearby Universe (Geha et al. 2017; Mao et al. 2021).

4.2. Cutoff Shape Posteriors

With one complete satellite population (e.g., for a complete census of satellites only within the MW), the shape of the occupation fraction is difficult to recover in either galaxy formation cutoff scenario we consider. In particular, using one complete host in either the “strong” or “weak” cutoff scenario, our f_{gal} reconstruction falls off less slowly with decreasing halo mass than the input model over most of the allowed parameter space (see the right panel of Figure 3). This is consistent with the tail toward low \mathcal{M}_{50} and the broad \mathcal{S}_{gal} posterior in the one-host case (see Appendix C.1). Adding an additional host brings our f_{gal} reconstruction into agreement with the input model at

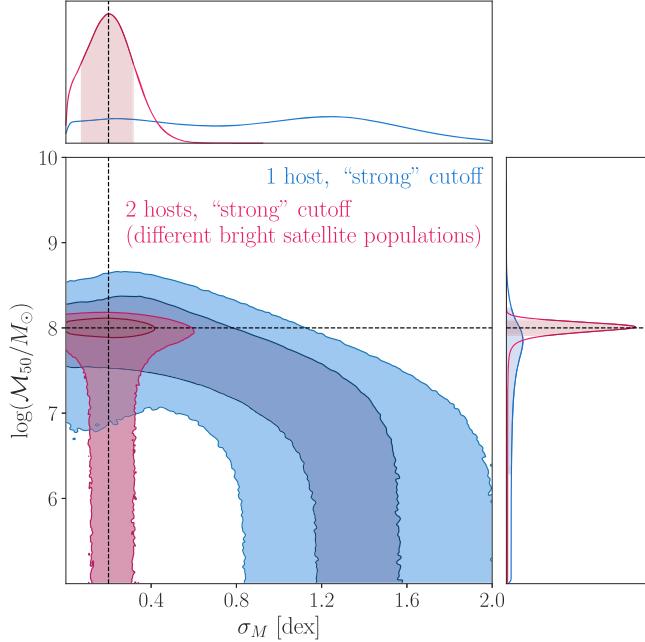


Figure 4. Marginalized posterior for \mathcal{M}_{50} vs. σ_M from our forecast using one (blue) and two (red) complete satellite populations of MW-mass hosts in the presence of a “strong” galaxy formation cutoff. Dashed lines mark the input values of \mathcal{M}_{50} and σ_M ; two-dimensional contours represent 68% and 95% confidence intervals. The top and side panels show marginal posteriors, where shaded regions represent 68% confidence intervals. The two hosts shown by the red contour have significantly different bright satellite abundances, which markedly improves the recovery of σ_M and indirectly improves \mathcal{M}_{50} constraints.

the $\approx 1\sigma$ level; however, we still infer shallower occupation fractions than the input model at the 2σ level because the tail toward low \mathcal{M}_{50} and large \mathcal{S}_{gal} (corresponding to a flat occupation fraction as a function of halo mass) remains consistent with the mock data. Our results suggest that more than two complete satellite populations are needed to probe the functional dependence of the occupation fraction on halo mass at the $\gtrsim 1\sigma$ level.

4.3. Projections for Additional Hosts

As discussed in Section 3.2, it is computationally challenging to run our full MCMC forecasts for more than two hosts. Instead, we derive projected constraints on \mathcal{M}_{50} by calculating the likelihood as a function of \mathcal{M}_{50} with all remaining galaxy–halo connection parameters fixed to their input values. We perform this test as a function of N_{hosts} , sampling realizations of our galaxy–halo connection model and different host combinations for each cutoff scenario and choice of N_{hosts} . This procedure optimistically neglects galaxy–halo connection degeneracies to illustrate the ideal scaling of \mathcal{M}_{50} constraints for statistically limited measurements; the resulting projections can be interpreted as cosmic variance-limited constraints from satellite populations.

Figure 5 shows projected 1σ measurement uncertainties on $\log(\mathcal{M}_{50}/M_{\odot})$ as a function of N_{hosts} in the “strong” cutoff scenario. We derive these uncertainties from the Fisher information evaluated at the input value of \mathcal{M}_{50} . For large N_{hosts} , the 1σ uncertainty on $\log(\mathcal{M}_{50}/M_{\odot})$ scales as $N_{\text{hosts}}^{-1/2}$, as expected when adding independent measurements. With 32 complete hosts, we project $\approx 2\%$ uncertainties on $\log(\mathcal{M}_{50}/M_{\odot})$ at the 1σ level, assuming negligible galaxy–

halo connection systematics; different host combinations affect these projected constraints at the subpercent level. For reference, there are ≈ 30 MW-mass hosts in the Local Volume with partially characterized satellite populations (e.g., Carlsten et al. 2022). Our projections do not significantly change in the “weak” cutoff scenario, although they should only be regarded as upper limits at the 1σ level in this case.

In Appendix B, we show that there is a prominent tail toward low \mathcal{M}_{50} in our idealized likelihoods for the “strong” and “weak” cutoff scenarios; this tail persists for large N_{hosts} , particularly for the “weak” cutoff. Thus, \mathcal{M}_{50} is difficult to measure at the 2σ level, even in the statistically limited regime, if subhalos with $M_{\text{peak}} \lesssim \mathcal{M}_{50}$ are not easily detectable. As a result, we do not expect a shallower survey than our fiducial assumption to detect either the “strong” or “weak” cutoff, even if it uses more hosts. On the other hand, constraining power qualitatively improves if satellites that do not pass our fiducial detection thresholds are included in the inference. For example, using more optimistic detection thresholds of $M_V < +2$ mag (comparable to the luminosity of the recently discovered Ursa Major III/UNIONS 1 dwarf galaxy candidate; Smith et al. 2023) and $\mu_V < 34$ mag arcsec $^{-2}$ (fainter than any known system and comparable to “stealth galaxies” hypothesized in previous studies; Bullock et al. 2010) disfavors models with low \mathcal{M}_{50} due to nondetections of extremely faint systems. Observations of satellites below our fiducial detection thresholds can therefore improve our projected galaxy formation cutoff constraints; we leave a study of these systems in the context of our modeling framework to future work.

4.4. Interpretation

Our forecasts indicate that complete measurements of two satellite populations (e.g., all satellites of the MW and M31) can probe a galaxy formation cutoff with $\mathcal{M}_{50} = 10^8 M_{\odot}$, at the 1σ level. Moreover, for statistically limited measurements, the uncertainty of this projected 1σ constraint shrinks as additional hosts are added. We emphasize that our forecasts only rely on halos that actually host galaxies. In particular, nondetections of faint, low surface brightness galaxies—which would primarily be hosted by halos with $M_{\text{peak}} \lesssim \mathcal{M}_{50}$ in the absence of a galaxy formation cutoff—play a crucial role in our inference. Thus, modeling observational incompleteness (e.g., following Drlica-Wagner et al. 2020 and Doliva-Dolinsky et al. 2022 for the MW and M31 systems, respectively), which can mimic such nondetections, will be necessary to robustly extract a galaxy formation cutoff signal from future dwarf galaxy surveys. Upcoming facilities like Rubin will make substantial progress in this regard (Mutlu-Pakdil et al. 2021); we do not specifically consider Rubin here because careful modeling of, e.g., star–galaxy separation is needed to accurately characterize its dwarf galaxy detection sensitivity (K. Tsiane et al. 2024, in preparation).

Note that mock satellite populations in our analyses range from ≈ 100 to 200 objects per host (e.g., Nadler et al. 2020b). Given that subhalo and satellite statistics are approximately Poissonian (e.g., Mao et al. 2015), the statistical uncertainty of our measurements is $\approx 10\%$ per host. Because faint-end galaxy–halo connection uncertainties are much larger than 10%, our projected measurements are limited by galaxy–halo connection systematics rather than statistical uncertainties.

Thus, using auxiliary data to constrain galaxy–halo connection parameters can reduce uncertainties in our forecasts and

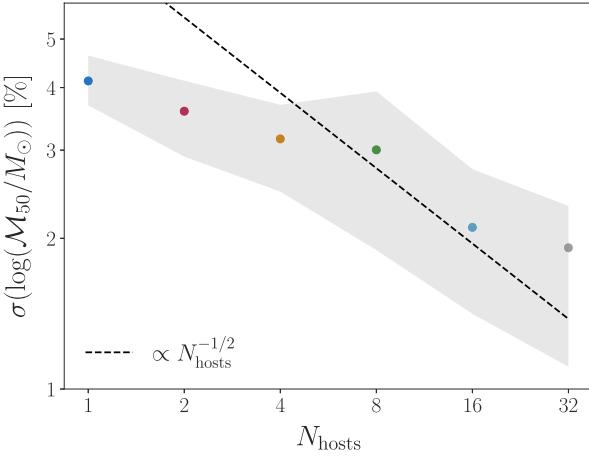


Figure 5. Projected 1σ uncertainties for statistically limited measurements of $\log(M_{50}/M_\odot)$ in the “strong” cutoff scenario, holding all remaining galaxy–halo connection parameters fixed. Points show mean measurement uncertainties as a function of the number of hosts combined in the inference for $N_{\text{hosts}} = 1$ (blue), 2 (red), 4 (orange), 8 (green), 16 (cyan), and 32 (gray); the gray band shows the 68% range of these uncertainties due to different host combinations. The black dashed line shows the $N_{\text{hosts}}^{-1/2}$ scaling expected when combining statistically limited measurements of independent hosts.

help realize the statistically limited measurements derived in Section 4.3. For example, constraints on the faint-end luminosity function slope, α , could be derived from observations of dwarf galaxies with luminosities comparable to the brightest satellites we consider here; forthcoming Roman data will detect such dwarf galaxies in a range of environments. Meanwhile, the galaxy–halo luminosity scatter, σ_M , can be probed more directly using dwarf galaxies’ inferred dynamical masses (e.g., from stellar velocity dispersion measurements; Simon et al. 2019). Finally, the efficiency of subhalo disruption due to baryons, \mathcal{B} , can be better constrained using large samples of bright satellite populations (e.g., from SAGA; Geha et al. 2017; Mao et al. 2021), and priors can be tightened with improved predictions for subhalo disruption in embedded-disk simulations (Y. Wang et al. 2024, in preparation).

5. Warm Dark Matter Forecasts

Next, we forecast measurements of the thermal-relic WDM particle mass, parameterized by the half-mode mass M_{hm} in our forecasts, using dwarf galaxy populations. We consider three scenarios, all of which include M_{hm} in the inference, but which differ in their assumptions regarding the underlying galaxy formation and DM physics as determined by the input values of M_{50} and M_{hm} :

1. Scenario A: CDM subhalo populations ($M_{\text{hm}} \rightarrow 0 M_\odot$) with a “weak” galaxy formation cutoff ($M_{50} = 3 \times 10^7 M_\odot$).
2. Scenario B: CDM subhalo populations ($M_{\text{hm}} \rightarrow 0 M_\odot$) with a “strong” galaxy formation cutoff ($M_{50} = 10^8 M_\odot$).
3. Scenario C: WDM subhalo populations ($M_{\text{hm}} = 10^8 M_\odot$, or $m_{\text{WDM}} = 4.9 \text{ keV}$) with a “strong” galaxy formation cutoff ($M_{50} = 10^8 M_\odot$).

Scenario A provides an optimistic forecast for the WDM constraining power of future dwarf galaxy surveys since it does not include an astrophysical signal that mimics the WDM cutoff. Scenarios B and C allow us to examine degeneracies between WDM and astrophysical cutoff parameters. For Scenario C, the WDM model we choose is disfavored by the

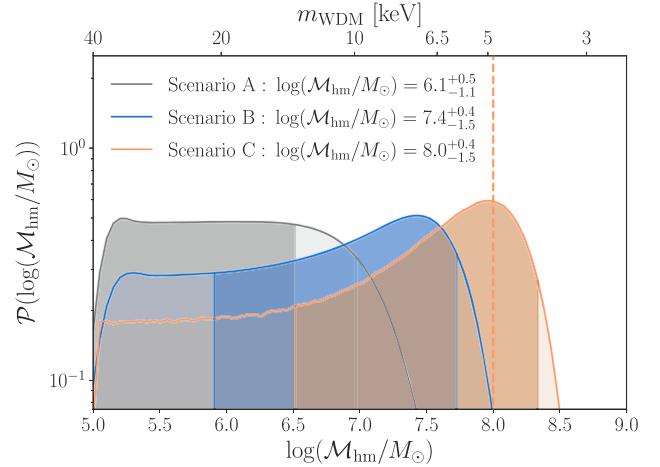


Figure 6. Marginalized posterior for the WDM half-mode mass, M_{hm} , derived from our forecasted measurements of one complete satellite population of a MW-mass system; top ticks show the corresponding values of m_{WDM} . Results are shown for forecasts that assume CDM subhalo abundances and a “weak” galaxy formation cutoff (Scenario A; gray), CDM subhalo abundances and a “strong” galaxy formation cutoff (Scenario B; blue), and WDM subhalo abundances and a “strong” galaxy formation cutoff (Scenario C; light red). Dark (light) regions show 68% (95%) confidence intervals, and the vertical dashed red line shows the value of $M_{\text{hm}} = 10^8 M_\odot$ ($m_{\text{WDM}} = 4.9 \text{ keV}$) assumed in Scenario C.

MW satellite population as well as by other small-scale structure probes like strong gravitational lensing (Gilman et al. 2020a; Hsueh et al. 2020) and the Ly α forest (Iršič et al. 2017, 2024), and combinations thereof (e.g., Nadler et al. 2021b; Enzi et al. 2021). Nevertheless, we chose this input model so that our predicted satellite populations are significantly affected by WDM, allowing us to study degeneracies between WDM and the galaxy–halo connection in detail.

5.1. WDM Particle Mass Posteriors

Figure 6 shows the marginalized posterior for M_{hm} from each of our WDM scenarios using one complete satellite population. Scenario A yields the strongest WDM constraint, with $M_{\text{hm}} < 10^{6.6} M_\odot$ ($m_{\text{WDM}} > 12.9 \text{ keV}$) from one host and $M_{\text{hm}} < 10^{6.0} M_\odot$ ($m_{\text{WDM}} > 19.5 \text{ keV}$) from two hosts at 68% confidence. In Scenario B, we find $M_{\text{hm}} < 10^{7.8} M_\odot$ ($m_{\text{WDM}} > 5.6 \text{ keV}$) from one host and $M_{\text{hm}} < 10^{6.1} M_\odot$ ($m_{\text{WDM}} > 17.0 \text{ keV}$) from two hosts. These limits are consistent with our expectation that WDM constraints are most stringent in the absence of a galaxy formation cutoff that mimics the WDM signal.

To contextualize these results, note that Nadler et al. (2020b) inferred $M_{\text{hm}} < 10^{7.3} M_\odot$ ($m_{\text{WDM}} > 7.9 \text{ keV}$) at 68% confidence from current DES and PS1 observations of the MW satellite population and simultaneously disfavor the “strong” cutoff scenario. Our forecasts are consistent with these bounds and indicate that observations of the remaining MW satellite population alone can improve WDM particle mass constraints by $\approx 60\%$ in the “weak” cutoff scenario. These projected one-host WDM constraints do not significantly change if we use MW-like hosts that contain Large Magellanic Cloud (LMC) analog subhalos; however, modeling the actual MW satellite population without accounting for LMC-associated satellites can bias inferred WDM constraints (e.g., see the discussion in Nadler et al. 2021b).

Our analysis of our full WDM forecast posteriors in Appendix C.2 reveals prominent degeneracies between M_{hm} and both \mathcal{M}_{50} and σ_M . In particular, M_{hm} and \mathcal{M}_{50} cannot simultaneously be large because they both suppress satellite abundances; meanwhile, M_{hm} and σ_M cannot simultaneously be large according to our discussion of the degeneracy between \mathcal{M}_{50} and σ_M in Section 4.1. Adding a second satellite population improves WDM constraints by both reducing statistical uncertainties and markedly improving the measurement of σ_M ; in turn, this improves the measurement of \mathcal{M}_{50} at 1σ and disfavors large values of M_{hm} . The $m_{\text{WDM}} \approx 20$ keV sensitivity of two complete hosts strongly motivates searches for remaining undiscovered satellites of the MW and M31.

In Scenario C, our M_{hm} posterior is consistent with the input value, indicating that the WDM extension of our inference framework is also self-consistent. We find that M_{hm} is only probed at the $\approx 1\sigma$ level using one host, mainly due to a significant degeneracy with \mathcal{M}_{50} . When two complete hosts are used, the nondetection of M_{hm} persists while the 1σ uncertainty on \mathcal{M}_{50} shrinks. In our model, \mathcal{M}_{50} and M_{hm} both suppress satellite abundances as a function of peak halo mass, and their observable effects only differ in the shape of the halo mass-dependent suppression imprinted on satellite populations. The nondetection of M_{hm} is consistent with our finding in Section 4.2 that the shape of the galaxy formation cutoff is difficult to measure, even when multiple hosts are combined. Note that we vary the shape of the galaxy formation cutoff but not of the WDM cutoff, which allows \mathcal{M}_{50} to be disentangled from \mathcal{S}_{gal} in the inference; it will be interesting to explore constraints on the functional dependence of SHMF suppression further in future work.

5.2. Projections for Additional Hosts

Following Section 4.3, we perform idealized tests with all galaxy–halo connection parameters held fixed to derive projections for WDM constraints with $N_{\text{hosts}} > 2$. Specifically, we calculate idealized likelihoods as a function of M_{hm} that sample over galaxy–halo connection model realizations and host combinations in the “strong” and “weak” cutoff scenarios, assuming CDM subhalo abundances. We use these to derive projected 2σ upper limits on M_{hm} , and we convert these to lower limits on m_{WDM} using Equation (4). Figure 7 shows the results of this test for the “strong” cutoff. We find that statistically limited measurements of 32 complete hosts yield $m_{\text{WDM}} \gtrsim 30$ keV at the 2σ level; different host combinations contribute a ± 10 keV uncertainty to this projection. These results do not significantly change in the “weak” cutoff scenario because galaxy–halo connection parameters are fixed; in Section 5.1, our 68% confidence WDM limits were more stringent for the “weak” versus the “strong” cutoff due to less prominent galaxy–halo connection degeneracies.

We derive projections for the scaling of WDM constraints with N_{hosts} from the dependence of the SHMF suppression on m_{WDM} . In particular, given that our mock dwarf galaxy population measurements probe the SHMF down to $M_{\text{peak}} \approx 10^8 M_{\odot}$, we use Equation (5) to predict how lower limits on m_{WDM} , denoted $m_{\text{WDM,lim}}$, depend on N_{hosts} . Assuming that uncertainties on the SHMF amplitude shrink as $N_{\text{hosts}}^{-1/2}$ and that SHMF suppression at $M_{\text{peak}} \approx 10^8 M_{\odot}$ determines the WDM constraint, we derive $m_{\text{WDM,lim}} \propto N_{\text{hosts}}^{3/20}$. Fixing the normalization of this calculation to the result of our

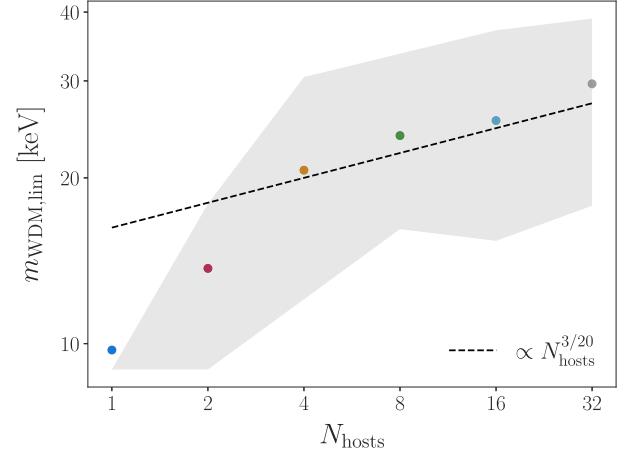


Figure 7. Projected 2σ lower limits on m_{WDM} in the “strong” cutoff scenario for statistically limited measurements of dwarf satellite populations, with all galaxy–halo connection parameters held fixed. Points show mean lower limits as a function of the number of hosts combined in the inference for $N_{\text{hosts}} = 1$ (blue), 2 (red), 4 (orange), 8 (green), 16 (cyan), and 32 (gray); the gray band shows the 68% range of these constraints due to different host combinations. The black dashed line shows the $N_{\text{hosts}}^{3/20}$ scaling expected from the functional dependence of SHMF suppression on m_{WDM} .

idealized test for $N_{\text{hosts}} = 4$, we project 2σ lower limits of

$$m_{\text{WDM,lim}} \gtrsim 20 \text{ keV} \times \left(\frac{N_{\text{hosts}}}{4} \right)^{3/20}, \quad (11)$$

which is valid for $N_{\text{hosts}} \geq 4$ and statistically limited measurements. Figure 7 confirms this scaling in our idealized tests.

The relatively shallow scaling between $m_{\text{WDM,lim}}$ and N_{hosts} in Equation (11) follows because, as m_{WDM} increases, the SHMF rapidly approaches that in CDM over the range of M_{peak} probed by satellite galaxies. Specifically, the (sub)halos that host dwarf galaxies in our forecasts span a particular range of peak virial masses from $\approx 10^8$ to $10^{10} M_{\odot}$. According to Equation (11), detecting all satellites of the several hundred MW-mass systems within ≈ 40 Mpc (e.g., Geha et al. 2017; Mao et al. 2021)—which may be enabled by future optical/infrared space telescopes such as the Habitable Worlds Observatory (for potential capabilities, see the LUVOIR report; The LUVOIR Team 2019)—would constrain models with $m_{\text{WDM}} \approx 40$ keV. Additional dwarf galaxy observables and probes of lower-mass halos below the galaxy formation threshold will be necessary to qualitatively improve this WDM sensitivity.

5.3. Interpretation

Our results indicate that future dwarf galaxy surveys can potentially place stringent constraints on WDM, comparable to the upcoming sensitivity of other small-scale structure probes (Drlica-Wagner et al. 2019). However, we emphasize that there are significant degeneracies between the effects of galaxy formation and DM physics on dwarf galaxy populations, which must be modeled accurately to for unbiased inference. Dwarf galaxy observables beyond those considered here, including star formation histories (which can be delayed in WDM-like models; e.g., Bozek et al. 2019) and stellar velocity dispersions (e.g., Kim & Peter 2021; Esteban et al. 2023) may help break these degeneracies. In parallel, combining dwarf galaxy populations with independent probes of low-mass halo abundances, such as strong lensing and stellar streams, is

Table 1
Summary of Galaxy Formation, WDM, and SHMF Forecast Results

Forecast	Parameter(s)	Input	1 Host	2 Hosts
Galaxy Formation: “weak” cutoff (Section 4)	$\log(\mathcal{M}_{50}/M_{\odot})$	7.5	$7.1^{+0.3}_{-1.6^{**}}$	$5.3^{+1.4}_{-0.2^{**}}$
Galaxy Formation: “strong” cutoff (Section 4)	$\log(\mathcal{M}_{50}/M_{\odot})$	8.0	$7.9^{+0.4}_{-1.6^{**}}$	$8.0^{+0.1}_{-0.1^{**}}$
WDM: Scenario A (Section 5)	$\log(M_{\text{hm}}/M_{\odot})$	5.0	$6.1^{+0.5}_{-1.1^{**}}$	$5.0^{+1.0}_{-0.0^{**}}$
WDM: Scenario B (Section 5)	$\log(M_{\text{hm}}/M_{\odot})$	5.0	$7.4^{+0.4}_{-1.5^{**}}$	$5.0^{+1.2}_{-0.0^{**}}$
WDM: Scenario C (Section 5)	$\log(M_{\text{hm}}/M_{\odot})$	8.0	$8.0^{+0.4}_{-1.5^{**}}$	$7.6^{+0.3}_{-1.8^{**}}$
SHMF: “weak” cutoff (Section 6)	ξ_8, ξ_9, ξ_{10}	0.0, 0.0, 0.0	$0.2^{+1.2^{**}}_{-0.3}, 0.1^{+0.5}_{-0.3}, 0.1^{+0.3}_{-0.4}$	$0.5^{+0.4}_{-0.3}, 0.3^{+0.2}_{-0.1}, 0.1^{+0.2}_{-0.2}$
SHMF: “strong” cutoff (Section 6)	ξ_8, ξ_9, ξ_{10}	0.0, 0.0, 0.0	$0.1^{+0.9}_{-1.3^{**}}, 0.1^{+0.7}_{-0.8}, 0.0^{+0.3}_{-1.2^{**}}$	$0.2^{+0.8}_{-0.3}, 0.1^{+0.2}_{-0.2}, 0.0^{+0.2}_{-0.2}$

Note. The first column describes the forecast, the second column lists the parameter(s) of interest, the third column lists parameters’ input value(s), and the fourth (fifth) columns list 68% confidence interval(s) from our one (two)-host forecasts. One asterisk (two asterisks) denote limits that are consistent with our priors at 68% (95%) confidence (see Appendix A).

another promising path forward (e.g., Banik et al. 2021a; Nadler et al. 2021b; Enzi et al. 2021).

Our MCMC forecasts’ sensitivity to WDM masses of $\approx 10\text{--}20$ keV can be understood in terms of the corresponding SHMF suppression. In particular, the left panel of Figure 2 shows $\approx 20\text{--}50\%$ suppression of the SHMF at peak virial masses of $M_{\text{peak}} \approx 10^8 M_{\odot}$ in the most extreme WDM models that are marginally consistent with our mock CDM data. This is consistent with the fact that our CDM forecasts are sensitive to a “strong” galaxy formation cutoff over the range of peak halo masses indicated by the blue band. Interestingly, our most optimistic constraint of $m_{\text{WDM}} > 17.0$ keV in Scenario B corresponds to SHMF suppression of only $\approx 10\%$, highlighting the importance of breaking additional galaxy–halo connection degeneracies in this case.

6. Subhalo Mass Function Forecasts

Finally, we forecast DM model-independent constraints on the amplitude of the SHMF relative to CDM. In particular, we forecast SHMF constraints for the “weak” and “strong” galaxy formation cutoff scenarios introduced above. In both scenarios, we assume underlying CDM subhalo abundances, corresponding to input values of $\xi_8 = 0$, $\xi_9 = 0$, and $\xi_{10} = 0$ (recall that ξ_i parameterizes the logarithmic deviation of the SHMF relative to CDM in mass decade i).

6.1. SHMF Amplitude Posteriors

In both the “strong” and “weak” cutoff scenarios, our ξ_8 , ξ_9 , and ξ_{10} posteriors are consistent with CDM, validating our inference framework when applied to generalized SHMFs. Furthermore, as summarized in Table 1, we find that SHMF suppression can be constrained to $\approx 70\%$, 60% , and 50% that in CDM at peak virial masses of $10^8 M_{\odot}$, $10^9 M_{\odot}$, and $10^{10} M_{\odot}$, respectively, for one complete host in the “weak” cutoff scenario. Meanwhile, enhanced SHMFs are only constrained relative to CDM at factors of ≈ 20 , 4, and 3 in these mass decades. Uniformly suppressed (enhanced) SHMFs corresponding to the 68% confidence limits from these “weak” cutoff forecasts are shown by the dotted red (dashed blue) lines in the right panel of Figure 2.

Our analysis of the full posteriors in Appendix C.3 reveals degeneracies between SHMF amplitude and galaxy–halo connection parameters including the faint-end luminosity

function slope α and luminosity scatter σ_M , particularly at the low-mass end ($M_{\text{peak}} \approx 10^8 M_{\odot}$). These degeneracies are expected because variations in the underlying SHMF can mimic or counteract the effects of our galaxy–halo connection model for the faintest satellite we consider. At higher masses ($M_{\text{peak}} \approx 10^{10} M_{\odot}$), we forecast that the SHMF can be measured reasonably precisely.

A complete census of satellites around two hosts improves the precision of these constraints. For example, at $M_{\text{peak}} = 10^{10} M_{\odot}$, 68% confidence limits on SHMF enhancement change from three to two times that in CDM, and limits on SHMF suppression change from 50% to 80% that in CDM; this level of improvement is similar at lower halo masses (see Table 1). Furthermore, with two complete hosts, both the galaxy formation cutoff and underlying SHMF are recovered more confidently at the 1σ level. Note that our projected SHMF enhancement constraints are not very sensitive to the assumed galaxy formation cutoff; however, SHMF suppression constraints are slightly weaker in our “strong” cutoff scenario because suppression in the underlying SHMF mimics the assumed galaxy formation cutoff. In the one-host case, these degeneracies substantially degrade the measurement of \mathcal{M}_{50} relative to our galaxy formation cutoff forecasts that assume CDM subhalo abundances in Section 4. These results underscore the importance of combining multiple satellite populations, particularly when assumptions about underlying subhalo populations are relaxed.

Figure 8 illustrates the results of our one- and two-host SHMF forecasts in the “strong” cutoff scenario relative to the SHMF from our MW-mass zoom-in simulation suite. In the $M_{\text{peak}} = 10^9 M_{\odot}$ and $10^{10} M_{\odot}$ decades, our forecasted SHMF constraints are comparable to the host-to-host scatter about the mean SHMF among our 40 zoom-in simulations (see Mao et al. 2015 for a detailed discussion of the host-to-host scatter). At $M_{\text{peak}} = 10^8 M_{\odot}$, SHMF suppression is probed at a level comparable to the host-to-host scatter, while SHMF enhancement that significantly exceeds the range of CDM predictions is compatible with the mock data.

These results imply that secondary properties beyond host halo mass that correlate with SHMF amplitude (e.g., concentration; Mao et al. 2015) can be probed using multiple satellite populations. On the other hand, accurately interpreting SHMF measurements from future dwarf galaxy surveys will require modeling these secondary host halo properties,

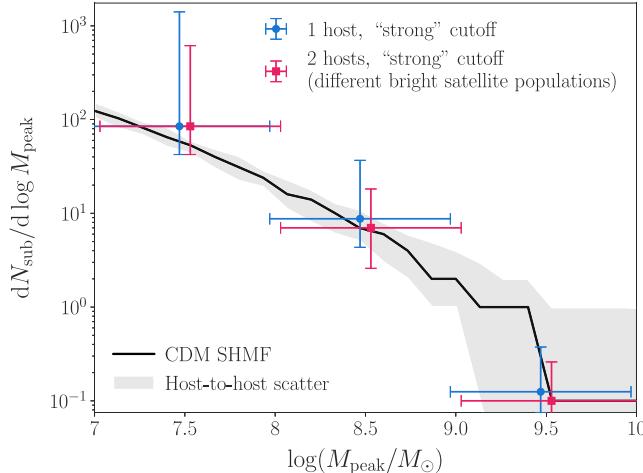


Figure 8. Mean and 16th–84th percentile host-to-host scatter of the SHMF (black line and band) from our 40 CDM MW-mass zoom-in simulations. Blue circles (red squares) and error bars show the mean and 68% confidence intervals from our SHMF forecast using one (two) complete satellite populations of MW-mass hosts, assuming a “strong” galaxy formation cutoff. For clarity, the one- and two-host error bars are offset horizontally.

particularly for models that suppress the SHMF at low masses. We emphasize that our forecasted upper limits on the SHMF are very weak at low masses; the precision of these upper limits is mainly determined by galaxy–halo connection degeneracies for two or more hosts. Thus, we expect surveys of dwarfs in a range of environments, at the bright end of the satellite populations we consider here, to reduce galaxy–halo connection uncertainties and help determine the SHMF.

6.2. Projections for Additional Hosts

We derive projected uncertainties on ξ_8 , ξ_9 , and ξ_{10} by calculating the likelihood of each parameter for fixed galaxy–halo connection parameters in the “strong” cutoff scenario. For large N_{hosts} , these projected SHMF constraints shrink according to $N_{\text{hosts}}^{-1/2}$, as expected from Poisson statistics; furthermore, they do not change significantly in the “weak” cutoff scenario. We project that statistically limited measurements of 4 (32) complete hosts yield SHMF uncertainties of $\approx 5\%$ ($\approx 2\%$) relative to CDM. Unlike our results in Section 6.1, these projected uncertainties are roughly symmetric and do not depend strongly on peak halo mass because they are not limited by galaxy–halo connection degeneracies.

6.3. Interpretation

Our results indicate that future dwarf galaxy surveys can set upper limits on the SHMF at peak subhalo masses from 10^9 to $10^{10} M_{\odot}$, while simultaneously placing a robust lower limit at $10^8 M_{\odot}$. The upper limits on the SHMF amplitude at high masses result from the fact that, if high-mass subhalos are more abundant than in the input CDM model, observed luminosity functions are overpredicted. In particular, it is difficult to “hide” the galaxies that occupy high-mass subhalos by varying galaxy–halo connection parameters without drastically reducing the abundance of lower-mass galaxies and resulting in worse agreement with the mock CDM data (for details, see our analysis of the full SHMF forecast posteriors in Appendix C.3). At lower peak halo masses, and particularly near $M_{\text{peak}} \approx 10^8 M_{\odot}$ where our fiducial galaxy formation cutoffs take effect, models with suppressed low-mass subhalo

abundances underpredict faint-end luminosity functions; however, models with enhanced low-mass subhalo abundances can be “hidden” by the galaxy formation cutoff. More precise constraints on the mass scale and shape of the cutoff would therefore improve our projected SHMF enhancement measurements.

Our SHMF forecasts can be translated to any DM or early-Universe physics scenario that suppresses or enhances the SHMF relative to CDM. This is most often realized by suppressing or enhancing the formation of (sub)halos relative to CDM; it will therefore be interesting to generalize our modeling framework to directly constrain the linear matter power spectrum, which often dictates the resulting suppression or enhancement of subhalo abundances in such scenarios (e.g., Gilman et al. 2022). Meanwhile, if the SHMF is altered at late times (e.g., via subhalo disruption due to new DM physics), our forecasts are applicable if the abundance of satellite galaxies is correspondingly altered. The connection between subhalo and satellite disruption is nontrivial and requires detailed study in the context of specific DM models (e.g., see Nadler et al. 2020a, 2021a and Mau et al. 2022 for examples in the context of self-interacting and decaying DM); we leave studies of such scenarios to future work.

Our SHMF suppression results are relevant for a variety of warm, interacting, and ultralight DM models (e.g., Bechtol et al. 2022). Meanwhile, DM models that enhance the SHMF more easily evade constraints from dwarf galaxy population statistics because their effects can be counteracted by galaxy formation physics. Such models include, e.g., ultralight axions that undergo parametric resonance in the early Universe (Arvanitaki et al. 2020; Cyncynates et al. 2022) or DM produced during a period of early matter domination (Erickcek & Sigurdson 2011). We expect that dwarf galaxy observables beyond luminosity, size, and distance (e.g., velocity dispersions; Kim & Peter 2021; Esteban et al. 2023) will help probe DM physics that enhances (sub)halo abundances. Distinct small-scale structure probes like strong gravitational lensing and stellar streams, which are more sensitive to subhalos density profiles than dwarf galaxy population statistics, will complement this effort.

7. Future Outlook

We now discuss our results in the context of near-future observational capabilities (Section 7.1), the long-term landscape of dwarf galaxy measurements (Section 7.2), and areas for future modeling work (Section 7.3).

7.1. Near-term Observational Capabilities

Our assumption that complete satellite populations of several MW-mass systems can be detected is clearly optimistic given current and near-future observational capabilities. In particular, this level of completeness will be difficult to achieve even for the MW and M31 in the next decade. Preliminary dwarf galaxy detection sensitivity estimates in Drlica-Wagner et al. (2019) indicate that Rubin is sensitive to satellites with $M_V < 0$ mag and $\mu < 32$ mag arcsec $^{-2}$ in the Southern Hemisphere, out to the MW virial radius. However, these projections may be optimistic due to the difficulties of star–galaxy separation; realistic mock simulations of dwarf recovery in Rubin data will be needed to accurately derive its sensitivity. Meanwhile, for M31, Doliva-Dolinsky et al. (2022) showed that current

PAndAS data is sensitive to systems with $M_V \lesssim -5.5$ mag and $\mu \lesssim 30$ mag arcsec $^{-2}$; for context, Mutlu-Pakdil et al. (2021) found that Rubin is sensitive to even fainter dwarf galaxies at comparable distances, although M31 is not within its field of view.

Thus, upcoming observations are unlikely to achieve complete coverage of dwarf galaxy populations down to our assumed detection limits. As a result, the “weak” cutoff scenario will likely remain undetectable in the near future. Nonetheless, we emphasize that the MW satellite population already probes the “strong” cutoff scenario (Nadler et al. 2020b), and we expect these constraints to improve with additional data. Beyond observations of MW and M31 satellites, Euclid and Roman will also detect and help confirm the nature of faint dwarf galaxies throughout the Local Volume, both by discoveries with in situ data and by providing distance measurements for dwarf candidates identified in Rubin imaging data. As Euclid has already launched, a focused study of its stand-alone and combined dwarf galaxy detection and characterization capabilities is also timely. For small areas of the sky, it has recently been demonstrated that JWST can detect very faint, low surface brightness dwarf galaxies, including at cosmological distance (e.g., Conroy et al. 2023), which will help characterize the dwarf galaxy population below the limit of ground-based photometric surveys.

An assumption underlying our forecasts is that properties of the MW-mass centrals that host the dwarf satellites we consider will be well characterized. Accurate and precise measurements of host halo masses are important for determining expected subhalo and satellite populations, and measurements of secondary properties like host concentration are also helpful (e.g., Mao et al. 2015; Fielder et al. 2019). At present, even the MW’s halo mass remains difficult to measure precisely; although, proper-motion data from Gaia have significantly reduced this uncertainty (e.g., see Gardner et al. 2021 for a review). For more distant systems, velocity dispersion measurements of globular clusters and of dwarf satellites themselves will likely provide the most precise host mass estimates. Jointly inferring host halo properties and the galaxy–halo connection for their satellites will be necessary to claim a detection of the astrophysical and DM signatures we have explored. Extending our model to marginalize over host halo properties is a natural avenue for future work. We do not expect marginalizing over host properties to change our qualitative conclusions (e.g., the value of combining hosts with different bright satellite populations to break galaxy–halo connection degeneracies); however, our quantitative results concerning the number of hosts needed to achieve given galaxy formation cutoff, WDM, or SHMF constraints will depend on these uncertainties in detail.

We note that our forecasts adopt broad priors on the galaxy–halo connection that are likely to be improved using measurements of brighter dwarf galaxies. As noted above, adopting tighter priors on, e.g., the luminosity function slope and galaxy–halo connection scatter shrinks our systematic error budget and therefore directly improves the precision of our galaxy formation cutoff, WDM, and SHMF constraints. Current surveys like ELVES (Carlsten et al. 2022) and SAGA (Geha et al. 2017; Mao et al. 2021) have already detected large numbers of dwarf satellites throughout the Local Volume and out to ≈ 40 Mpc, and DESI is supplementing these efforts over a very large area (Darragh-Ford et al. 2023). Combining our

framework with data from these surveys will improve galaxy–halo connection constraints (e.g., Danieli et al. 2023). Furthermore, upcoming facilities including Roman will likely detect large numbers of dwarfs in both the (partially) resolved regimes and in integrated light with distances measured via, e.g., surface brightness fluctuations (Greco et al. 2021). In parallel, ongoing neutral hydrogen surveys like WALLABY are expected to reach HI detection thresholds of $\approx 10^5 M_\odot$ (Westmeier et al. 2022), probing the gas content of low-mass dwarfs and providing a complementary handle on the faint-end galaxy–halo connection.

7.2. Long-term Observational Outlook

In the longer term, next-generation optical/infrared space telescopes have the potential to revolutionize our ability to detect faint dwarf galaxies. For example, The LUVOIR Team (2019) estimated that a 15 m optical/infrared space telescope can resolve individual stars out to tens of megaparsecs and potentially detect dwarf galaxies with $M_V < 0$ mag at distances of ~ 15 Mpc. Because of the relatively small field of view of such telescope concepts, targeting satellites around known hosts at these distances will be crucial to maximize dwarf galaxy discovery. Such detection sensitivity easily satisfies the assumptions of our forecasts; for example, there are ≈ 30 MW-mass systems to survey within 12 Mpc (e.g., Carlsten et al. 2022). As the community develops the detailed science requirements for the Astro2020-prioritized Habitable Worlds Observatory, the imaging and spectroscopic detection and characterization of faint dwarf galaxies, through both diffuse light and direct stellar-cluster or stellar identifications, is a capability that should be explored. This research will be a key step toward the galaxy formation and DM physics measurements we have forecasted.

Leveraging complementary dwarf galaxy observables will also help identify signatures of a galaxy formation cutoff and new DM physics. For example, stellar velocity dispersion measurements of nearby dwarfs probe their present-day dynamical halo masses. These measurements are necessary to disambiguate dwarf galaxies from star clusters, which is becoming increasingly challenging and may represent a limiting systematic uncertainty of satellite population measurements (e.g., Cerny et al. 2023). At the same time, they provide valuable information about the galaxy–halo connection (e.g., Simon et al. 2019) and probe DM physics that affects (sub) halos’ inner density profiles, including WDM-like cutoffs (e.g., Kim & Peter 2021) or DM self-interactions (e.g., see Tulin & Yu 2018 and Adhikari et al. 2022 for reviews). Excitingly, next-generation spectroscopic facilities will pursue such measurements (e.g., Chakrabarti et al. 2022). In parallel, improved constraints on dwarf galaxies’ star formation histories from space telescopes like Roman will help determine the detailed impact of reionization on the faint end of the galaxy–halo connection (e.g., Weisz et al. 2014a, 2014b), providing important complementary constraints on the galaxy formation cutoff.

7.3. Modeling Work

Further modeling work will improve our ability to extract galaxy formation and DM physics from future dwarf galaxy surveys. Most importantly, including a broader range of host masses will yield mock dwarf galaxy populations that better

span those probed by future surveys. For example, Symphony’s LMC-mass suite contains hosts with $M_{\text{host}} \approx 10^{11} M_{\odot}$, and its Group suite contains hosts with $M_{\text{host}} \approx 10^{13} M_{\odot}$. Incorporating these suites will require recalibrating aspects of our galaxy–halo connection framework that are currently tuned to MW-mass systems, including our subhalo disruption model.

In addition to the effects of host halo mass, optimally characterizing the correlated information contained in satellite population statistics and secondary host properties remains challenging. Cosmological zoom-in simulations indicate that the strength of these correlations peaks at the MW host halo mass scale (Nadler et al. 2023). Thus, accurate measurements of the masses and secondary properties of MW-mass hosts can significantly reduce uncertainties on their underlying SHMFs. To achieve this, the following theoretical uncertainties should be addressed:

1. Even in a dark-matter-only setting, the physical origin of correlations between secondary host properties and SHMF amplitude remains unclear. Nadler et al. (2023) showed that the mass of the largest subhalo correlates strongly with SHMF amplitude in simulations; there is a similar trend in SAGA data, which shows that the luminosity of the brightest observed satellite correlates with satellite luminosity function amplitude more strongly than central galaxy properties (Mao et al. 2021). For the MW, this correlation partially arises because the LMC brings its own satellites into the MW (e.g., Kallivayalil et al. 2018; Patel et al. 2020). In fact, Nadler et al. (2020b) showed that it is difficult to match the spatial anisotropy in the observed MW satellite population without modeling the LMC system. Because the LMC accreted recently and is near pericenter today, this effect is likely transient (e.g., D’Souza & Bell 2021; Barry et al. 2023). Understanding the origin of the correlation between the mass of the largest subhalo and SHMF amplitude (or the luminosity of the brightest satellite and satellite luminosity function amplitude) would therefore be fruitful. This is particularly relevant because, as shown in Section 4, combining hosts with differing bright satellite abundances can markedly improve galaxy formation constraints.
2. Beyond dark-matter-only predictions, the presence of a central galaxy has been claimed to affect surviving subhalo and satellite abundances at the factor of ≈ 2 level (e.g., Garrison-Kimmel et al. 2017b; Kelley et al. 2019; Richings et al. 2020). The strength of this effect is debated (e.g., Webb & Bovy 2020; Green et al. 2022) and may be related to artificial disruption in cosmological simulations more generally (e.g., van den Bosch & Ogiya 2018; van den Bosch et al. 2018; Mansfield et al. 2023). Current hydrodynamic simulations indicate that subhalo disruption is not highly sensitive to central galaxy properties beyond total mass. Because satellite populations will be measured around a variety of central galaxies in the near future, improved theoretical predictions for subhalo and satellite disruption around central galaxies with a range of properties will be helpful. Combining isolated and satellite dwarf populations will also reduce uncertainties associated with subhalo and satellite disruption, which are conservatively marginalized over in our forecasts.

Although our simulations capture a representative range of MW-mass host halos and formation histories (Mao et al. 2015; Nadler et al. 2023), we have not explicitly modeled the impact of host mass and cosmic environment on the galaxy–halo connection itself. In particular, our galaxy–halo connection model populates subhalos according to the same procedure in all MW-mass hosts we consider. This assumption likely does not capture the complexity of real dwarf galaxy populations, in which galaxy formation and evolution are environmentally dependent (e.g., Christensen et al. 2024; also see Danieli et al. 2023 for discussion in the context of Local Volume satellite populations), particularly for low-mass galaxies affected by photoionization (e.g., Benson et al. 2003). Modeling environmental effects is particularly relevant because the large-scale environment of the Local Volume may be unusual (e.g., Neuzil et al. 2020; McAlpine et al. 2022), thereby impacting the expected number of nearby dwarf galaxies and the effects of photoionization on these systems (e.g., Benson et al. 2002; Busha et al. 2010; Li et al. 2014; Dixon et al. 2018).

We plan to incorporate host halo property and environmental dependence in our galaxy–halo connection framework in future work. In a similar spirit, applying our inference framework to dwarf galaxy populations predicted by hydrodynamic simulations and SAMs, including in beyond-CDM scenarios—rather than mock data generated by applying our empirical galaxy–halo connection model to CDM simulations—would be a useful test of our framework’s flexibility.

8. Prospects for Detecting “Dark” Halos

Our forecasts indicate that future dwarf galaxy population measurements can probe galaxy formation cutoffs that manifest at peak virial masses of $\approx 10^8 M_{\odot}$. By sampling from our forecasts’ posteriors, we confirm that subhalos with peak virial masses of $\approx 10^8 M_{\odot}$ host observable satellites in our mock observations. Although slightly lower-mass halos may host galaxies in our “strong” and “weak” cutoff scenarios, it is difficult to unambiguously classify such systems as dark matter-dominated galaxies since they are predicted to form fewer than 100 stars, on average (e.g., see Smith et al. 2023 for a recent example of such a system).

Thus, we assume that a peak subhalo virial mass of $3 \times 10^7 M_{\odot}$ —corresponding to the minimum observable halo mass in our “weak” cutoff scenario—sets an operational limit below which halos are effectively “dark.” Note that the present-day masses of surviving subhalos in cosmological simulations are reduced by $\approx 50\%$ or more on average relative to their peak masses (e.g., Nadler et al. 2023), although particle-tracking models (Mansfield et al. 2023) and idealized simulations (Errani & Navarro 2021) indicate that CDM subhalos often survive to much lower masses. Meanwhile, the present-day masses of isolated halos are roughly equal to their peak masses.

This peak mass threshold of $3 \times 10^7 M_{\odot}$ provides a benchmark for the detection of dark (sub)halos using gravitational DM probes. For example, Gilman et al. (2019) predicted that subhalos with present-day masses down to $\approx 10^7 M_{\odot}$ can be probed using flux ratio statistics for a sample of 50 strong gravitational lenses. Facilities including Rubin and Roman are expected to identify thousands of strong lenses in the coming years (e.g., Oguri & Marshall 2010; Collett 2015). It should be feasible to obtain the deeper follow-up data necessary for DM substructure studies for hundreds of these systems over the next decade, following efforts that are already underway using

JWST (Nierenberg et al. 2024). In this context, our forecasts indicate that a statistical detection of dark line-of-sight halos with masses of $\approx 3 \times 10^7 M_\odot$, may be achieved by combining upcoming dwarf galaxy and strong-lensing data. Meanwhile, an unambiguous statistical detection of dark subhalos will likely require sensitivity to substructure with stripped masses below $\approx 10^7 M_\odot$, although it will still be necessary to model the line-of-sight contribution when interpreting lensing substructure (e.g., Despali et al. 2018; Sengül et al. 2022).

Measuring the gravitational effects of low-mass (sub)halos at an individual level is more challenging. Current gravitational imaging analyses have identified halos with masses of $\approx 10^9 M_\odot$ (Vegetti et al. 2010, 2012; Hezaveh et al. 2016); due to its excellent angular resolution, very long baseline interferometric (VLBI) observations using the MIT-Green Bank Very Large Array are beginning to probe even lower-mass (sub)halos (e.g., Powell et al. 2022). It is therefore plausible that individual (sub)halos below the galaxy formation threshold will be detected in future VLBI data, e.g., from the Next-Generation Very Large Array (Selina et al. 2018; Kadler et al. 2023). Intriguingly, recent JWST observations indicate that highly magnified and strongly lensed stars can also potentially probe extremely low-mass DM substructure (Diego et al. 2023).

Low-mass subhalos can also be detected via their gravitational effects on stellar streams. For example, current analyses suggest that the GD-1 stream was impacted by an object with mass below $\approx 10^8 M_\odot$ (Bonaca et al. 2019) and that power spectra of stream density fluctuations are sensitive to subhalo populations of similar masses (e.g., Banik et al. 2021a). Rubin is expected to increase the precision of such measurements, potentially reaching a perturber mass sensitivity of $\approx 10^6 M_\odot$ (Drlica-Wagner et al. 2019). Upcoming stream measurements should provide an independent constraint on the abundance and spatial distribution of low-mass MW subhalos within the next decade.

The gravitational DM probes discussed above are sensitive to both halo abundances and density profiles (e.g., Bonaca et al. 2019; Gilman et al. 2020b), whereas the dwarf population statistics we have focused on mainly probe the galaxy formation cutoff and SHMF. As a result, simultaneously inferring the galaxy formation cutoff, the abundance of (sub)halos near and below the galaxy formation threshold, and the density profiles of these low-mass objects will maximize the constraining power of future searches for dark halos in the context of models beyond CDM. Note that specific DM models often affect halo abundances and density profiles in a correlated fashion; for example, suppression of the linear matter power spectrum imprinted by WDM prevents the formation of low-mass halos and simultaneously reduces the concentrations of higher-mass halos. Measurements of low-mass halo populations near and below the galaxy formation threshold will therefore contain rich, multidimensional information about galaxy formation and DM physics.

9. Summary

We have forecasted the galaxy formation and DM constraints that future surveys of dwarf galaxy populations around MW-mass hosts can deliver. We find the following:

1. A galaxy formation cutoff at peak virial subhalo masses of $\approx 10^8 M_\odot$ can be constrained by a survey of all dwarf

satellite galaxies around one MW-mass system; adding another host improves this sensitivity at the 1σ level (Figure 3).

2. A weaker galaxy formation cutoff at peak virial subhalo masses of $\approx 3 \times 10^7 M_\odot$ can be constrained from above, but is difficult to detect even with observations of multiple complete satellite populations. Gravitational probes of low-mass halos will be needed to help detect the cutoff in this scenario.
3. Combining observations of multiple complete satellite populations can dramatically improve constraints on galaxy–halo connection scatter, particularly when combining hosts with differing abundances of bright satellites in the presence of a “strong” galaxy formation cutoff (Figure 4).
4. Complete measurements of one (two) satellite populations can yield lower limits on the WDM particle mass of ≈ 10 keV (≈ 20 keV), assuming CDM subhalo abundances. However, it is difficult to disentangle WDM and galaxy formation cutoff signals using dwarf galaxy populations alone (Figure 6).
5. Future dwarf galaxy surveys can measure the SHMF from peak virial masses of 10^8 to $10^{10} M_\odot$; SHMF suppression can be constrained to $\approx 50\%$ that in CDM, while enhancement can only be constrained at the factor of ≈ 3 level times that in CDM due to galaxy–halo connection degeneracies (Figure 8).

These results indicate that next-generation dwarf galaxy surveys will probe unexplored galaxy formation and DM physics parameter space. Thus, searches for faint dwarf galaxies throughout the Local Volume and low-redshift Universe are timely. In parallel, developing simulation and galaxy–halo connection frameworks that accurately and flexibly model the specific dwarf galaxy populations observed by future surveys will critically enable a detection of new physics using the faintest galaxies.

Acknowledgments

Our galaxy–halo connection framework is publicly available at github.com/eonadler/subhalo_satellite_connection.

We are grateful to the referee for providing constructive feedback. We thank Keith Bechtol and Alex Drlica-Wagner for comments on the manuscript and Niusha Ahvazi, Rui An, Shany Danieli, Ivan Esteban, and Josh Simon for helpful discussions related to this work.

V.G. acknowledges support from NASA through the Astrophysics Theory Program, award No. 21-ATP21-0135 and from the National Science Foundation (NSF) CAREER grant No. PHY-2239205. The work of LAM was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with NASA. This work received additional support from the U.S. Department of Energy under contract No. DE-AC02-76SF00515 to SLAC National Accelerator Laboratory and from the Kavli Institute for Particle Astrophysics and Cosmology at Stanford University. This work was performed in part at the Aspen Center for Physics, which is supported by NSF grant PHY-2210452.

The computations presented here were conducted through Carnegie’s partnership with the Resnick High Performance Computing Center, a facility supported by the Resnick Sustainability Institute at the California Institute of

Technology. This work used data from the Symphony suite of simulations (<http://web.stanford.edu/group/gfc/symphony/>), which was supported by the Kavli Institute for Particle Astrophysics and Cosmology at Stanford University and SLAC National Accelerator Laboratory, and by the U.S. Department of Energy under contract No. DE-AC02-76SF00515 to SLAC National Accelerator Laboratory.

This research used <https://arXiv.org> and NASA's Astrophysics Data System for bibliographic information.

Software: ChainConsumer (Hinton 2016), emcee (Foreman-Mackey et al. 2013), getdist (Lewis 2019), Jupyter (jupyter.org), Matplotlib (Hunter 2007), NumPy (van der Walt et al. 2011), Scikit-Learn (Pedregosa et al. 2011), SciPy (Jones et al. 2001), Seaborn (seaborn.pydata.org).

Appendix A Galaxy–Halo Connection and Beyond-CDM Priors

Table 2 lists the priors distributions for our galaxy–halo connection and beyond-CDM forecasts. The first eight rows list our galaxy–halo connection priors (Section 2.2), the ninth row lists our WDM prior (Section 2.3.1), and the last three rows list our SHMF amplitude priors (Section 2.3.2). Note that certain subsets of these parameters are varied in each of our forecasts. In particular, our galaxy formation cutoff forecasts vary the eight galaxy–halo connection parameters (Section 4), our WDM forecasts vary the eight galaxy–halo connection parameters plus M_{hm} (Section 5), and our SHMF forecasts vary the eight galaxy–halo connection parameters plus ξ_8 , ξ_9 , and ξ_{10} (Section 6).

Table 2
Prior Distributions for the Parameters Varied in Our Galaxy Formation, WDM, and SHMF Forecasts

Free Parameter	Prior Distribution	Motivation
Faint-end slope	$\arctan \alpha \sim \text{unif}(-1.1, -0.9)$	Jeffreys prior over wide range of $-2 < \alpha < -1.2$
Luminosity scatter	$\sigma_M \sim \text{unif}(0, 2) \text{ dex}$	Conservative upper limit (Garrison-Kimmel et al. 2017a; Lehmann et al. 2017)
50% occupation mass	$\log(M_{50}/M_{\odot}) \sim \text{unif}(5, 10)$	Brackets sims. & SAMs (Munshi et al. 2021; Kravtsov & Manwadkar 2022)
Disruption efficiency	$\ln(\mathcal{B}) \sim \mathcal{N}(\mu = 1, \sigma = 0.5)$	Centered on simulations (Garrison-Kimmel et al. 2017b; Nadler et al. 2018)
Occupation shape	$\mathcal{S}_{\text{gal}} \sim \text{unif}(0, 1)$	Brackets sims. & SAMs (Munshi et al. 2021; Kravtsov & Manwadkar 2022)
Size amplitude	$\mathcal{A} \sim \text{unif}(0, 0.5) \text{ kpc}$	Brackets empirical galaxy–halo size relation (Kravtsov 2013)
Size scatter	$\sigma_{\log R} \sim \text{unif}(0, 2) \text{ dex}$	Brackets empirical galaxy–halo size relation (Kravtsov 2013)
Size power-law index	$n \sim \mathcal{N}(\mu = 1, \sigma = 0.5)$	Centered on empirical galaxy–halo size relation (Kravtsov 2013)
WDM half-mode mass	$\log(M_{\text{hm}}/M_{\odot}) \sim \text{unif}(5, 10)$	Brackets current constraints (Gilman et al. 2020a; Banik et al. 2021b)
SHMF deviation at $10^8 M_{\odot}$	$\xi_8 \sim \text{unif}(-2, 2)$	Brackets current constraints (Banik et al. 2021a; Nadler et al. 2021b)
SHMF deviation at $10^9 M_{\odot}$	$\xi_9 \sim \text{unif}(-2, 2)$	Brackets current constraints (Banik et al. 2021a; Nadler et al. 2021b)
SHMF deviation at $10^{10} M_{\odot}$	$\xi_{10} \sim \text{unif}(-2, 2)$	Brackets current constraints (Banik et al. 2021a; Nadler et al. 2021b)

Note. Here $\mathcal{N}(\mu, \sigma)$ denotes a normal distribution with mean μ and standard deviation σ .

Appendix B Idealized \mathcal{M}_{50} Likelihoods

Here, we study \mathcal{M}_{50} likelihoods in our idealized tests that fix all remaining galaxy–halo connection parameters and vary the number of hosts combined in the inference. Figure 9 shows \mathcal{M}_{50} likelihoods for $N_{\text{hosts}} \in [1, 2, 4, 8, 16, 32]$ in “weak” (left), “strong” (middle), and “very strong” ($\mathcal{M}_{50} = 10^9 M_{\odot}$; right) cutoff scenarios. Below, we analyze the likelihoods working from the strongest to weakest signal. In all panels, we use a fixed combination of hosts for each choice of N_{hosts} , and thus do not illustrate the combination-to-combination scatter shown in Figure 5.

In the “very strong” cutoff scenario, \mathcal{M}_{50} is sharply constrained from below because models with $\mathcal{M}_{50} \ll 10^9 M_{\odot}$ predict more satellites than the input model at a level that is significant even when using one or two hosts. \mathcal{M}_{50} is less well constrained from above, particularly when using one or two hosts, because the input model only yields a handful of detectable satellites per hosts; thus, models with $\mathcal{M}_{50} \gg 10^9 M_{\odot}$ that (on average) yield zero detectable

satellites are only constrained when enough hosts are combined such that Poisson uncertainties are small.

In the “strong” cutoff scenario, \mathcal{M}_{50} is sharply constrained from above because models with $\mathcal{M}_{50} \gg 10^8 M_{\odot}$ predict significantly fewer satellites than the input model, even when using one or two hosts. However, the likelihood has a prominent tail toward low \mathcal{M}_{50} that does not diminish as N_{hosts} increases. This is consistent with the results of our full MCMC forecasts in Section 4 and follows because models with $\mathcal{M}_{50} \lesssim 10^8 M_{\odot}$ do not predict significantly more satellites than the “strong” cutoff model. As N_{hosts} increases and Poisson uncertainties shrink, the upper limit on \mathcal{M}_{50} becomes more stringent, but the tail toward low \mathcal{M}_{50} persists.

Finally, \mathcal{M}_{50} likelihoods in the “weak” cutoff scenario are qualitatively similar to the “strong” cutoff results but display a weaker preference for the input model. This is again consistent with our forecasts in Section 4 and follows because the “weak” cutoff suppresses an even smaller fraction of the mock satellite populations in our inference, given our fiducial detection thresholds.

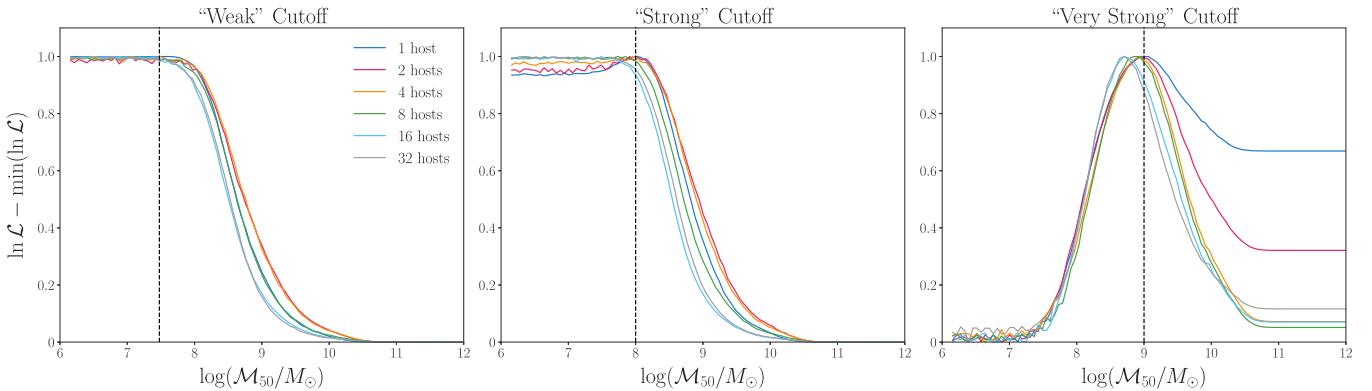


Figure 9. Likelihoods as a function of $\log(\mathcal{M}_{50}/M_{\odot})$, from our idealized tests that hold all remaining galaxy–halo connection parameters fixed at their input values, in the “weak” (left panel), “strong” (middle panel), and “very strong” (right panel) cutoff scenarios. Each panel shows the likelihood for $N_{\text{hosts}} = 1$ (blue), 2 (red), 4 (orange), and 8 (green), 16 (cyan), and 32 (gray), for a fixed combination of hosts in each case. Each likelihood is normalized to its maximum value over the plotted range.

Appendix C

Posterior Distributions

Finally, we present the full posteriors from each forecast described in the main text. For all posteriors below, dark (light) contours represent 68% (95%) confidence intervals, and dashed black lines show the input value of each parameter; contours are colored blue (red) for our one-host (two-host) forecasts. Note that σ_M and $\sigma_{\log R}$ are reported in dex, \mathcal{M}_{50} is reported as $\log(\mathcal{M}_{50}/M_{\odot})$, \mathcal{A} is reported in parsecs, M_{hm} is reported as $\log(M_{\text{hm}}/M_{\odot})$, and the remaining galaxy–halo connection and beyond-CDM parameters are dimensionless. Furthermore, the two-host posteriors below always combine hosts with different bright satellite populations (i.e., classical satellite abundances that differ at a level comparable to the 1σ host-to-host scatter in our simulation suite).

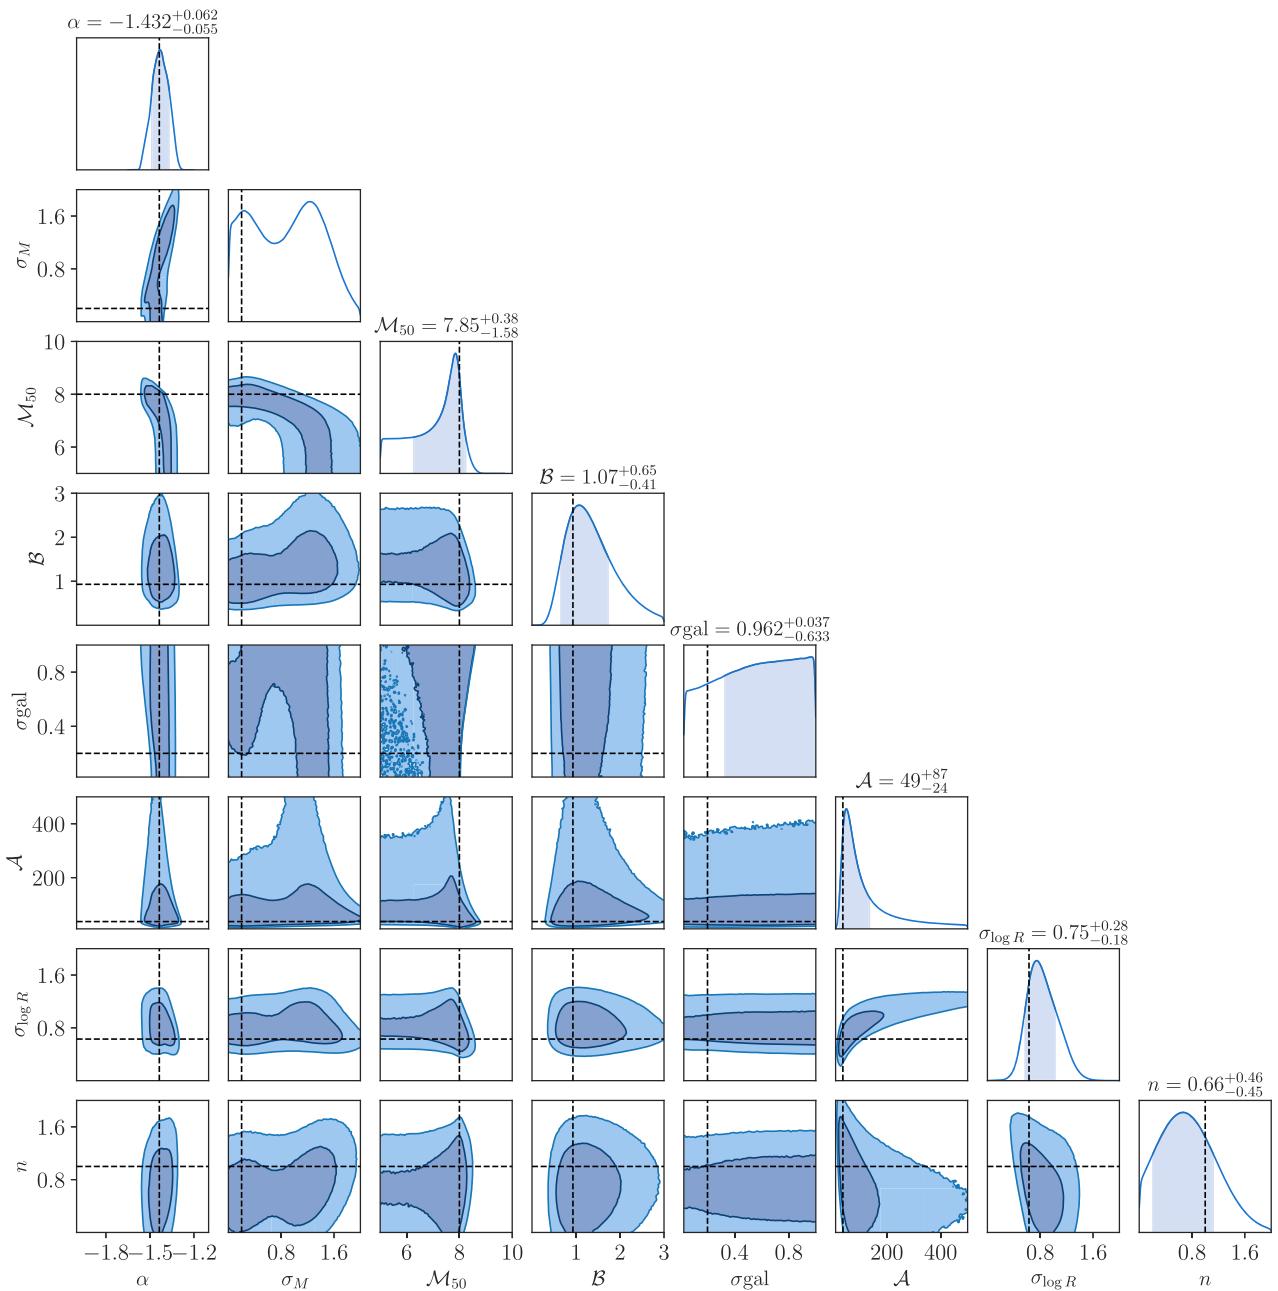


Figure 10. Posterior distribution for our “strong” galaxy formation cutoff forecast using one complete CDM satellite population.

C.1. Galaxy Formation Cutoff

The posteriors from our one- and two-host “strong” galaxy formation cutoff forecasts are shown in Figures 10 and 11, respectively. The parameters of interest for this forecast (i.e., σ_M , \mathcal{M}_{50} , and \mathcal{S}_{gal}) are not strongly correlated with the other galaxy–halo connection parameters; however, there is a prominent degeneracy between \mathcal{M}_{50} and σ_M (and, to a lesser extent, between \mathcal{S}_{gal} and σ_M). As described in Section 4, this follows because larger values of σ_M preferentially cause faint galaxies to up-scatter to observable luminosities, decreasing the average masses of halos inferred to host the faintest observed galaxies. Increasing \mathcal{M}_{50} cuts off galaxy formation more drastically in lower-mass halos, which is not allowed in regions of parameter space where the galaxies they host are crucial to explain the data, resulting in a negatively sloped degeneracy between \mathcal{M}_{50} and σ_M .

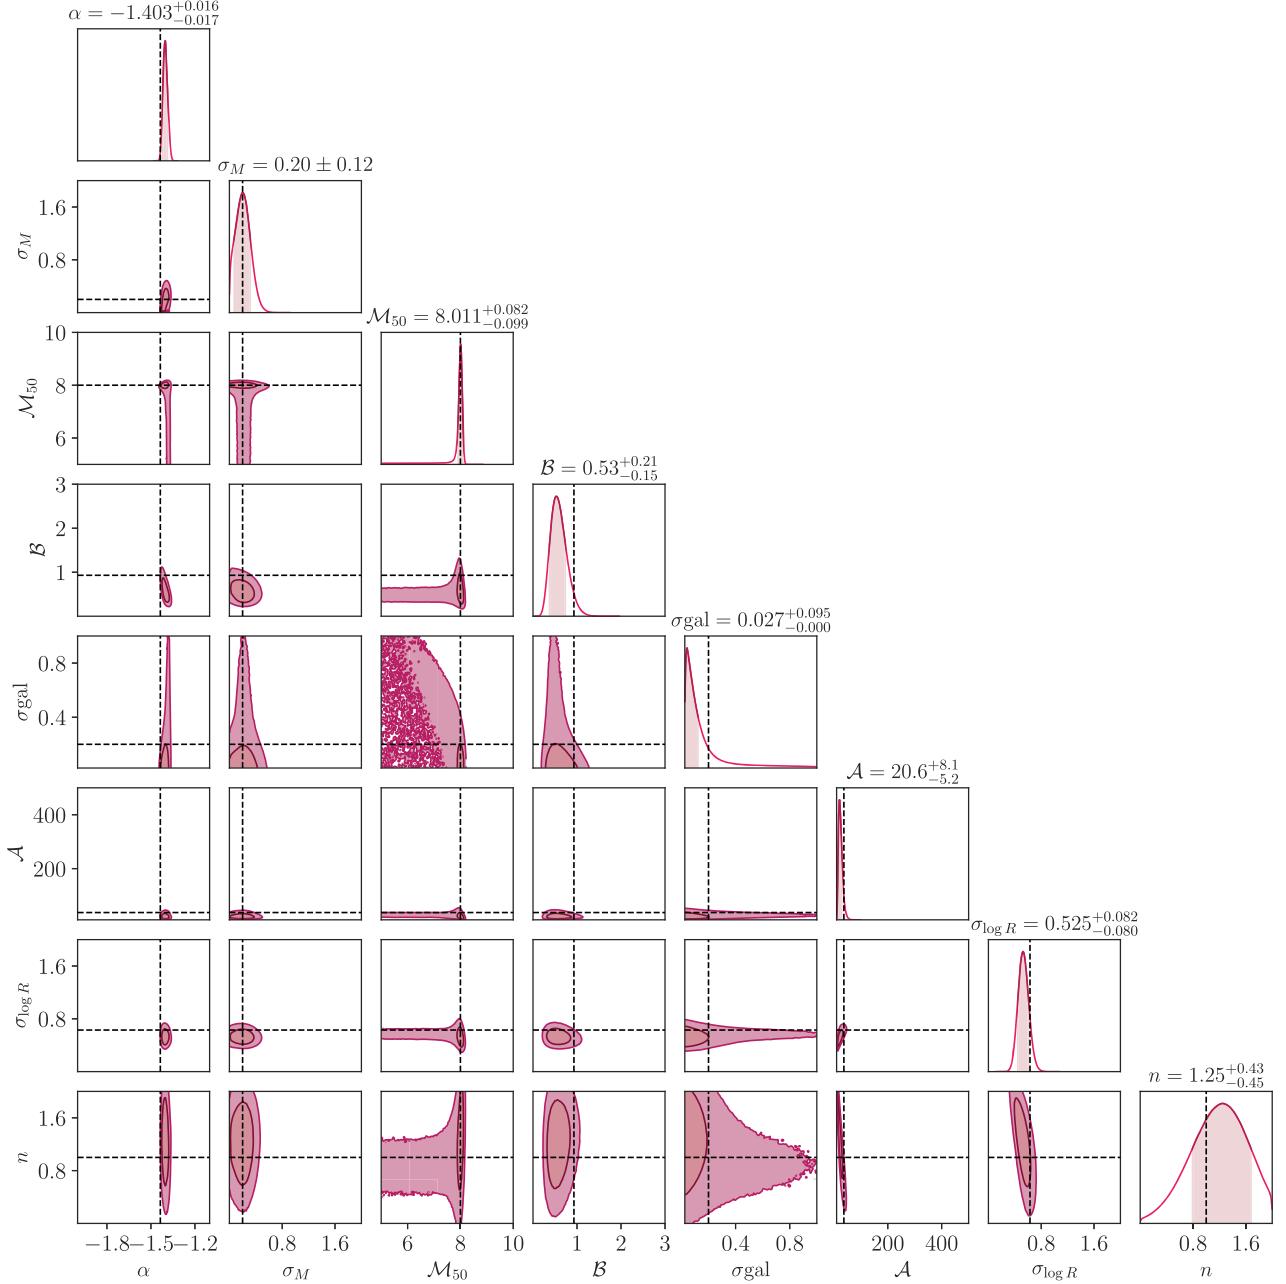


Figure 11. Posterior distribution for our “strong” galaxy formation cutoff forecast using two complete CDM satellite populations (compare to Figure 10).

Next, the posteriors from our one- and two-host “weak” galaxy formation cutoff forecasts are shown in Figures 12 and 13, respectively. In the one-host case, we only recover an upper limit on \mathcal{M}_{50} , and we observe weaker degeneracies between \mathcal{M}_{50} and the remaining galaxy–halo connection parameters. Although uncertainties significantly shrink in the two-host forecast, \mathcal{M}_{50} is still not constrained from below in this case, which is consistent with our discussion in Section 4. Note that the upper limit on \mathcal{M}_{50} in the two-host forecast is not in significant tension with its input value.

In all of these forecasts, we also observe degeneracies among the galaxy–halo connection size parameters (and particularly between \mathcal{A} and $\sigma_{\log R}$) that are consistent with those reported in Nadler et al. (2020b). For example, larger values of \mathcal{A} increase the average sizes of galaxies at all luminosities, pushing some

faint systems below our assumed surface brightness detectability threshold; this can be counteracted by increasing the size scatter, which preferentially causes small systems hosted by abundant, low-mass halos to up-scatter above our fiducial $r_{1/2} > 10$ pc cut and become detectable.

C.2. Warm Dark Matter

The posteriors from our one- and two-host WDM forecasts assuming a “weak” galaxy formation cutoff and CDM subhalo abundances are shown in Figures 14 and 15, and the corresponding “strong” galaxy formation cutoff runs are shown in Figures 16 and 17, respectively. The posteriors are qualitatively similar in these scenarios, and our description below applies to both.

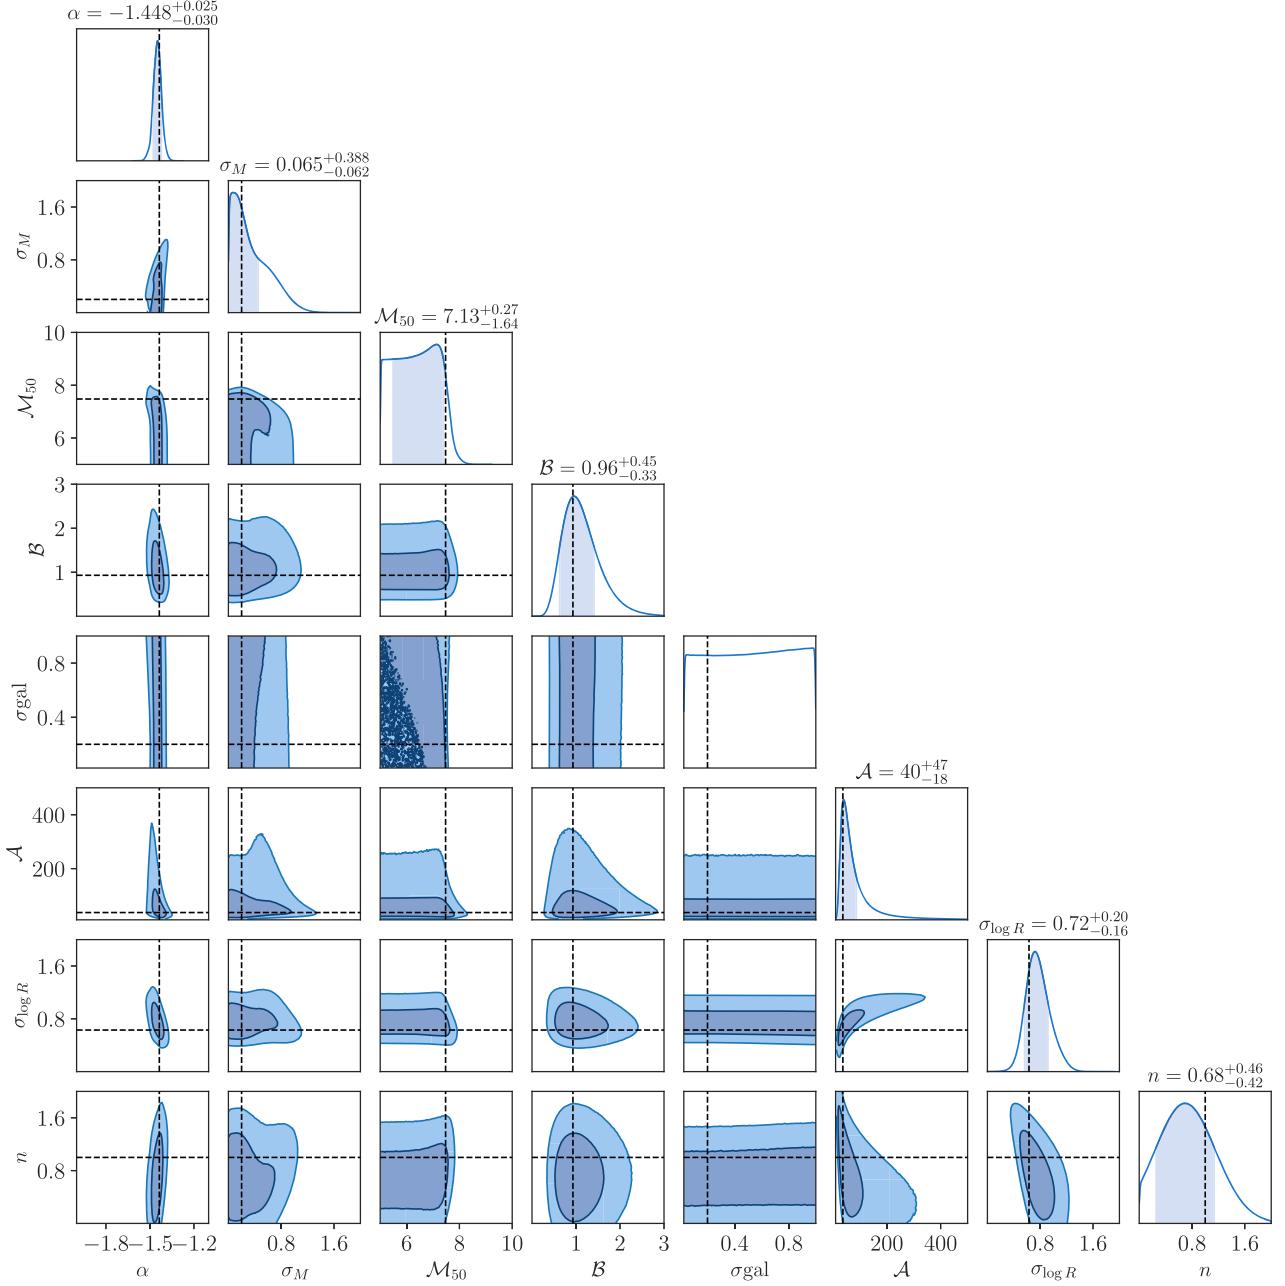


Figure 12. Posterior distribution for our “weak” galaxy formation cutoff forecast using one complete CDM satellite population.

All of the same degeneracies among galaxy–halo connection parameters noted in Appendix C.1 remain when M_{hm} is added to the inference. In addition to the significant degeneracies between M_{hm} and both \mathcal{M}_{50} and σ_M discussed in Section 5, these posteriors also illustrate a weak degeneracy between M_{hm} and \mathcal{S}_{gal} that persists even when two complete satellite populations with different bright satellite populations are used. Figures 16 and 17 clearly illustrate that the \mathcal{M}_{50} – σ_M degeneracy is broken when combining two complete satellite populations with differing classical satellite abundances, which in turn yields a much stronger upper limit on M_{hm} in the two-host case.

Figures 18 and 19, respectively, show the posteriors from our one- and two-host WDM forecasts assuming a “strong” galaxy formation cutoff and WDM subhalo abundances with

$M_{\text{hm}} = 10^8 M_{\odot}$ ($m_{\text{WDM}} = 4.9$ keV). The degeneracies described in the previous paragraph all remain in this case and often become more prominent. As discussed in Section 5, this follows because, in our model, astrophysical and DM-induced cutoffs affect dwarf galaxy abundances in a qualitatively similar fashion. Specifically, the only difference between the effects of these cutoffs in our inference is the shape of the suppression in dwarf galaxy abundances as a function of M_{peak} , dictated by the astrophysical and WDM cutoffs in Equations (2) and (5), respectively (see Figure 2).

The shape of the WDM SHMF cutoff has been studied in simulations (e.g., Angulo et al. 2013; Lovell et al. 2014; Bose et al. 2016; Stücker et al. 2022), and the shape of the astrophysical cutoff in galaxy abundances has been explored in SAMs (e.g., Benitez-Llambay & Frenk 2020; Kravtsov &

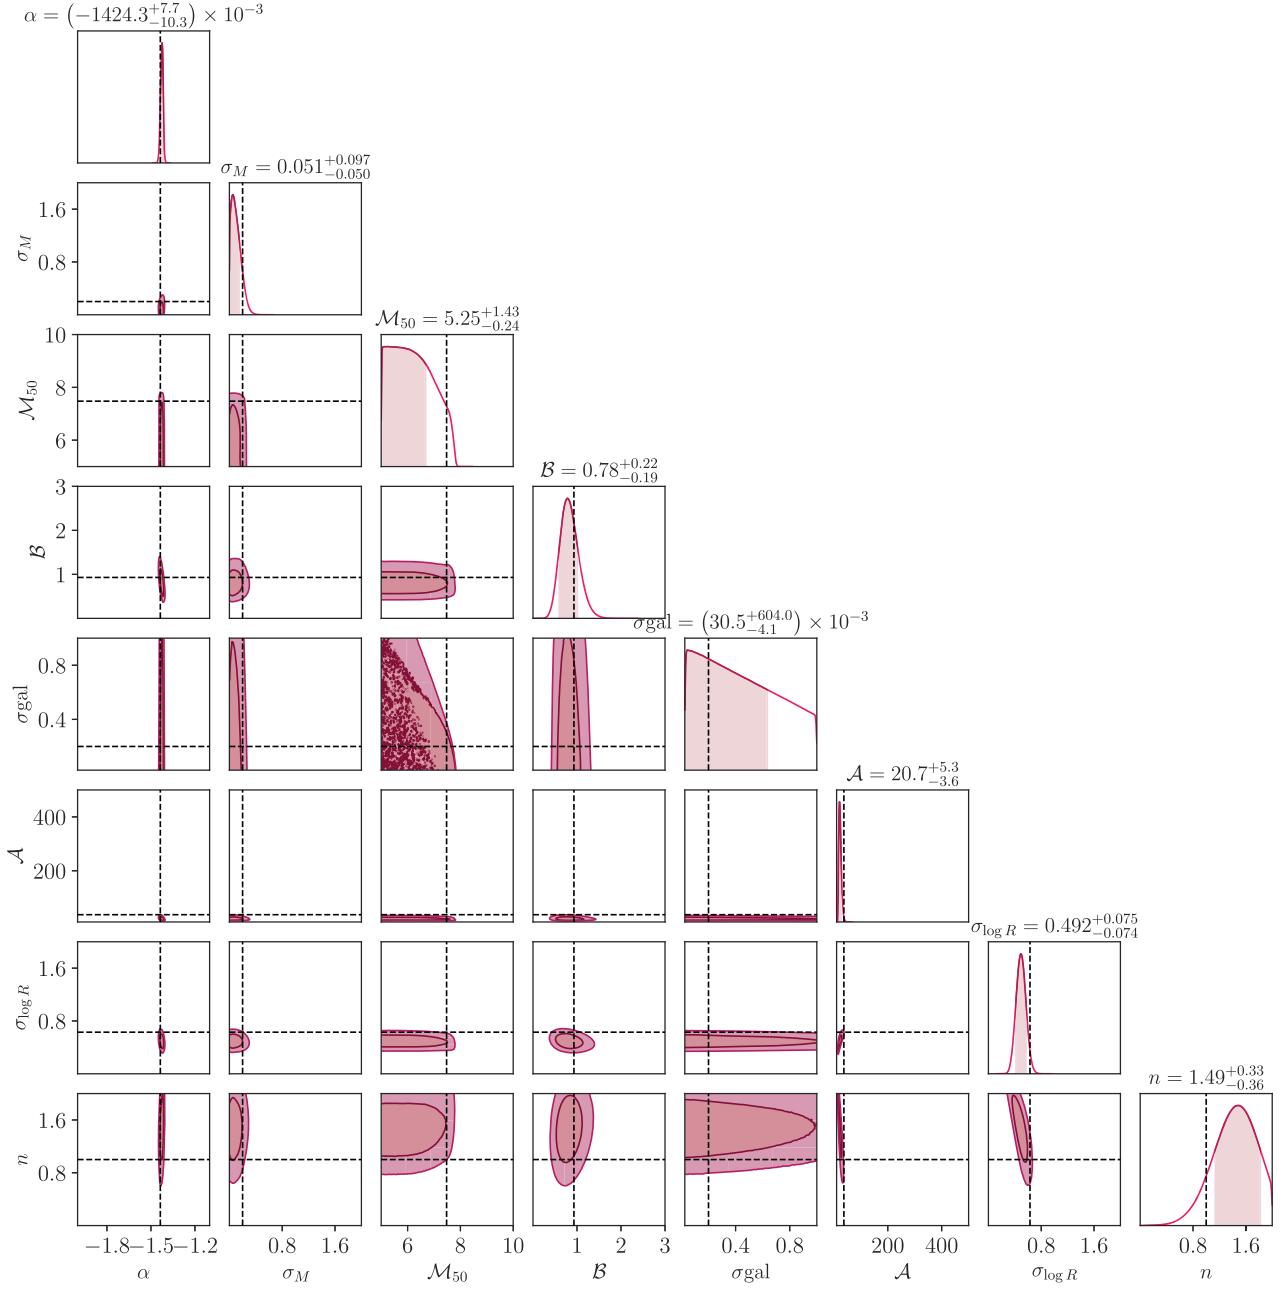


Figure 13. Posterior distribution for our “weak” galaxy formation cutoff forecast using two complete CDM satellite populations (compare to Figure 12).

Manwadkar 2022; Ahvazi et al. 2024) and hydrodynamic simulations (e.g., Sawala et al. 2016; Fitts et al. 2017; Munshi et al. 2021). Refining predictions for the shapes of both cutoffs and their dependence on halo properties beyond M_{peak} is an interesting area for future study and will help constrain our model. For example, halo masses evaluated during the epoch of reionization are expected to correlate more strongly with the galaxy occupation fraction than M_{peak} (e.g., Benitez-Llambay & Frenk 2020), which is typically achieved after reionization. We expect that incorporating such physically motivated parameterizations will help differentiate astrophysical and DM-induced cutoffs.

In addition to degeneracies with parameters related to our satellite luminosity model, galaxy–halo size connection parameters are generally constrained less precisely than in the

scenario without a WDM cutoff. For example, degeneracies between the size amplitude \mathcal{A} or size scatter $\sigma_{\log R}$ and M_{50} are exacerbated by the larger uncertainty on M_{hm} in the presence of both astrophysical and DM cutoffs. Physically, degeneracies with the size model arise because our mock observations are surface-brightness-limited, such that increasing the sizes of all galaxies at fixed galaxy–halo connection parameters can mimic a cutoff by making these systems too spatially extended to be detectable at a given luminosity (see Nadler et al. 2020b for related discussion).

C.3. Subhalo Mass Function

The posteriors from our one- and two-host SHMF forecasts assuming a “weak” galaxy formation cutoff and CDM subhalo abundances are shown in Figures 20 and 21, and the

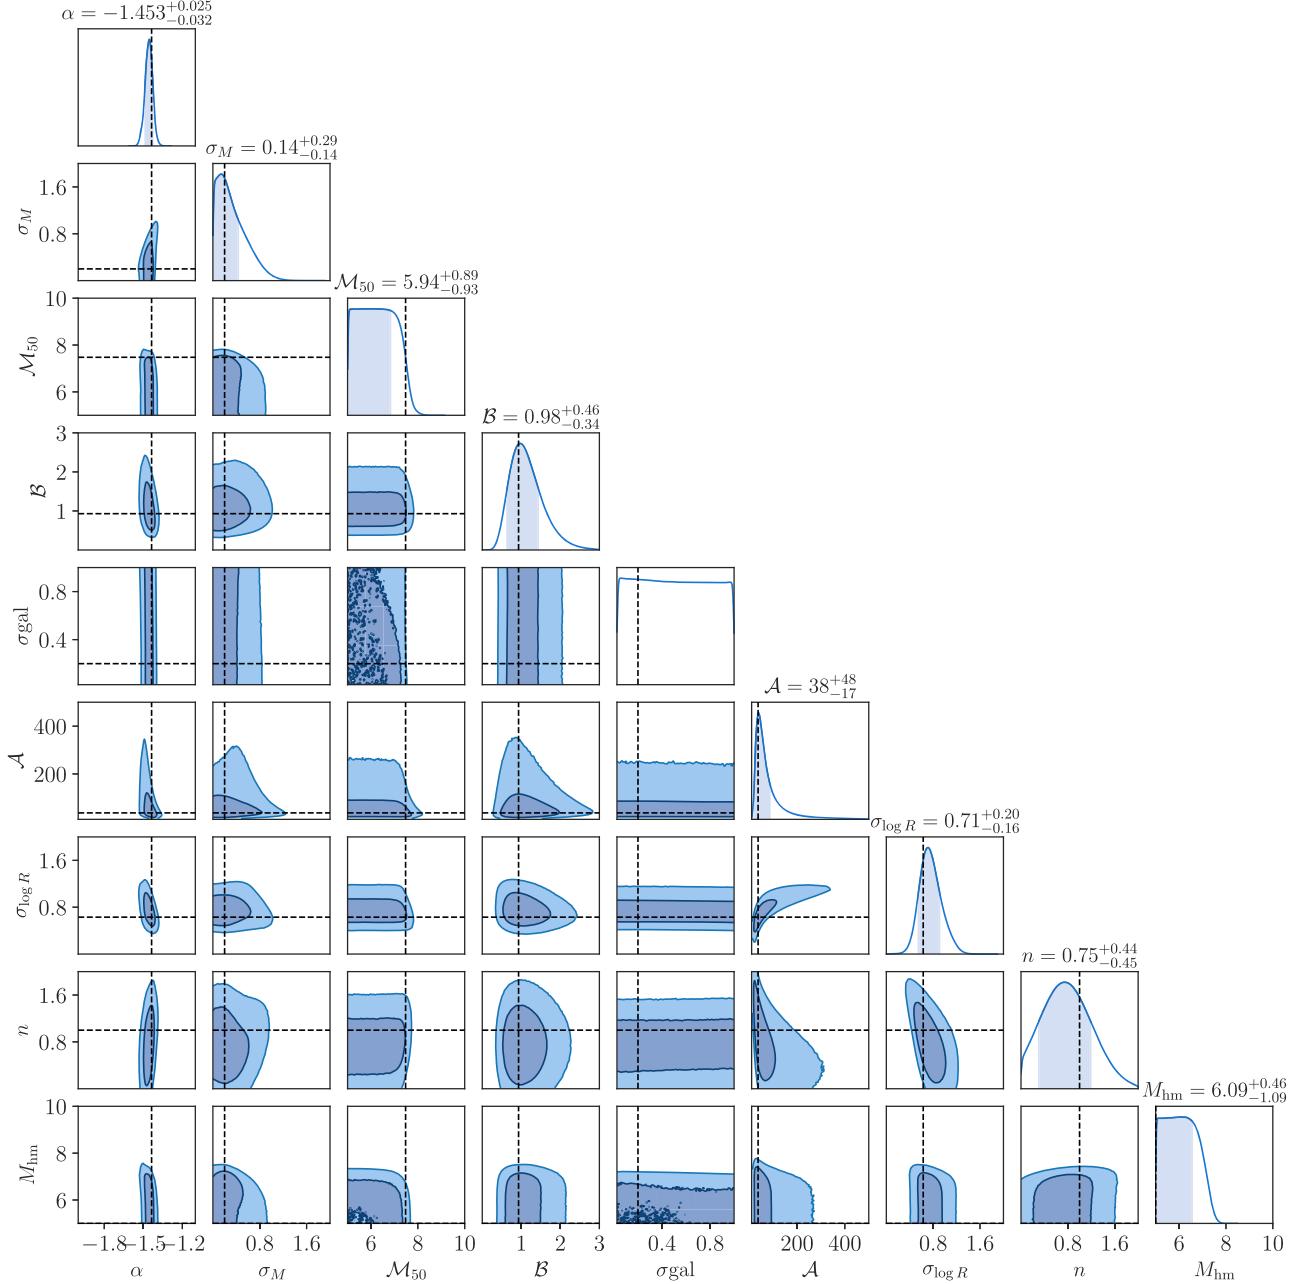


Figure 14. Posterior distribution for our WDM forecast with a “weak” galaxy formation cutoff assuming CDM subhalo abundances (Scenario A) using one complete satellite population.

corresponding “strong” galaxy formation cutoff runs are shown in Figures 22 and 23, respectively; the description below applies to both scenarios, which display similar degeneracies. We find that it is particularly challenging to sample the posterior using our MCMC method for SHMF forecasts that combine two hosts; as a result, parameters like α are biased at the 1σ – 2σ level in our two-host forecasts (see Figures 21 and 23). However, these biases do not affect the interpretation of degeneracies presented below or the precision of our projected SHMF constraints.

The SHMF amplitude at $M_{\text{peak}} = 10^{10} M_{\odot}$, ξ_{10} , does not show noticeable degeneracies with galaxy–halo connection or other SHMF parameters. As described in Section 6, this follows because (i) galaxy–halo connection parameters affect satellites that occupy lower-mass halos most strongly, (ii) faint satellites

are more abundant and thus contribute more to the likelihood, and (iii) the effects of varying ξ_{10} cannot be hidden by nondetections, since satellites hosted by the most massive subhalos are always detectable given our observational assumptions.

Several degeneracies appear between ξ_8 , ξ_9 , and the galaxy–halo connection model. Both of these SHMF amplitude parameters are positively correlated with the faint-end luminosity function slope. In our model, α only affects the galaxy–halo connection for $M_V > -13$ mag, which roughly corresponds to halos with peak virial masses between $10^9 M_{\odot}$ and $10^{10} M_{\odot}$ (e.g., Nadler et al. 2020b). Since less negative values of α yield fewer faint satellites that occupy such halos, increasing the underlying SHMF amplitude counteracts this effect. In turn, the faint-end slope is measured less precisely in our SHMF forecasts relative to

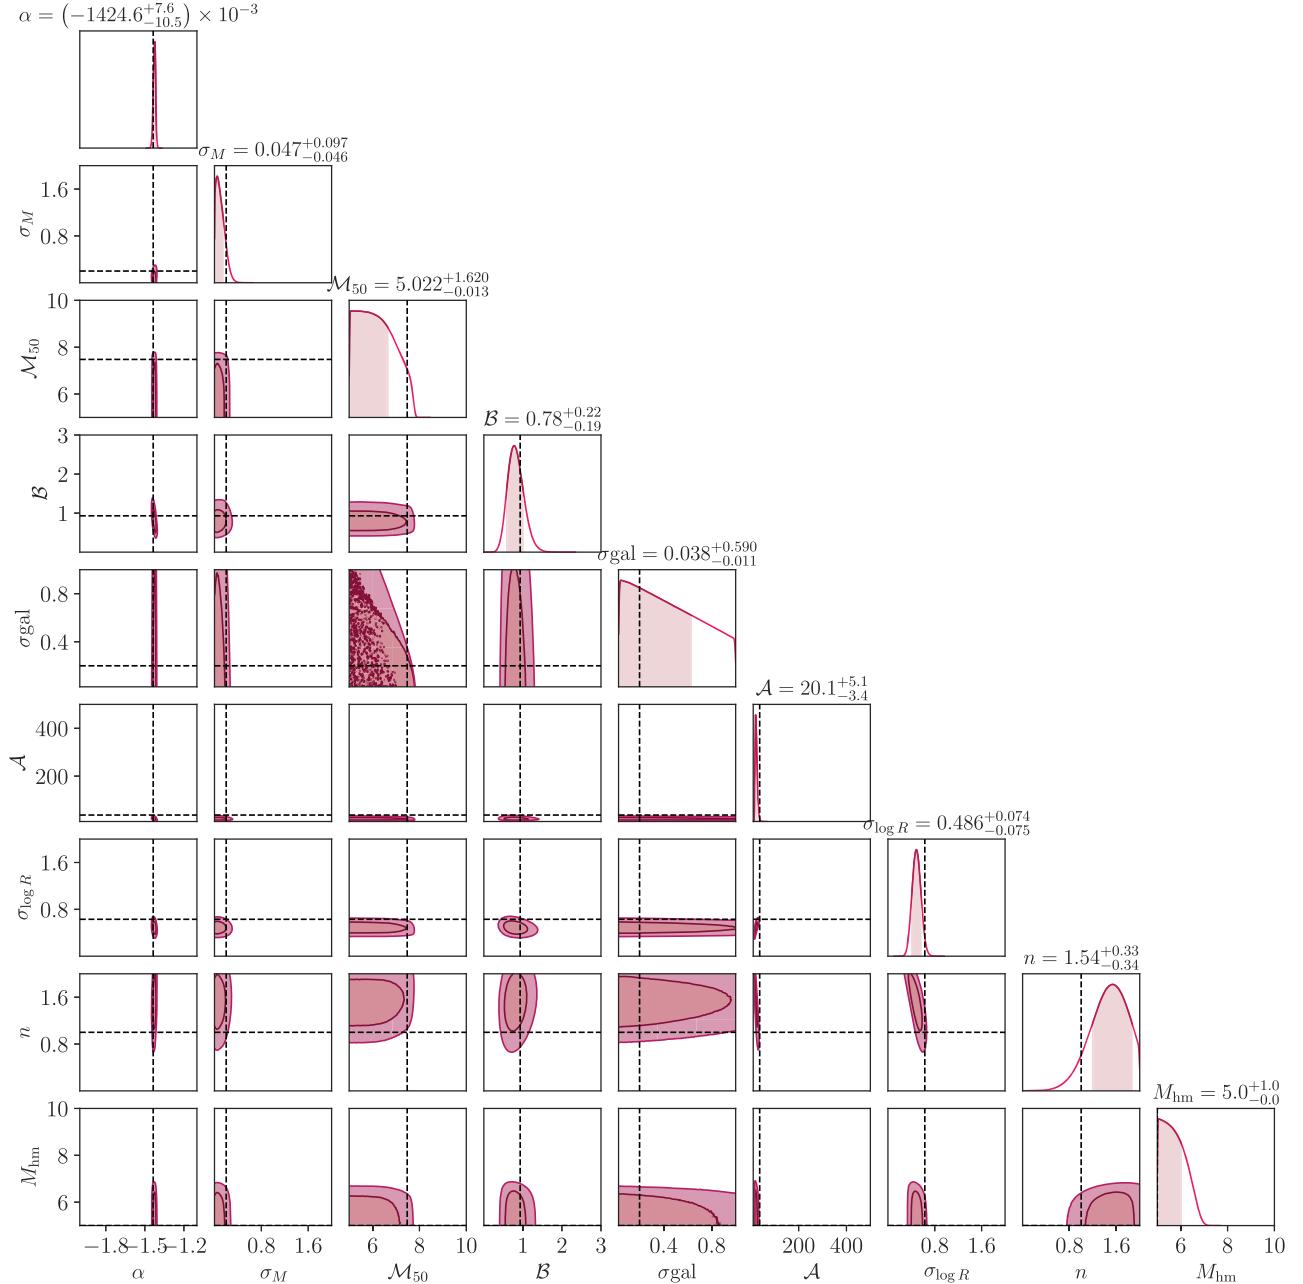


Figure 15. Posterior distribution for our WDM forecast with a “weak” galaxy formation cutoff assuming CDM subhalo abundances (Scenario A) using two complete satellite populations (compare to Figure 14).

our galaxy formation cutoff and WDM forecasts, and also displays mild degeneracies with galaxy–halo connection parameters like σ_M when the SHMF is allowed to vary.

The SHMF posteriors clearly show that ξ_8 is only constrained weakly from above, consistent with our discussion in Section 6. As a result, measurements of M_{50} degrade relative to our other forecasts when the SHMF is allowed to

vary. Interestingly, ξ_8 and ξ_9 are also mildly degenerate, which follows because our galaxy–halo connection model can produce satellites with similar properties across this decade of halo mass due to, e.g., luminosity and size scatter. This tail toward large ξ_8 persists in the two-host posteriors, even though uncertainties on remaining SHMF and galaxy–halo connection parameters are reduced.

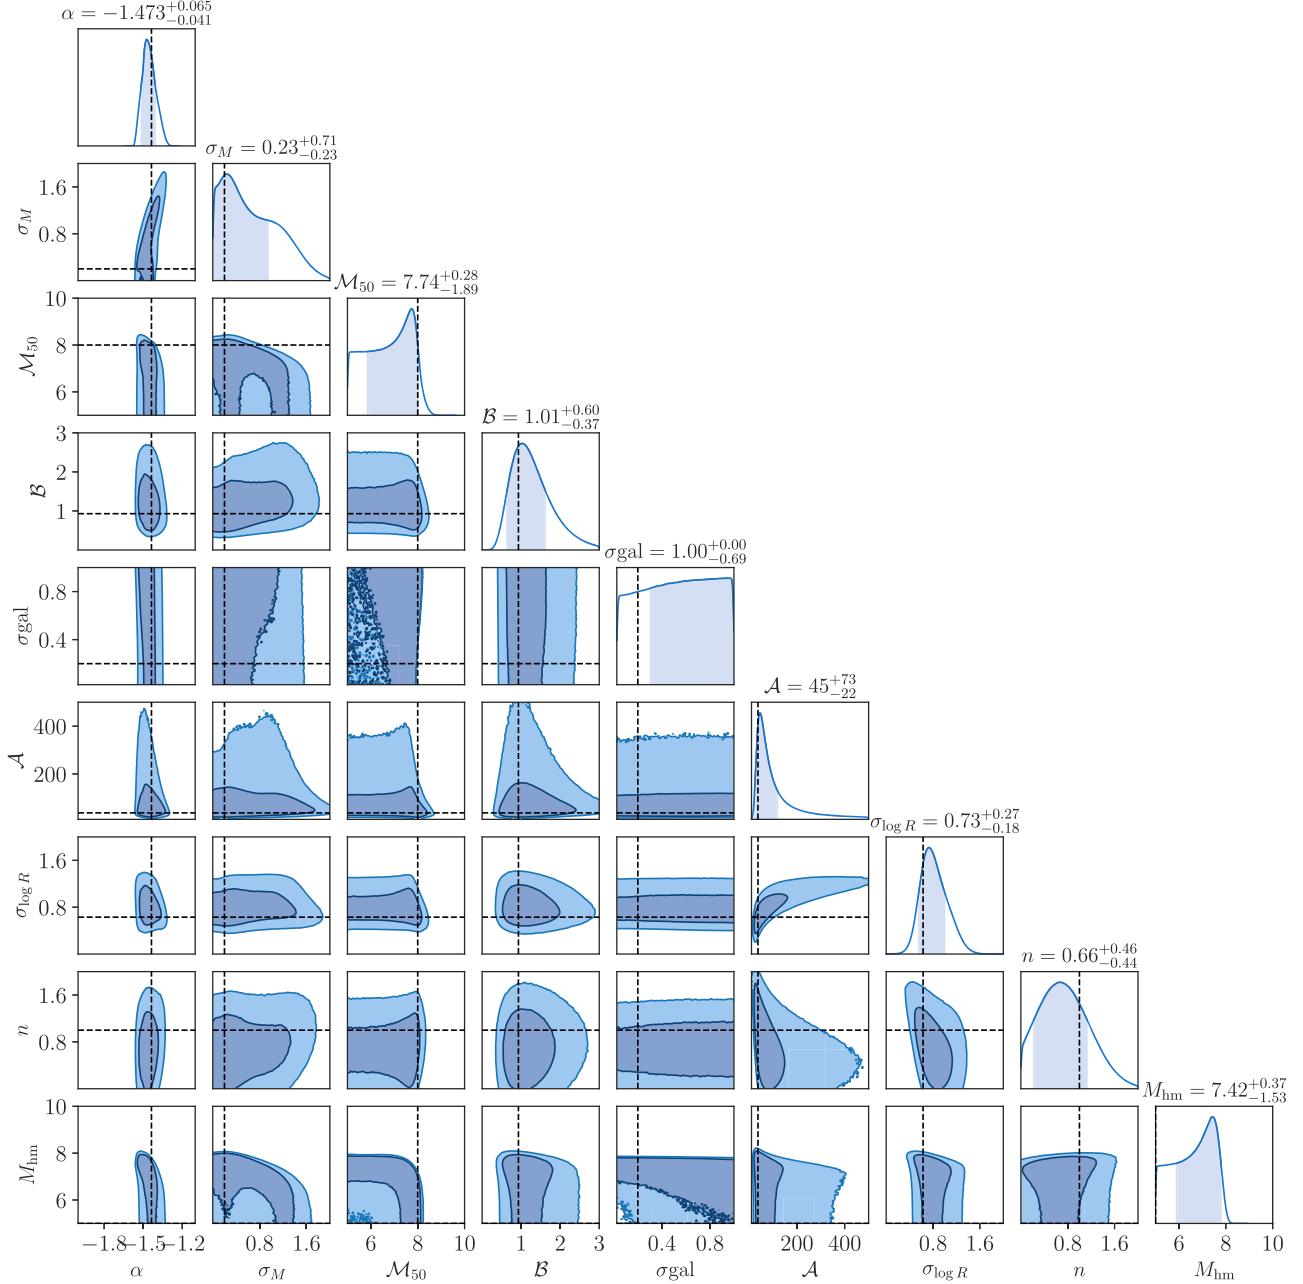


Figure 16. Posterior distribution for our WDM forecast with a “strong” galaxy formation cutoff assuming CDM subhalo abundances (Scenario B) using one complete satellite population.

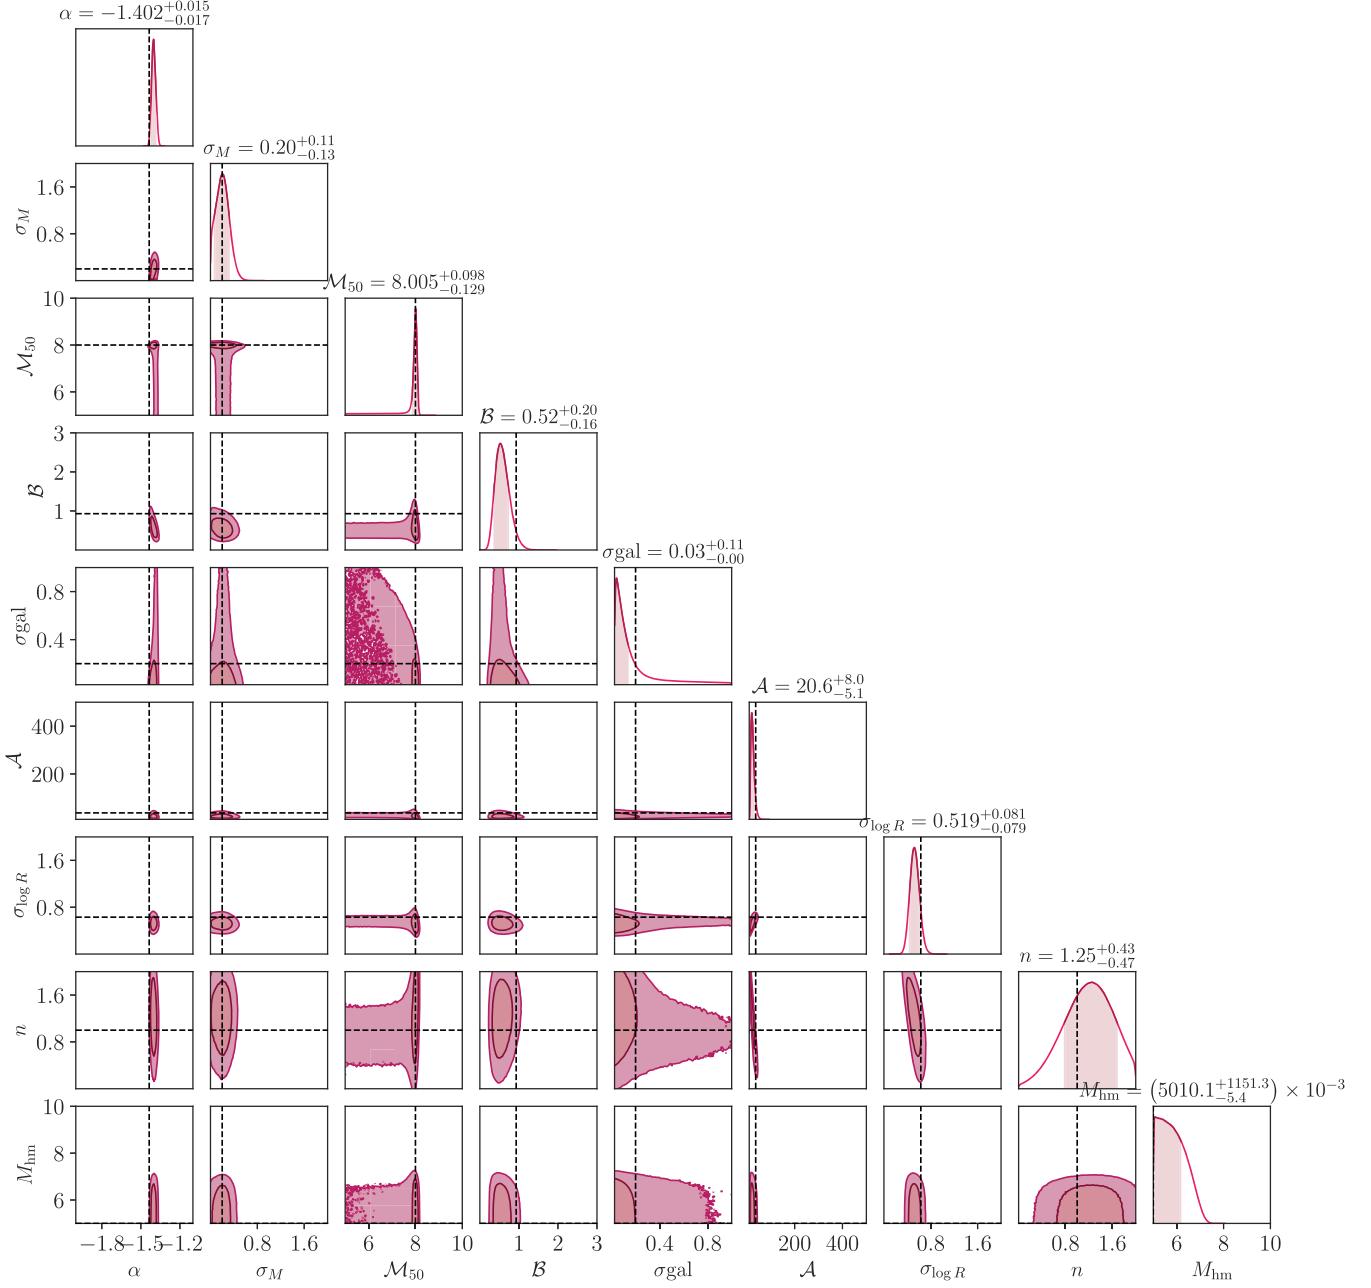


Figure 17. Posterior distribution for our WDM forecast with a “strong” galaxy formation cutoff assuming CDM subhalo abundances (Scenario B) using two complete satellite populations (compare to Figure 16).

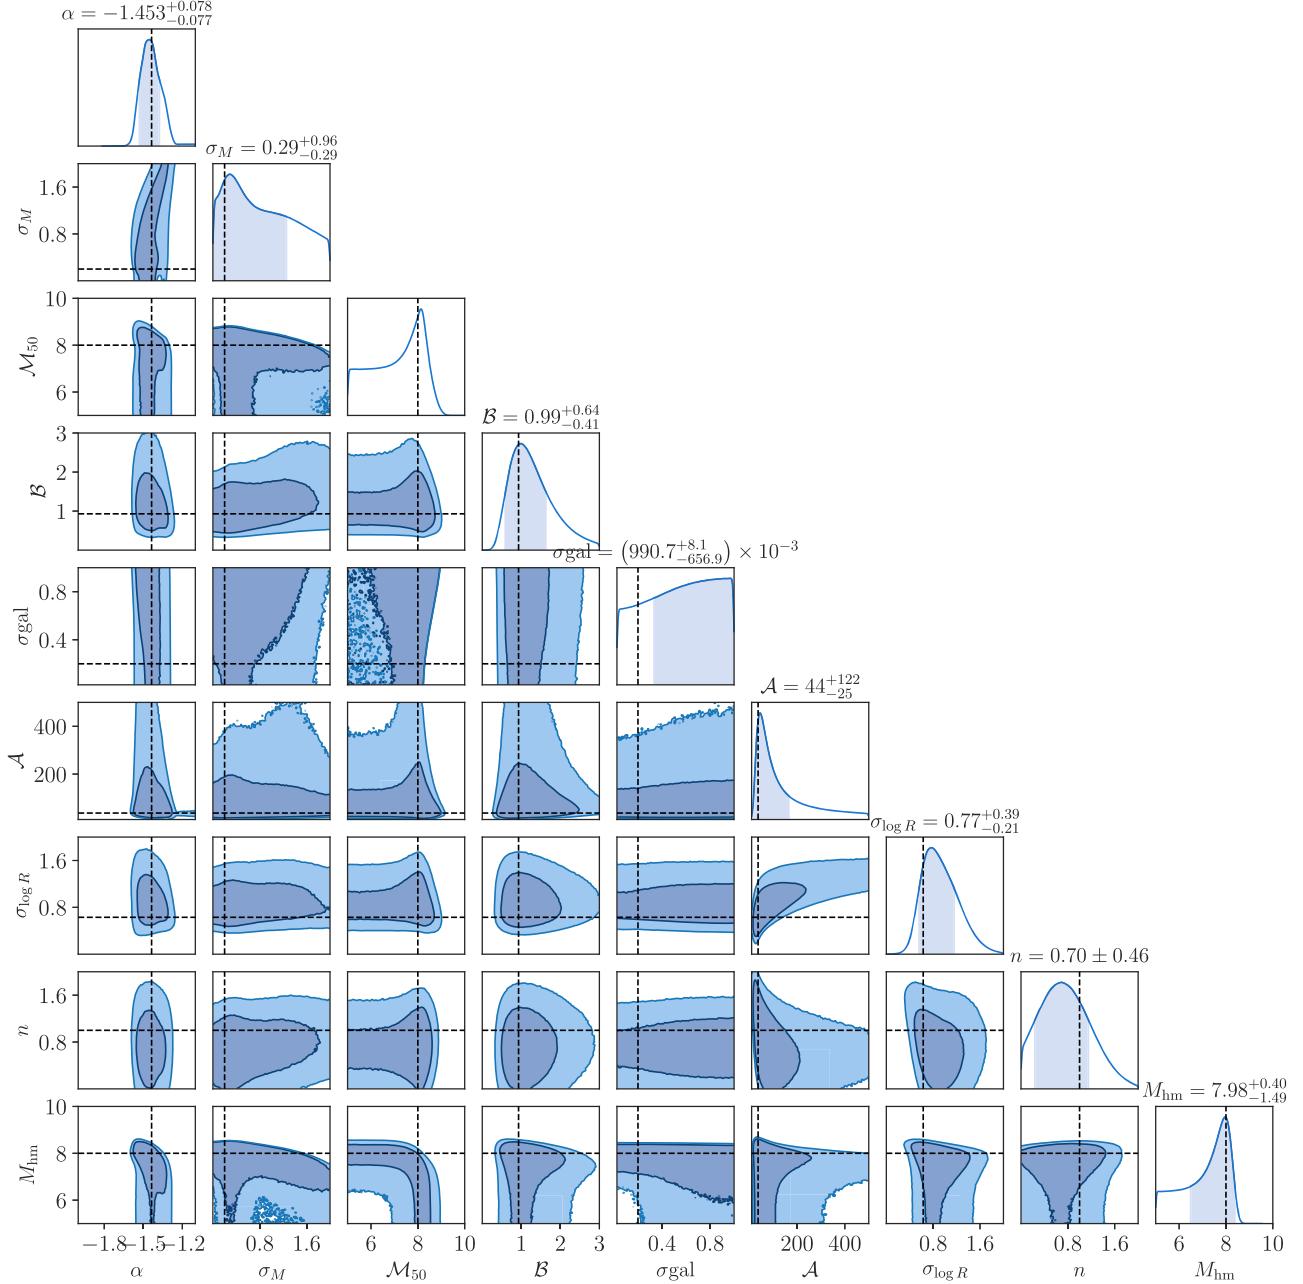


Figure 18. Posterior distribution for our WDM forecast with a “strong” galaxy formation cutoff assuming WDM subhalo abundances with $M_{\text{hm}} = 10^8 M_{\odot}$, corresponding to $m_{\text{WDM}} = 4.9 \text{ keV}$ (Scenario C), using one complete satellite population.

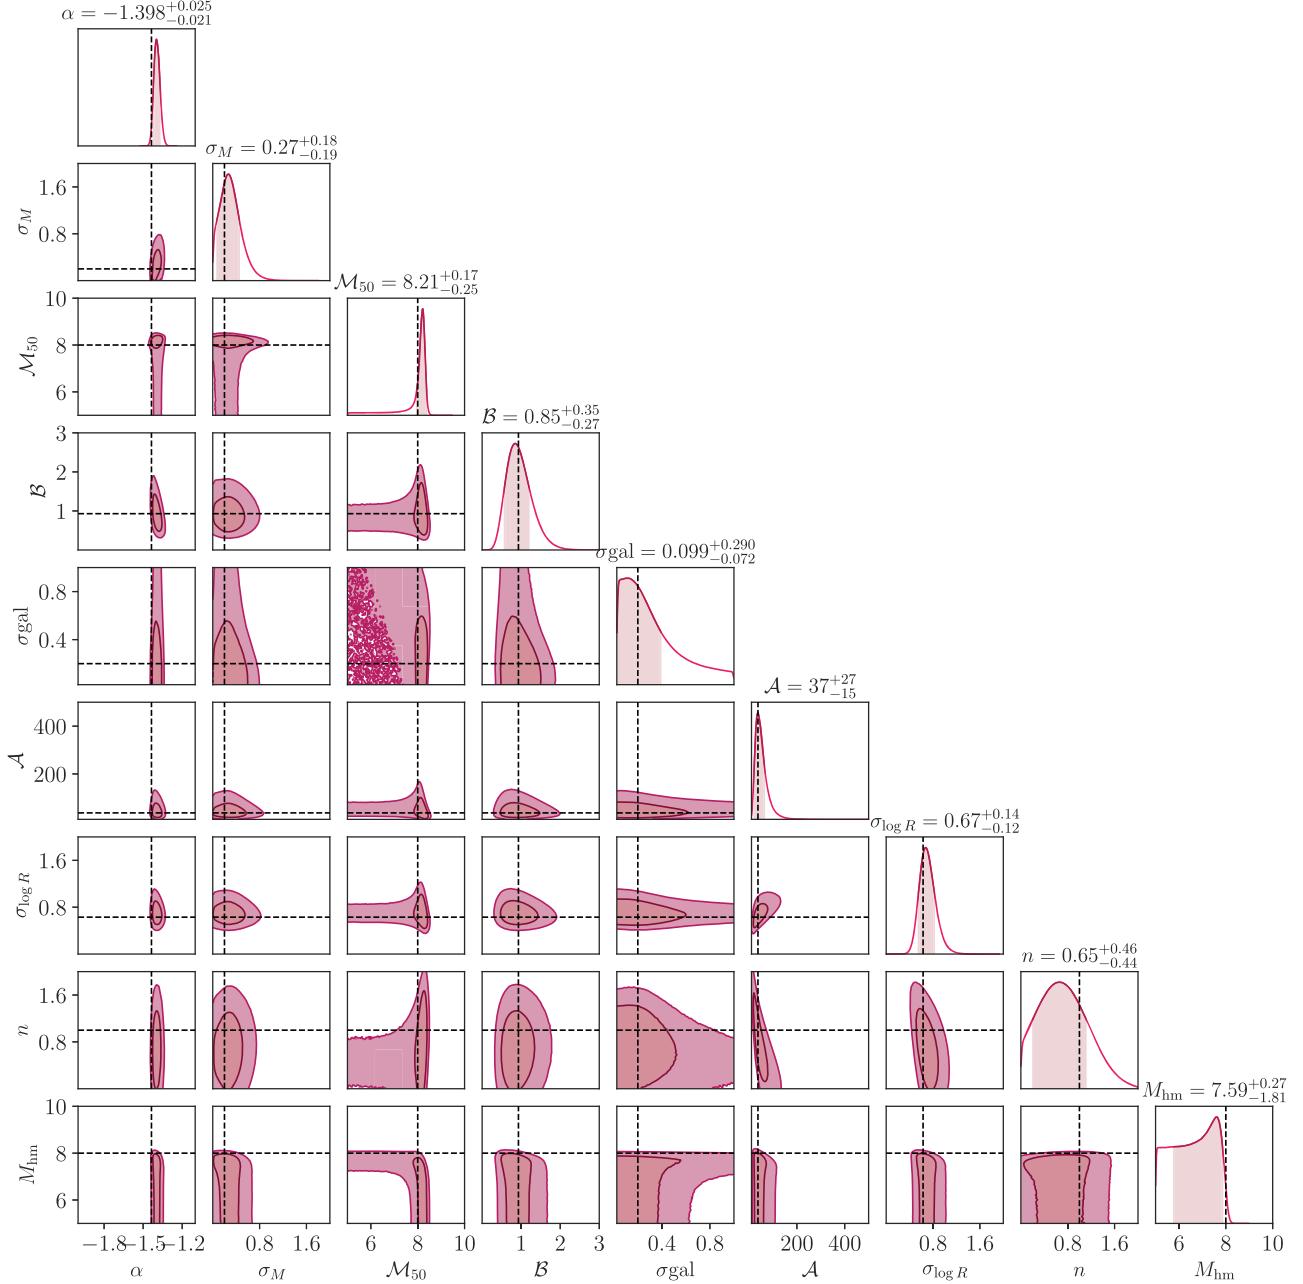


Figure 19. Posterior distribution for our WDM forecast with a “strong” galaxy formation cutoff, assuming WDM subhalo abundances with $M_{\text{hm}} = 10^8 M_{\odot}$, corresponding to $m_{\text{WDM}} = 4.9 \text{ keV}$ (Scenario C), using two complete satellite populations (compare to Figure 18).

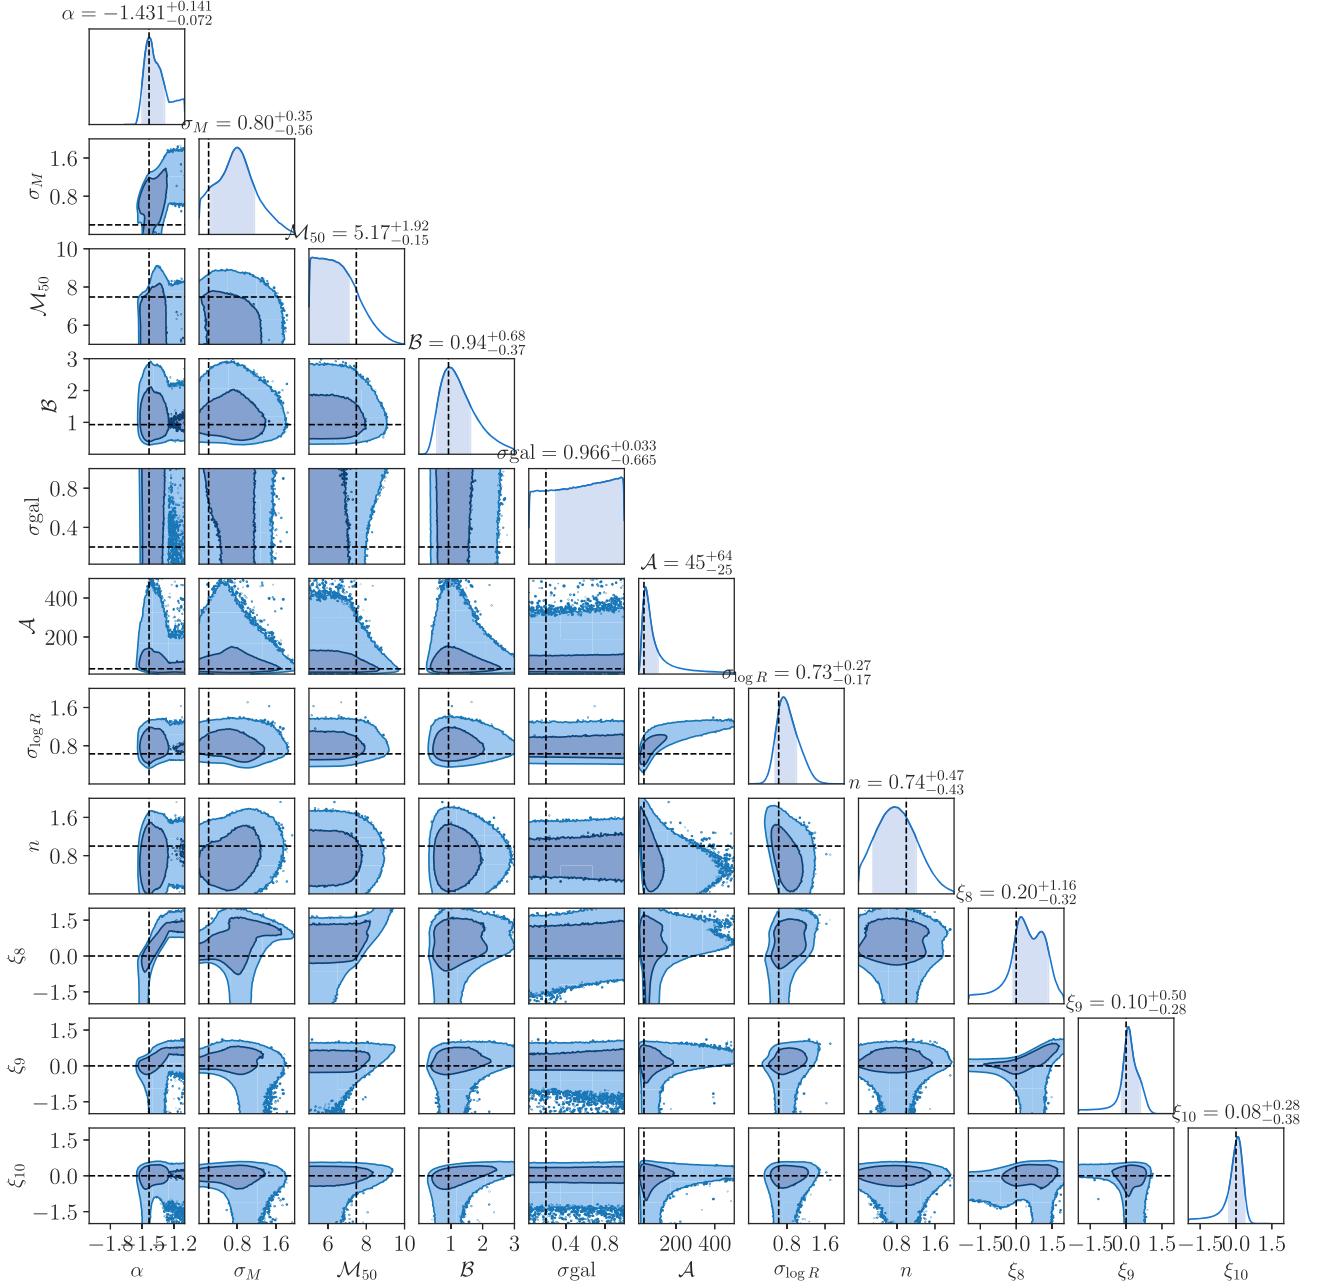


Figure 20. Posterior distribution for our SHMF forecast with a “weak” galaxy formation cutoff using one complete satellite population.

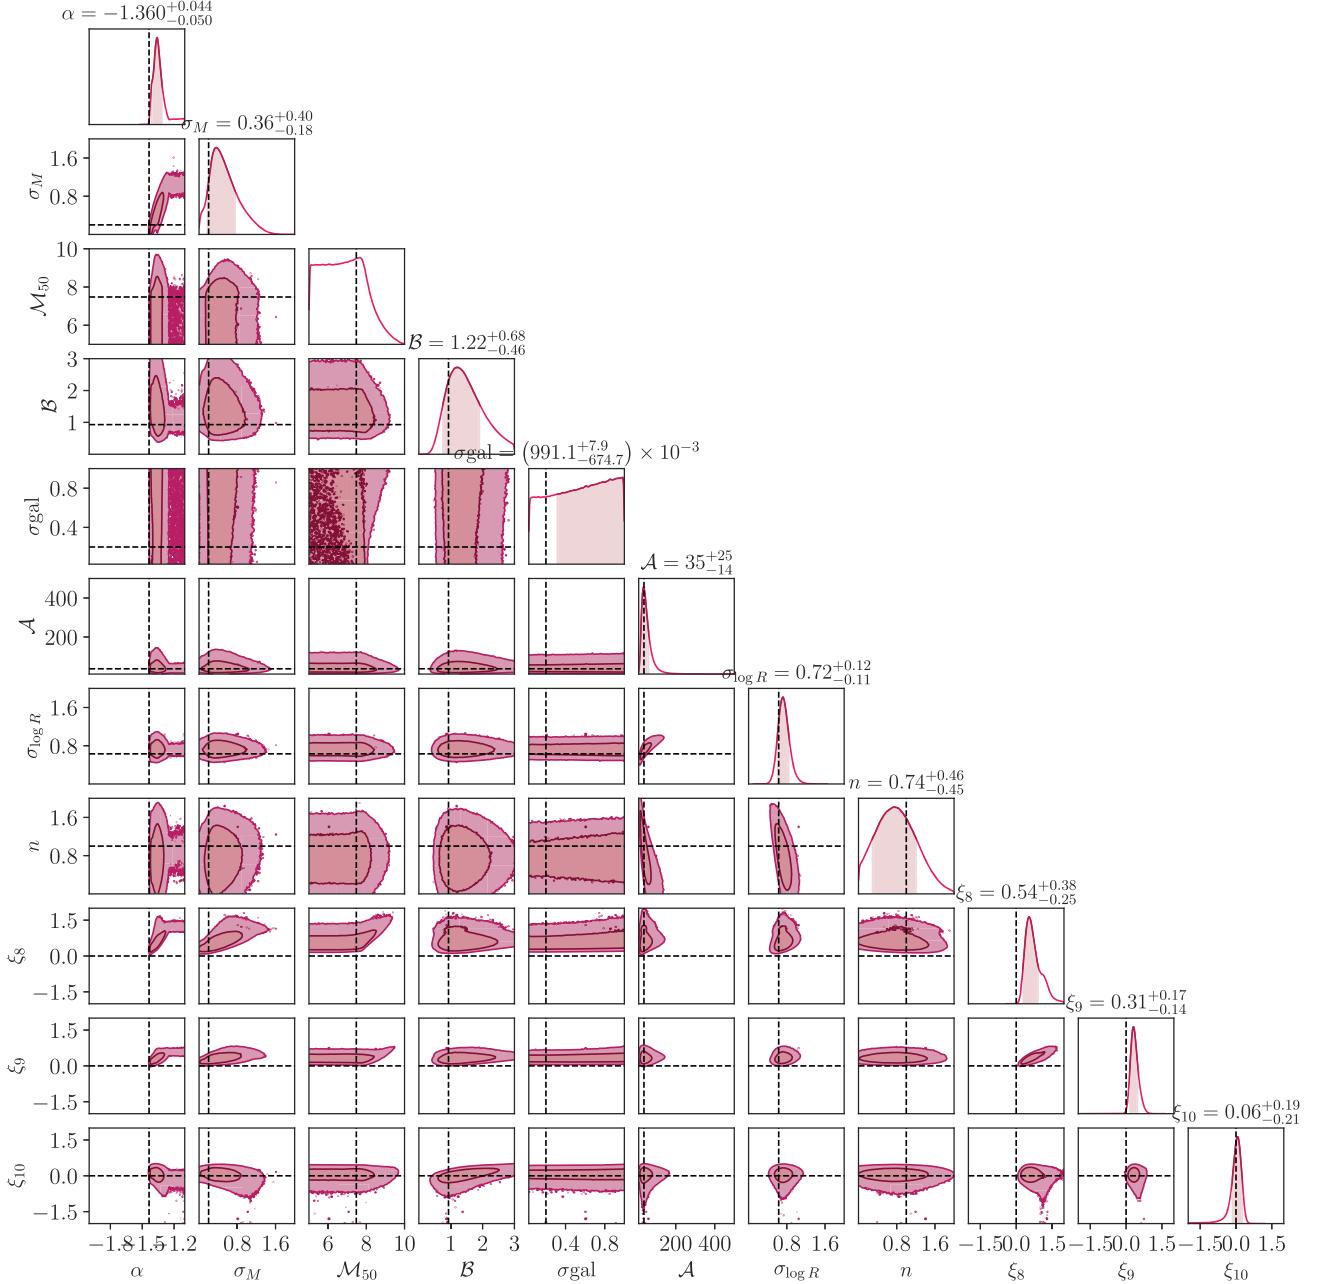


Figure 21. Posterior distribution for our SHMF forecast with a “weak” galaxy formation cutoff using two complete satellite populations (compare to Figure 20).

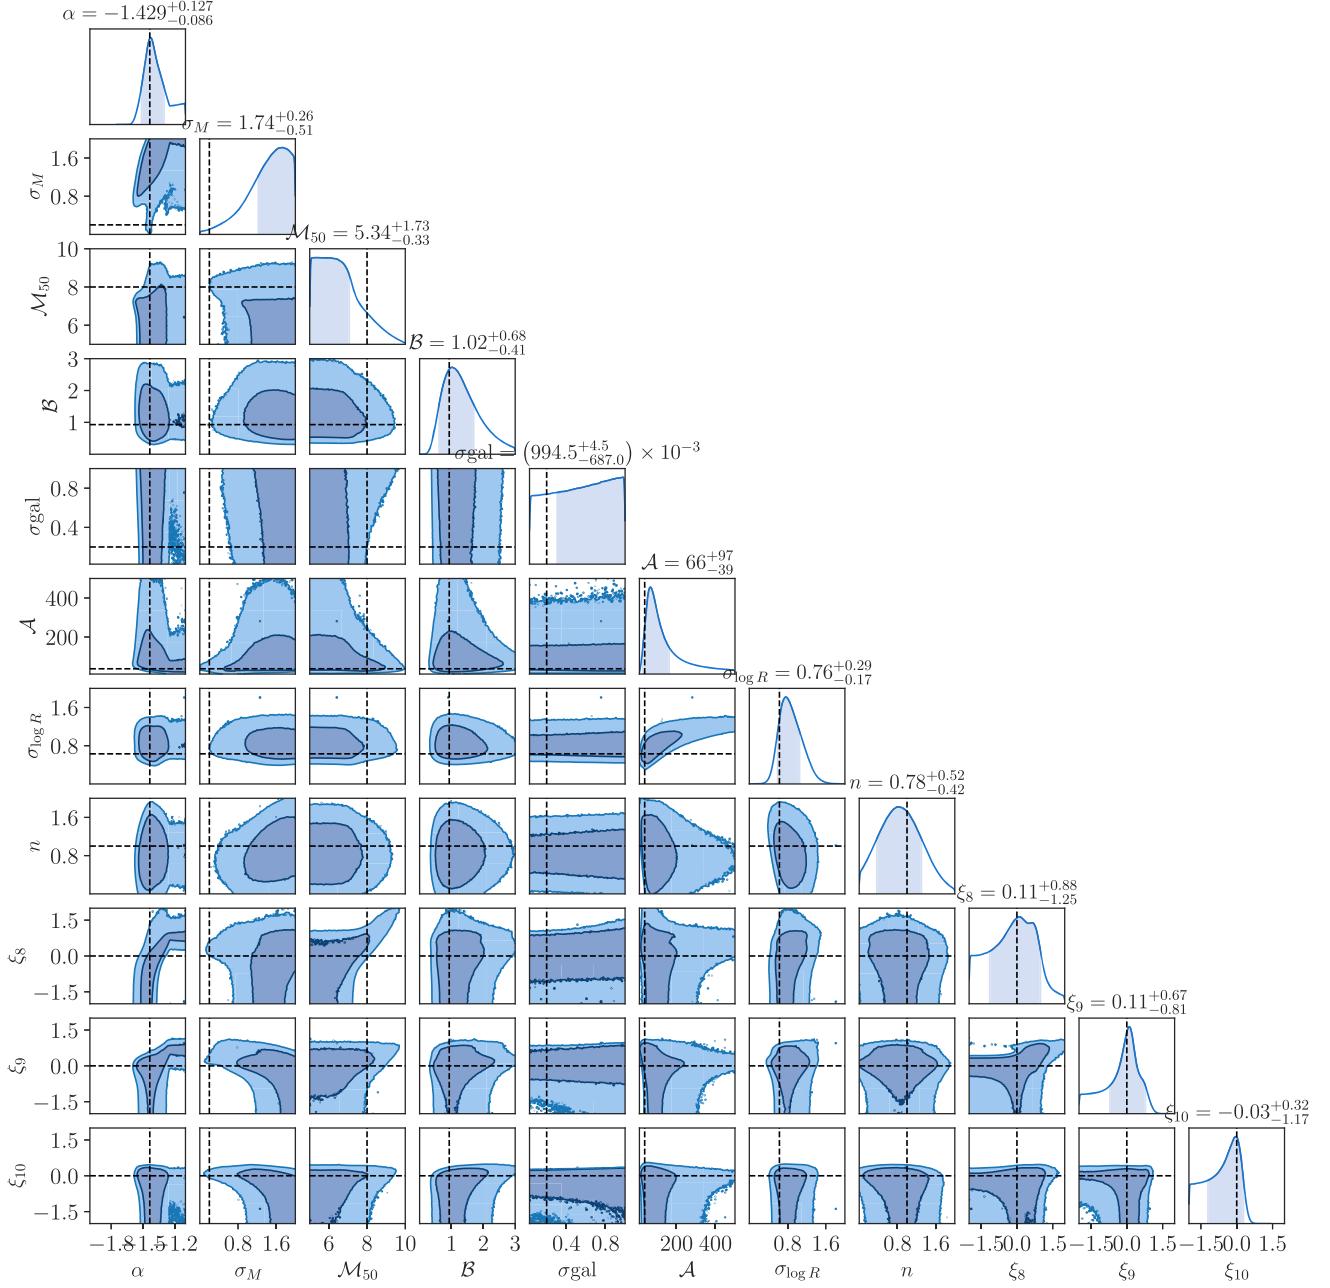


Figure 22. Posterior distribution for our SHMF forecast with a “strong” galaxy formation cutoff using one complete satellite population.

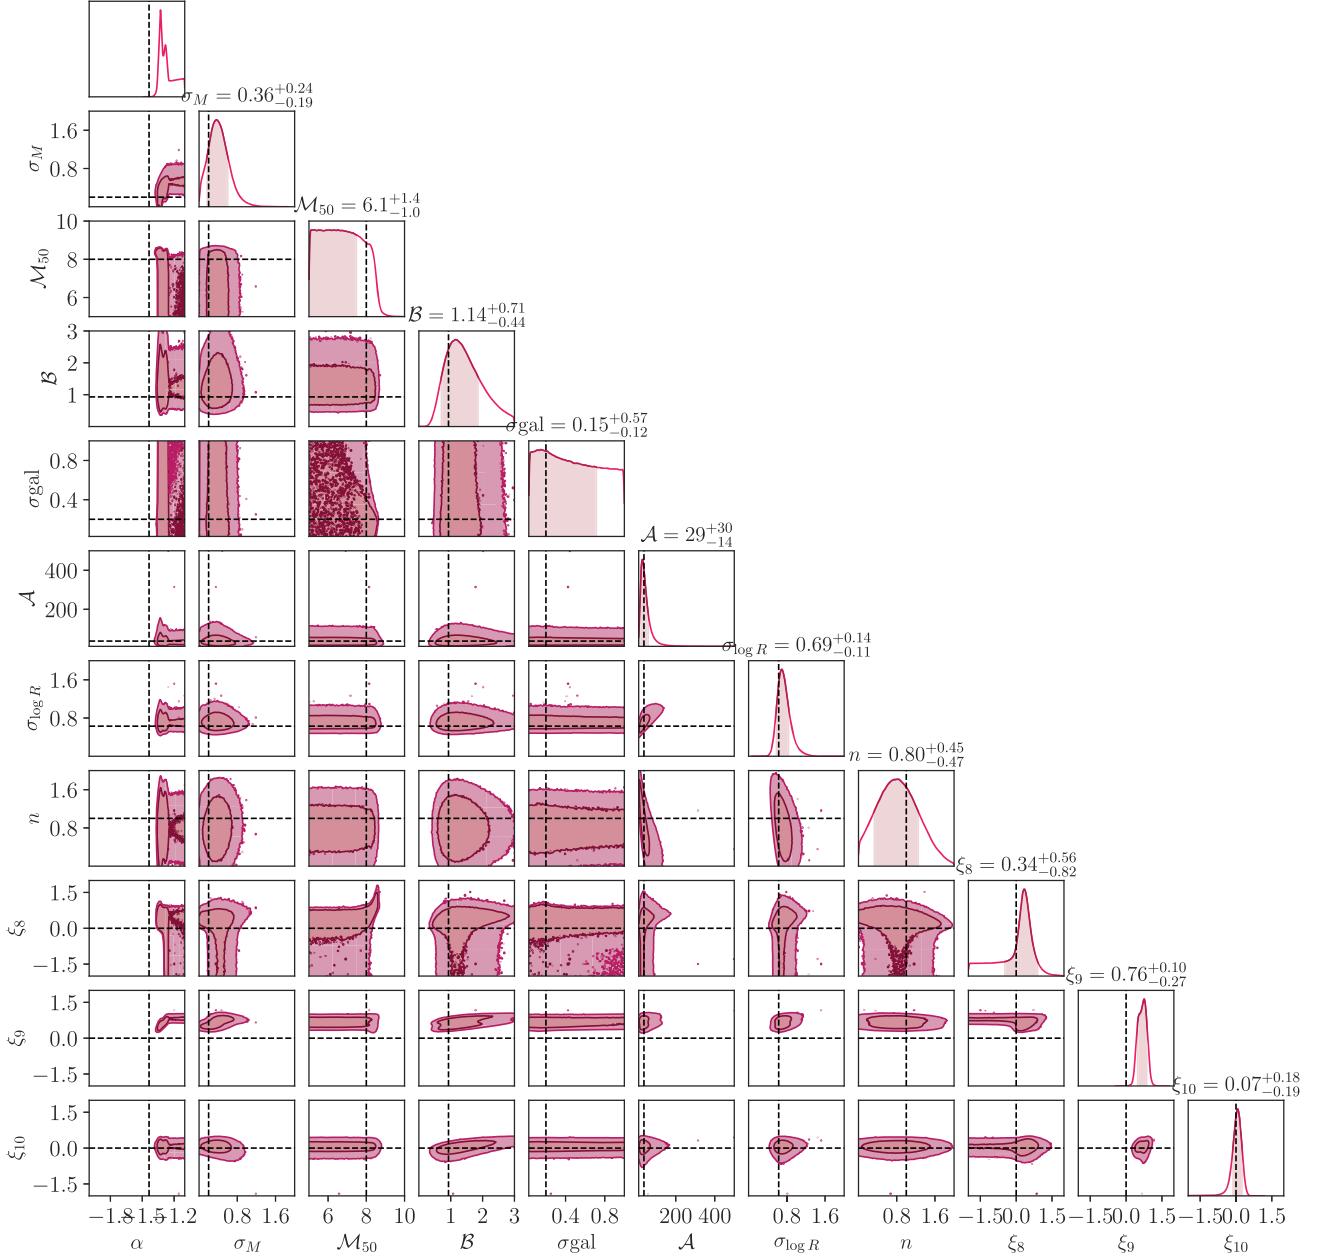


Figure 23. Posterior distribution for our SHMF forecast with a “strong” galaxy formation cutoff using two complete satellite populations (compare to Figure 22).

ORCID iDs

Ethan O. Nadler <https://orcid.org/0000-0002-1182-3825>
 Vera Gluscevic <https://orcid.org/0000-0002-3589-8637>
 Trey Driskell <https://orcid.org/0000-0001-9472-7179>
 Risa H. Wechsler <https://orcid.org/0000-0003-2229-011X>
 Leonidas A. Moustakas <https://orcid.org/0000-0003-3030-2360>
 Andrew Benson <https://orcid.org/0000-0001-5501-6008>
 Yao-Yuan Mao <https://orcid.org/0000-0002-1200-0820>

References

Abbott, T. M. C., Abdalla, F. B., Allam, S., et al. 2018, *ApJS*, **239**, 18
 Adhikari, S., Banerjee, A., Boddy, K. K., et al. 2022, arXiv:2207.10638
 Ahvazi, N., Benson, A., Sales, L. V., et al. 2024, *MNRAS*, **529**, 3387
 Angulo, R. E., Hahn, O., & Abel, T. 2013, *MNRAS*, **434**, 3337
 Arvanitaki, A., Dimopoulos, S., Galanis, M., et al. 2020, *PhRvD*, **101**, 083014
 Banik, N., Bovy, J., Bertone, G., Erkal, D., & de Boer, T. J. L. 2021a, *MNRAS*, **502**, 2364
 Banik, N., Bovy, J., Bertone, G., Erkal, D., & de Boer, T. J. L. 2021b, *JCAP*, **2021**, 043
 Barry, M., Wetzel, A., Chapman, S., et al. 2023, *MNRAS*, **523**, 428
 Bechtol, K., Birrer, S., Cyr-Racine, F.-Y., et al. 2022, arXiv:2203.07354
 Benitez-Llambay, A., & Frenk, C. 2020, *MNRAS*, **498**, 4887
 Benson, A. J., & Du, X. 2022, *MNRAS*, **517**, 1398
 Benson, A. J., Frenk, C. S., Baugh, C. M., Cole, S., & Lacey, C. G. 2003, *MNRAS*, **343**, 679
 Benson, A. J., Frenk, C. S., Lacey, C. G., Baugh, C. M., & Cole, S. 2002, *MNRAS*, **333**, 177
 Bode, P., Ostriker, J. P., & Turok, N. 2001, *ApJ*, **556**, 93
 Bonaca, A., Hogg, D. W., Price-Whelan, A. M., & Conroy, C. 2019, *ApJ*, **880**, 38
 Bond, J. R., & Szalay, A. S. 1983, *ApJ*, **274**, 443
 Bose, S., Deason, A. J., Belokurov, V., & Frenk, C. S. 2020, *MNRAS*, **495**, 743
 Bose, S., Hellwing, W. A., Frenk, C. S., et al. 2016, *MNRAS*, **455**, 318
 Bozek, B., Fitts, A., Boylan-Kolchin, M., et al. 2019, *MNRAS*, **483**, 4086
 Bryan, G. L., & Norman, M. L. 1998, *ApJ*, **495**, 80
 Bullock, J. S., & Boylan-Kolchin, M. 2017, *ARA&A*, **55**, 343
 Bullock, J. S., Stewart, K. R., Kaplinghat, M., Tollerud, E. J., & Wolf, J. 2010, *ApJ*, **717**, 1043
 Busha, M. T., Alvarez, M. A., Wechsler, R. H., Abel, T., & Strigari, L. E. 2010, *ApJ*, **710**, 408
 Carlsten, S. G., Greene, J. E., Beaton, R. L., Danieli, S., & Greco, J. P. 2022, *ApJ*, **933**, 47
 Cerny, W., Martínez-Vázquez, C. E., Drlica-Wagner, A., et al. 2023, *ApJ*, **953**, 1
 Chakrabarti, S., Drlica-Wagner, A., Li, T. S., et al. 2022, arXiv:2203.06200
 Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv:1612.05560
 Christensen, C. R., Brooks, A., Munshi, F., et al. 2024, *ApJ*, **961**, 236
 Collett, T. E. 2015, *ApJ*, **811**, 20
 Conroy, C., Johnson, B. D., van Dokkum, P., et al. 2023, arXiv:2310.13048
 Côté, B., Silvia, D. W., O’Shea, B. W., Smith, B., & Wise, J. H. 2018, *ApJ*, **859**, 67
 Cyncynates, D., Giurgica-Tiron, T., Simon, O., & Thompson, J. O. 2022, *PhRvD*, **105**, 055005
 Danieli, S., Greene, J. E., Carlsten, S., et al. 2023, *ApJ*, **956**, 6
 Darragh-Ford, E., Wu, J. F., Mao, Y.-Y., et al. 2023, *ApJ*, **954**, 149
 Decant, Q., Heisig, J., Hooper, D. C., & Lopez-Honorez, L. 2022, *JCAP*, **2022**, 041
 Dekker, A., Ando, S., Correa, C. A., & Ng, K. C. Y. 2022, *PhRvD*, **106**, 123026
 Despali, G., Vegetti, S., White, S. D. M., Giocoli, C., & van den Bosch, F. C. 2018, *MNRAS*, **475**, 5424
 Diego, J. M., Sun, B., Yan, H., et al. 2023, *A&A*, **679**, A31
 Dixon, K. L., Iliev, I. T., Gottlöber, S., et al. 2018, *MNRAS*, **477**, 867
 Doliva-Dolinsky, A., Martín, N. F., Thomas, G. F., et al. 2022, *ApJ*, **933**, 135
 Drlica-Wagner, A., Bechtol, K., Mau, S., et al. 2020, *ApJ*, **893**, 47
 Drlica-Wagner, A., Mao, Y.-Y., Adhikari, S., et al. 2019, arXiv:1902.01055
 D’Souza, R., & Bell, E. F. 2021, *MNRAS*, **504**, 5270
 Elahi, P. J., Mahdi, H. S., Power, C., & Lewis, G. F. 2014, *MNRAS*, **444**, 2333
 Enzi, W., Murgia, R., Newton, O., et al. 2021, *MNRAS*, **506**, 5848
 Erickcek, A. L., & Sigurdson, K. 2011, *PhRvD*, **84**, 083503
 Errani, R., & Navarro, J. F. 2021, *MNRAS*, **505**, 18
 Esteban, I., Peter, A. H. G., & Kim, S. Y. 2023, arXiv:2306.04674
 Euclid Collaboration, Scaramella, R., Amiaux, J., et al. 2022, *A&A*, **662**, A112
 Fielder, C. E., Mao, Y.-Y., Newman, J. A., Zentner, A. R., & Licquia, T. C. 2019, *MNRAS*, **486**, 4545
 Fitts, A., Boylan-Kolchin, M., Bozek, B., et al. 2019, *MNRAS*, **490**, 962
 Fitts, A., Boylan-Kolchin, M., Elbert, O. D., et al. 2017, *MNRAS*, **471**, 3547
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
 Gardner, S., McDermott, S. D., & Yanny, B. 2021, *PrPnP*, **121**, 103904
 Garrison-Kimmel, S., Bullock, J. S., Boylan-Kolchin, M., & Bardwell, E. 2017a, *MNRAS*, **464**, 3108
 Garrison-Kimmel, S., Wetzel, A., Bullock, J. S., et al. 2017b, *MNRAS*, **471**, 1709
 Geha, M., Wechsler, R. H., Mao, Y.-Y., et al. 2017, *ApJ*, **847**, 4
 Gezari, S., Bentz, M., De, K., et al. 2022, arXiv:2202.12311
 Gilman, D., Benson, A., Bovy, J., et al. 2022, *MNRAS*, **512**, 3163
 Gilman, D., Birrer, S., Nierenberg, A., et al. 2020a, *MNRAS*, **491**, 6077
 Gilman, D., Du, X., Benson, A., et al. 2020b, *MNRAS*, **492**, L12
 Gilman, D., Birrer, S., Treu, T., Nierenberg, A., & Benson, A. 2019, *MNRAS*, **487**, 5721
 Gluscevic, V., Ali-Haimoud, Y., Bechtol, K., et al. 2019, *BAAS*, **51**, 134
 Graus, A. S., Bullock, J. S., Kelley, T., et al. 2019, *MNRAS*, **488**, 4585
 Greco, J. P., van Dokkum, P., Danieli, S., Carlsten, S. G., & Conroy, C. 2021, *ApJ*, **908**, 24
 Green, S. B., van den Bosch, F. C., & Jiang, F. 2021, *MNRAS*, **503**, 4075
 Green, S. B., van den Bosch, F. C., & Jiang, F. 2022, *MNRAS*, **509**, 2624
 Hezaveh, Y. D., Dalal, N., Marrone, D. P., et al. 2016, *ApJ*, **823**, 37
 Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, *ApJS*, **208**, 19
 Hinton, S. R. 2016, *JOSS*, **1**, 00045
 Hsueh, J. W., Enzi, W., Vegetti, S., et al. 2020, *MNRAS*, **492**, 3047
 Hunter, J. D. 2007, *CSE*, **9**, 90
 Iršič, V., Viel, M., Haehnelt, M. G., et al. 2017, *PhRvD*, **96**, 023522
 Iršič, V., Viel, M., Haehnelt, M. G., et al. 2024, *PhRvD*, **109**, 043511
 Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2008, *ApJ*, **873**, 111
 Jiang, F., Dekel, A., Kneller, O., et al. 2019, *MNRAS*, **488**, 4801
 Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open source scientific tools for Python, <http://www.scipy.org/>
 Kadler, M., Riechers, D. A., Bacsko, A. K., et al. 2023, arXiv:2311.10056
 Kallivayalil, N., Sales, L. V., Zivick, P., et al. 2018, *ApJ*, **867**, 19
 Kelley, T., Bullock, J. S., Garrison-Kimmel, S., et al. 2019, *MNRAS*, **487**, 4409
 Kim, S. Y., & Peter, A. H. G. 2021, arXiv:2106.09050
 Knebe, A., Arnold, B., Power, C., & Gibson, B. K. 2008, *MNRAS*, **386**, 1029
 Kravtsov, A., & Manwadkar, V. 2022, *MNRAS*, **514**, 2667
 Kravtsov, A. V. 2013, *ApJL*, **764**, L31
 Lehmann, B. V., Mao, Y.-Y., Becker, M. R., Skillman, S. W., & Wechsler, R. H. 2017, *ApJ*, **834**, 37
 Lewis, A. 2019, arXiv:1910.13970
 Li, T. Y., Alvarez, M. A., Wechsler, R. H., & Abel, T. 2014, *ApJ*, **785**, 134
 Loveday, J., Norberg, P., Baldry, I. K., et al. 2015, *MNRAS*, **451**, 1540
 Lovell, M. R., Cautun, M., Frenk, C. S., Hellwing, W. A., & Newton, O. 2021, *MNRAS*, **507**, 4826
 Lovell, M. R., Frenk, C. S., Eke, V. R., et al. 2014, *MNRAS*, **439**, 300
 Mansfield, P., Darragh-Ford, E., Wang, Y., Nadler, E. O., & Wechsler, R. H. 2023, arXiv:2308.10926
 Mao, Y.-Y., Geha, M., Wechsler, R. H., et al. 2021, *ApJ*, **907**, 85
 Mao, Y.-Y., Williamson, M., & Wechsler, R. H. 2015, *ApJ*, **810**, 21
 Mau, S., Nadler, E. O., Wechsler, R. H., et al. 2022, *ApJ*, **932**, 128
 McAlpine, S., Helly, J. C., Schaller, M., et al. 2022, *MNRAS*, **512**, 5823
 Munshi, F., Brooks, A. M., Applebaum, E., et al. 2021, *ApJ*, **923**, 35
 Mutlu-Pakdil, B., Sand, D. J., Crnojević, D., et al. 2021, *ApJ*, **918**, 88
 Nadler, E. O. 2021, PhD thesis, Stanford Univ.
 Nadler, E. O., Banerjee, A., Adhikari, S., Mao, Y.-Y., & Wechsler, R. H. 2020a, *ApJ*, **896**, 112
 Nadler, E. O., Wechsler, R. H., Bechtol, K., et al. 2020b, *ApJ*, **893**, 48
 Nadler, E. O., Banerjee, A., Adhikari, S., Mao, Y.-Y., & Wechsler, R. H. 2021a, *ApJL*, **920**, L11
 Nadler, E. O., Birrer, S., Gilman, D., et al. 2021b, *ApJ*, **917**, 7
 Nadler, E. O., Drlica-Wagner, A., Bechtol, K., et al. 2021c, *PhRvL*, **126**, 091101
 Nadler, E. O., Gluscevic, V., Boddy, K. K., & Wechsler, R. H. 2019a, *ApJL*, **878**, L32
 Nadler, E. O., Mao, Y.-Y., Green, G. M., & Wechsler, R. H. 2019b, *ApJ*, **873**, 34
 Nadler, E. O., Mansfield, P., Wang, Y., et al. 2023, *ApJ*, **945**, 159

Nadler, E. O., Mao, Y.-Y., Wechsler, R. H., Garrison-Kimmel, S., & Wetzel, A. 2018, *ApJ*, **859**, 129

Neuzil, M. K., Mansfield, P., & Kravtsov, A. V. 2020, *MNRAS*, **494**, 2600

Newton, O., Leo, M., Cautun, M., et al. 2021, *JCAP*, **2021**, 062

Nierenberg, A. M., Keeley, R. E., Sluse, D., et al. 2024, *MNRAS*, **530**, 2960

Oguri, M., & Marshall, P. J. 2010, *MNRAS*, **405**, 2579

Patel, E., Kallivayalil, N., Garavito-Camargo, N., et al. 2020, *ApJ*, **893**, 121

Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, *JMLR*, **12**, 2825

Powell, D. M., Vegetti, S., McKean, J. P., et al. 2022, *MNRAS*, **516**, 1808

Qu, H., Yuan, Z., Doliva-Dolinsky, A., et al. 2023, *MNRAS*, **523**, 876

Racca, G. D., Laureijs, R., Stagnaro, L., et al. 2016, *Proc. SPIE*, **9904**, 99040O

Richings, J., Frenk, C., Jenkins, A., et al. 2020, *MNRAS*, **492**, 5780

Sales, L. V., Wetzel, A., & Fattah, A. 2022, *NatAs*, **6**, 897

Sawala, T., Frenk, C. S., Fattah, A., et al. 2016, *MNRAS*, **456**, 85

Selina, R. J., Murphy, E. J., McKinnon, M., et al. 2018, *Proc. SPIE*, **10700**, 1070010

Sengül, A. Ç., Dvorkin, C., Ostdiek, B., & Tsang, A. 2022, *MNRAS*, **515**, 4391

Simon, J., Bechtol, K., Drlica-Wagner, A., et al. 2019, *BAAS*, **51**, 409

Simon, J. D. 2019, *ARA&A*, **57**, 375

Smith, S. E. T., Cerny, W., Hayes, C. R., et al. 2023, arXiv:2311.10147

Spergel, D., Gehrels, N., Baltay, C., et al. 2015, arXiv:1503.03757

Stücker, J., Angulo, R. E., Hahn, O., & White, S. D. M. 2022, *MNRAS*, **509**, 1703

The LUVOIR Team 2019, arXiv:1912.06219

Tulin, S., & Yu, H.-B. 2018, *PhR*, **730**, 1

van den Bosch, F. C., Ogiya, G., Hahn, O., & Burkert, A. 2018, *MNRAS*, **474**, 3043

van den Bosch, F. C., & Ogiya, G. 2018, *MNRAS*, **475**, 4066

van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22

Vegetti, S., Koopmans, L. V. E., Bolton, A., Treu, T., & Gavazzi, R. 2010, *MNRAS*, **408**, 1969

Vegetti, S., Lagattuta, D. J., McKean, J. P., et al. 2012, *Natur*, **481**, 341

Vogel, C. M., & Abazajian, K. N. 2023, *PhRvD*, **108**, 043520

Webb, J. J., & Bovy, J. 2020, *MNRAS*, **499**, 116

Wechsler, R. H., & Tinker, J. L. 2018, *ARAA*, **56**, 435

Weisz, D. R., Dolphin, A. E., Skillman, E. D., et al. 2014a, *ApJ*, **789**, 147

Weisz, D. R., Johnson, B. D., & Conroy, C. 2014b, *ApJL*, **794**, L3

Westmeier, T., Deg, N., Spekkens, K., et al. 2022, *PASA*, **39**, e058

Wheeler, C., Hopkins, P. F., Pace, A. B., et al. 2019, *MNRAS*, **490**, 4447