

Alma Mater Studiorum - Università di Bologna

Department of Physics and Astronomy
Second-cycle degree in Astrophysics and Cosmology

**Internal Structure and Environmental
Dependence of Subhaloes in AIDA–TNG**

Master thesis

Presented by:
**Anna Mariia
Okeanova**

Supervisor:
Prof. Giulia Despali
Co-supervisor:
Prof. Lauro Moscardini

Academic year 2024-2025

Abstract

This thesis investigates how the structure and abundance of dark-matter subhaloes differ between collisionless cold dark matter (CDM) and velocity-dependent self-interacting dark matter (vSIDM) in the AIDA–TNG simulation suite, a set of cosmological simulations that combines alternative dark-matter models with state-of-the-art galaxy formation physics. The comparison is carried out in both dark-matter-only (DMO) and full-physics hydrodynamical (HYDRO) runs in order to assess which signatures are most plausibly linked to dark-matter microphysics and which are altered by baryonic physics and host environment.

The analysis uses subhalo structural and demographic diagnostics, including the catalogue-based (R_{\max}, V_{\max}) relation, stacked dark-matter density and circular-velocity profiles, cumulative radial distributions, and the subhalo mass function (SHMF). This comparison is motivated by the fact that self-interacting dark matter has been proposed as a possible explanation for several small-scale tensions between CDM and astrophysical observations, many of which are linked to the structure and survival of satellite systems.

The strongest CDM–vSIDM differences are found in the density profiles. In the DMO runs, lower-mass vSIDM subhaloes are typically less centrally dense than their CDM counterparts over the resolved radial range, while the corresponding circular-velocity profiles show the same ordering with smaller amplitude. By contrast, the catalogue-based (R_{\max}, V_{\max}) relation shows only weak and sign-changing offsets. In HYDRO, the structural trend is not obviously erased, but it becomes substantially harder to isolate because baryons broaden the diversity of the surviving subhalo population. At the population level, cumulative radial distributions show only weak model dependence, indicating that they are less sensitive to CDM–vSIDM differences than internal-structure diagnostics. The SHMF shows a modest low-mass excess of vSIDM relative to CDM when subhaloes are selected by present-day bound mass, although this is difficult to interpret uniquely because baryonic effects, altered stripping, and redistribution between mass bins remain degenerate in this statistic. In contrast, warm dark matter models show a clearer demographic signature through suppression of low-mass subhaloes. The results also show that environment is a major confounding variable in HYDRO: subhaloes at smaller present-day host-centric distances are systematically denser than those at larger distances in both CDM and vSIDM, indicating that environmental processing can produce structural differences qualitatively comparable to the residual model-to-model offset.

Overall, the results suggest that subhalo internal structure is the most sensitive probe of CDM versus vSIDM in the present dataset, but that baryonic physics, selection effects, and host environment must be controlled carefully before such differences can be interpreted as clean signatures of dark-matter microphysics.

Contents

1	Dark matter	2
1.1	Evidence and the Standard Picture	2
1.2	Particle Candidates and Their Limits	3
1.3	Small-Scale Challenges to Λ CDM	4
1.4	Primordial power spectrum and small-scale structure	7
1.5	Dark matter models	8
1.5.1	Cold Dark Matter (CDM)	8
1.5.2	Self-Interacting Dark Matter (SIDM1)	8
1.5.3	Velocity-Dependent Self-Interacting Dark Matter (vSIDM)	8
1.5.4	Warm Dark Matter (WDM3)	10
2	Cosmological simulations	13
2.1	From analog models to cosmological N-body simulations	13
2.2	From gravity-only to hydrodynamic simulations	15
2.3	Identifying structure: FoF and SUBFIND	15
2.4	The AIDA-TNG suite	17
2.4.1	Effect of baryonic physics on subhalo populations	18
2.5	Thesis scope and structure	19
3	Internal Structure and Abundance of Subhaloes in AIDA-TNG	22
3.1	Data Products, Definitions, and Analysis Strategy	22
3.1.1	Halo and subhalo catalogues	22
3.1.2	Diagnostic quantities	23
3.1.3	DMO versus HYDRO comparison strategy	24
3.1.4	Sample binning and profile stacking	25
3.2	Internal Structure: Density Profiles, Circular Velocity, and the (R_{\max}, V_{\max}) Plane	26
3.3	Subhalo Mass Function	28
4	Environmental Dependence of Subhalo Structure	39
4.1	Subhalo Antibias and Orbital Environment	39
4.2	Cumulative Radial Distributions	40
4.3	Distance-Dependent Density Profiles in the HYDRO Runs	45
4.3.1	CDM HYDRO	46
4.3.2	vSIDM HYDRO	48
4.3.3	Environmental variation versus model variation	49

5	Discussion and Conclusions	51
5.1	Summary of the main results	51
5.2	Interpretation	53
5.3	Limitations of the present analysis	55
5.4	Future directions	55
5.5	Final remarks	56

Chapter 1

Dark matter

1.1 Evidence and the Standard Picture

Comprising approximately 85% of the matter density of the Universe, dark matter is invisible in the electromagnetic range and can only be detected through its gravitational influence on visible matter. The gravitational evidence for dark matter is compelling precisely because it emerges independently across a wide range of astrophysical systems and scales.

In disk galaxies, observed rotation curves remain nearly flat at large radii, well beyond the extent of the luminous stellar disk, first systematically documented by Rubin et al. (1980) and Bosma (1981). This implies that the enclosed mass scales approximately as $M(< r) \propto r$, requiring an extended non-luminous mass component — the dark matter halo — to provide the additional gravitational pull. On larger scales, galaxy groups and clusters show even more striking evidence: galaxy velocity dispersions within clusters, first noted by (Zwicky, 1933), combined with weak and strong gravitational lensing measurements, consistently demand total masses that far exceed the sum of stars and hot X-ray-emitting gas. Gravitational lensing is especially powerful here because it measures mass directly through the deflection of light, independent of any assumptions about the dynamical state or composition of the cluster, making it one of the cleanest probes of the total matter distribution.

At cosmological scales, the Cosmic Microwave Background (CMB) serves as a cornerstone of evidence for dark matter. Observations of CMB temperature anisotropies show that at the time of recombination ($z \sim 1100$), density fluctuations were extremely small, of order $\delta \approx 10^{-5}$. If the Universe consisted only of baryonic matter, these fluctuations would have grown too slowly to produce the non-linear structures observed today. Dark matter, being non-electromagnetically coupled, decoupled from the cosmic fluid much earlier and began growing its density perturbations long before baryons did; once baryons decoupled, they were gravitationally pulled into these pre-existing potential wells, allowing structures to reach the observed non-linear amplitudes. Precision measurements by the Planck satellite confirm that dark matter accounts for approximately 26% of the total energy density, while baryons contribute only about 5% Planck Collaboration et al. (2020).

The Λ CDM model, which assumes dark matter is cold (non-relativistic at decoupling) and collisionless, has proven remarkably successful on large scales ($\gtrsim 1$ Mpc), reproducing the CMB angular power spectrum, the clustering of galaxies, the

cosmic web of filaments and voids, and the overall expansion history of the Universe. However, on sub-galactic and dwarf-galaxy scales, a number of persistent discrepancies between CDM predictions and observations have motivated the exploration of alternative dark matter models.

1.2 Particle Candidates and Their Limits

The leading particle physics candidate for cold dark matter has long been the Weakly Interacting Massive Particle (WIMP). WIMPs are stable particles with masses typically in the range of 10 GeV to a few TeV and interaction strengths on the order of the weak force. Their enduring appeal stems from the so-called “WIMP miracle”: such particles, if they existed in thermal equilibrium in the early Universe, would naturally freeze out with a relic density remarkably close to the observed dark matter abundance. This theoretical convergence, coupled with their promising detectability through direct scattering off nuclei, indirect annihilation signals, and production at colliders, established WIMPs as the dominant paradigm in dark matter research for over two decades.

However, direct-detection experiments have so far returned only non-detections and increasingly stringent upper limits. For example, the latest multi-tonne noble-liquid searches such as LZ [McCarthy & Lux-Zeplin Team \(2024\)](#) report no statistically significant WIMP signal, translating into world-leading exclusions over a broad mass range under standard halo assumptions. Axions and axion-like particles provide a conceptually different candidate, with searches like ADMX probing new mass intervals with sensitivity reaching well-motivated QCD-axion benchmarks, yet similarly reporting no detection to date.

Crucially, the absence of a direct-detection signal does not rule out astrophysically relevant self-interaction cross-sections: in SIDM and ν SIDM frameworks, the strength of dark-sector self-scattering can be largely independent of the coupling that controls interactions with Standard Model targets. This motivates a complementary approach — modifying the microphysics of dark matter itself — and shifts the most direct empirical constraints toward gravitational and dynamical probes such as lensing substructure, perturbations in stellar streams, and the internal structure of low-mass haloes.

Beyond WIMPs, theoretical and observational considerations motivate two broad classes of alternative. *Warm dark matter* (WDM) replaces the cold, non-relativistic particle with a lighter thermal relic whose residual free-streaming velocity erases density perturbations below a characteristic scale, suppressing the formation of low-mass haloes and reducing their central concentrations. *Self-interacting dark matter* (SIDM) retains the cold, massive character of CDM but introduces elastic scattering between dark matter particles via a dark-sector force, redistributing mass within halo interiors and producing flat, isothermal density cores. A velocity-dependent generalisation (ν SIDM) makes the cross-section a decreasing function of relative collision speed, reconciling the strong self-interactions required to produce cores in dwarf spheroidals with the much weaker interactions permitted by cluster-scale lensing observations. The microphysical implementation of each model and its specific parameters are described in detail in Section 1.5; here we introduce them conceptually to frame the small-scale tensions discussed below.

This thesis investigates the differences between different dark matter models by analyzing subhalo populations in the AIDA-TNG simulation suite Despali et al. (2025), which provides matched runs across multiple dark-matter models at consistent resolution and with both dark-matter-only (DMO) and hydrodynamic (HYDRO) variants. The central goal is to determine whether microphysics-driven differences in subhalo structure and abundance persist once baryonic physics is included, and to identify which diagnostics are most robust for disentangling the two effects.

1.3 Small-Scale Challenges to Λ CDM

The small-scale challenges to Λ CDM can be grouped into two broad categories: problems with the *abundance* of low-mass structures, and problems with their *internal structure*. While baryonic physics can partially address some of these tensions, no purely baryonic solution has proven fully satisfactory across all mass scales, which motivates the exploration of alternative dark matter models.

The **missing satellites problem** refers to the large excess of low-mass subhaloes predicted by CDM simulations relative to the number of observed satellite galaxies around the Milky Way Klypin et al. (1999). CDM predicts thousands of gravitationally bound subhaloes within the virial volume, yet even after the discovery of ultra-faint dwarf galaxies through surveys such as SDSS and DES brought the known census of Milky Way satellites to approximately 50, the discrepancy remains stark Nadler et al. (2021). The most natural Λ CDM explanation is that galaxy formation becomes increasingly inefficient in low-mass haloes, leaving many predicted subhaloes effectively starless Bullock & Boylan-Kolchin (2017). Reionisation can suppress gas accretion into shallow potential wells, and abundance matching arguments suggest that the observed satellite population is consistent with CDM if only a small fraction of low-mass haloes form stars. However, this explanation shifts rather than resolves the problem: it requires a large population of completely dark subhaloes whose existence remains unconfirmed by direct observation.

The **too-big-to-fail problem** Boylan-Kolchin et al. (2011) poses a more direct challenge. The most massive subhaloes predicted by CDM simulations are systematically too centrally dense to host the brightest observed Milky Way satellite galaxies: the stellar kinematics of classical dwarfs imply central masses substantially lower than the CDM prediction for subhaloes of the same peak circular velocity. These massive subhaloes are too large to have remained dark — they should have formed stars — yet no observed counterparts match their predicted density profiles. Baryonic feedback is largely ineffective in this regime because the stellar masses of the affected dwarfs ($M_* \lesssim 10^6 M_\odot$) are too low to drive the outflows needed to significantly alter the dark matter distribution Bullock & Boylan-Kolchin (2017).

On the observational side, satellite samples are incomplete because of limited sky coverage and surface-brightness selection, and inferred inner masses from stellar kinematics can carry substantial uncertainties, particularly for the faintest dwarfs Lovell et al. (2014).

The **cusp-core problem** arises because CDM simulations robustly produce cuspy inner density profiles scaling approximately as $\rho \propto r^{-1}$, while kinematic observations of many dwarf galaxies favour flatter, approximately constant-density cores. Supernova-driven feedback can generate cores in galaxies with sufficient star

formation: FIRE simulations Wetzel et al. (2016) show that repeated gas outflows can flatten the central dark-matter profile, but only above a threshold stellar mass of $M_* \approx 10^6 M_\odot$, corresponding to a halo mass of $M_{\text{vir}} \approx 10^{10} M_\odot$ under standard abundance-matching assumptions. Below this threshold there is simply not enough star-formation energy to significantly alter the cuspy NFW profile, leaving the cusp–core problem open for the faintest dwarfs.


This issue is closely related to the velocity scale at which the too-big-to-fail problem is most acute, namely systems with $V_{\text{max}} \sim 30\text{--}40 \text{ km s}^{-1}$, where CDM-only simulations predict subhaloes that are substantially denser than the classical Milky Way dwarf satellites. Because baryonic feedback is expected to be inefficient in this low-stellar-mass regime, the persistence of low inferred central densities remains a significant challenge for collisionless CDM and helps motivate alternatives such as self-interacting dark matter, which can generate cores without relying on stellar feedback alone. In addition, the observed diversity of central densities among Milky Way satellites may indicate that a single universal inner profile is insufficient, although such interpretations remain sensitive to modelling assumptions and observational uncertainties Read et al. (2018).

The **rotation-curve diversity problem** Oman et al. (2015) highlights the large observed spread in dwarf galaxy rotation curve shapes at fixed halo mass. Kaplinghat et al. (2014) quantified this as approximately a factor of four in central velocity $V_{2\text{kpc}}$ at a fixed maximum rotation velocity of $V_{\text{max}} \sim 70\text{--}80 \text{ km s}^{-1}$. This spread is difficult to reproduce in CDM simulations even when baryonic effects are included: models that generate cores through feedback (such as NIHAO, Wang et al. (2015)) can produce slowly rising rotation curves but struggle to simultaneously reproduce the compact, steeply rising curves seen in other galaxies, and no single baryonic model has reproduced the full observed range.


Recent work argues that part of the observed rotation-curve diversity may be recoverable with revised baryonic feedback prescriptions within Λ CDM, although modified dark-matter microphysics remains an alternative explanation in the literature Cruz et al. (2025).

Taken together, these problems share a common thread: they are most severe in the lowest-mass systems, precisely where baryonic feedback is least effective. This points toward a modification of dark matter microphysics as a complementary — or in some regimes necessary — solution. Velocity-dependent self-interacting dark matter naturally addresses the structural diversity problems by producing large cores in some systems while allowing gravothermal core collapse to super-cuspy states in others, providing a unified explanation for both diffuse satellites and anomalously dense structures such as the GD-1 stellar stream perturber and certain strong-lensing subhaloes Nadler et al. (2025). The SIDM framework does not directly reduce the total number of predicted subhaloes, so it does not resolve the missing satellites problem on its own; however, combined with standard galaxy formation inefficiency arguments, it addresses the structural mismatch that purely baryonic models struggle to explain Bullock & Boylan-Kolchin (2017). These small-scale tensions motivate using dark-matter substructure as a direct testbed for alternative dark-matter physics. Subhaloes can be constrained not only through satellite galaxy counts, but also through probes that do not require the subhalo to host visible stars, such as strong gravitational lensing (e.g. flux-ratio anomalies and arc perturbations) and

Classical small-scale challenges to collisionless CDM: Satellite abundance, inner density structure, and central mass of dwarf systems



Panel A: Missing satellites
The tension is strongest at the low-mass end, where predictions are very sensitive to the faint-end slope



Panel B: Cusp-core
The observed rotation curves rise slower than the cuspy NFW expectation

Panel C: Too-big-to-fail
Even the largest subhalos in CDM can look too dense compared to observed dwarf spheroidals.
Not only a satellite issue; similar tension also appears for some field dwarfs

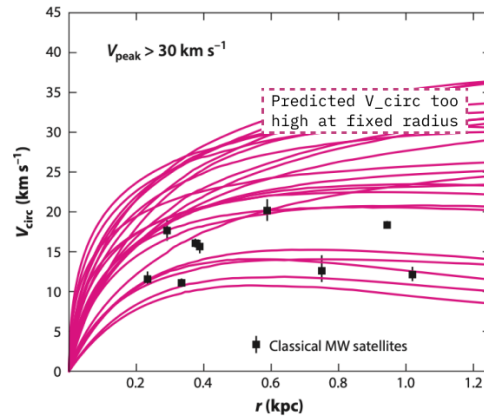


Figure 1.1: Illustration of the three classical small-scale tensions of collisionless cold dark matter (CDM): the missing-satellites problem, the cusp-core problem, and the too-big-to-fail problem. Compared with observations of dwarf galaxies, CDM generically predicts a larger abundance of faint subhaloes, steeper inner density profiles, and larger circular velocities at fixed radius. Adapted from Bullock & Boylan-Kolchin (2017).

perturbations of cold stellar streams in the Milky Way. Because these observables respond to both the abundance and the internal structure of subhaloes, they provide a natural bridge between particle-physics models and measurable non-linear signatures. Subhalo populations are therefore a uniquely powerful laboratory for several interconnected reasons. First, the total number and mass function of subhaloes directly encodes the microphysics of dark matter: warm dark matter free-streaming suppresses small-scale power and produces a characteristic truncation in the subhalo mass function, while self-interactions alter internal structure without necessarily changing abundance, allowing the two effects to be disentangled observationally. Second, because different baryonic feedback prescriptions — stellar winds, supernovae, and AGN-driven outflows — modify the subhalo population in ways that can mimic or mask these microphysical signatures, the subhalo census simultaneously serves as a discriminant between galaxy formation models. Third, gravitational lensing and stellar stream perturbations extend this discriminating power to the regime of completely dark subhaloes, probing the underlying matter distribution in a way that is insensitive to whether star formation ever occurred. Together, these threads explain why the missing satellites discrepancy has never been fully resolved by baryonic arguments alone: confirming or ruling out a large population of dark, low-mass subhaloes requires precisely the kind of multi-probe, microphysics-aware analysis that this thesis pursues.

1.4 Primordial power spectrum and small-scale structure

In the Λ CDM model, cosmic structure is seeded by primordial adiabatic fluctuations and grows by gravitational instability in an expanding background Bullock & Boylan-Kolchin (2017). A compact description of the initial conditions is provided by the primordial power spectrum $P(k)$ as a function of comoving wavenumber k , which is nearly scale invariant and is often approximated by a power law $P(k) \propto k^{n_s}$ with $n_s \simeq 1$ Bullock & Boylan-Kolchin (2017). CMB observations constrain this spectrum precisely on large scales; the final full-mission Planck analysis finds a scalar spectral index $n_s = 0.965 \pm 0.004$ and reports no evidence for deviations from a purely power-law primordial spectrum Planck Collaboration et al. (2020).

As these initial overdensities grow, regions that become non-linear eventually stop expanding with the Hubble flow, turn around, and collapse, converting potential energy into kinetic energy in the process. The outcome is a population of virialized dark-matter haloes and subhaloes whose abundance and internal structure (e.g. $\rho(r)$, $V_{\text{circ}}(r)$, and the $R_{\text{max}}-V_{\text{max}}$ relation) encode both their assembly history and, potentially, the underlying dark-matter microphysics.

Although the Λ CDM paradigm is well tested from horizon scales down to approximately megaparsec scales, observations below ~ 1 Mpc are more challenging to interpret Bullock & Boylan-Kolchin (2017). At these sub-galactic scales, both the limited direct observational access and the hard-to-model impact of baryonic physics complicate first-principles predictions and the interpretation of apparent tensions. This motivates treating the small-scale matter distribution in simulations as a sensitive arena for probing both baryonic effects (via DMO versus HYDRO

comparisons) and dark-matter microphysics (via CDM versus vSIDM/WDM comparisons).

1.5 Dark matter models

This thesis compares four dark matter models available in the AIDA-TNG suite: CDM, SIDM1, vSIDM, and WDM3. All models share identical initial conditions and baryonic physics within each box, so any differences in subhalo properties can be attributed directly to the dark matter microphysics. We describe each model in turn.

1.5.1 Cold Dark Matter (CDM)

CDM serves as the baseline for all comparisons in this thesis. Dark matter particles are assumed to be cold (non-relativistic at decoupling), collisionless, and interacting only through gravity. The CDM model produces cuspy NFW-like density profiles Navarro et al. (1997) and a hierarchical subhalo mass function with a steep low-mass slope ($dN/dM \propto M^{-\alpha}$, $\alpha \simeq 1.9$). All model-dependent signatures discussed in this thesis are measured as deviations from this CDM baseline.

1.5.2 Self-Interacting Dark Matter (SIDM1)

SIDM1 assumes purely elastic, isotropic scattering between dark matter particles with a constant cross-section per unit mass $\sigma/m = 1 \text{ cm}^2 \text{ g}^{-1}$ Spergel & Steinhardt (2000). Interactions are implemented via a Monte Carlo approach in which colliding particle pairs exchange momenta with a randomly drawn scattering direction. The constant cross-section means that all haloes, regardless of mass or internal velocity, experience the same interaction rate per unit density. This produces isothermal cores in all systems, with more massive haloes developing larger and deeper cores due to their higher central densities and longer interaction times Despali et al. (2025). While SIDM1 alleviates the cusp-core and too-big-to-fail problems at dwarf scales, its mass-independent nature leads to tension with cluster-scale observations, which require the effective cross-section to be substantially smaller at high velocities. SIDM1 is available in the L50 DMO configuration and is included in this thesis as a reference point for comparing against the more physically motivated velocity-dependent model.

1.5.3 Velocity-Dependent Self-Interacting Dark Matter (vSIDM)

The vSIDM model addresses the tension between dwarf and cluster constraints by making the cross-section a function of the relative scattering velocity. The implementation follows the prescription of Correa (2021), in which the differential cross-section takes the Yukawa-type form

$$\frac{d\sigma}{d\Omega} = \frac{\sigma_0}{\left(1 + \left(\frac{v}{w}\right)^2 \sin^2\left(\frac{\theta}{2}\right)\right)^2}, \quad (1.1)$$

where v is the relative collision velocity, w is a characteristic turnover velocity separating the high- and low-interaction regimes, σ_0 is the cross-section amplitude, and θ is the scattering angle. The velocity dependence arises physically from a Yukawa-type interaction mediated by a new dark-sector gauge boson, which naturally suppresses the cross-section at high relative velocities. In practice, the effective cross-section is computed as the viscosity cross-section σ_V , which accurately reproduces the results of full angular-dependent calculations at reduced computational cost.

The parameters are calibrated to satisfy observational constraints across mass scales simultaneously Despali et al. (2025): at low velocities characteristic of dwarf spheroidals ($v \sim 10 \text{ km s}^{-1}$) the effective cross-section reaches $\sigma/m \sim 100 \text{ cm}^2 \text{ g}^{-1}$, sufficient to produce large cores and alleviate the cusp-core and too-big-to-fail problems; at cluster scales ($v \sim 1000 \text{ km s}^{-1}$), it drops below $0.1 \text{ cm}^2 \text{ g}^{-1}$, consistent with lensing arc statistics and cluster halo shape constraints. The vSIDM cross-section equals that of SIDM1 at relative velocities of approximately 300 km s^{-1} .

This transition from core formation to collapse-like behaviour is also seen in dedicated high-resolution SIDM suites spanning LMC-scale haloes up to clusters, which show that the structural outcome depends jointly on the velocity-dependent cross-section and the halo assembly history. Such suites motivate using both internal-structure diagnostics (e.g. $\rho_{\text{DM}}(r)$, $V_{\text{circ}}(r)$, $R_{\text{max}}-V_{\text{max}}$) and population-level statistics (radial survival and subhalo abundances) when attempting to separate microphysics-driven effects from baryonic processing.

This mass-dependent interaction rate produces a rich diversity of structural outcomes. Low-mass subhaloes with low internal velocities experience strong scattering, which initially drives core formation but can eventually trigger gravothermal core collapse if the interaction rate is sufficiently high and the subhalo has enough time to evolve. In this collapsed state, the inner halo becomes denser than the CDM prediction, forming a so-called super-cusp. Gravothermal collapse does not occur in all low-mass subhaloes within a Hubble time: the core-collapsed fraction peaks at roughly 50–60% for subhaloes near $V_{\text{peak}} \approx 60 \text{ km s}^{-1}$, with the remainder remaining in the core-forming phase Nadler et al. (2025). High-mass haloes, by contrast, experience negligible self-interactions at their characteristic velocities and retain profiles close to CDM. The simultaneous presence of core-forming and collapse-candidate subhaloes in the same population produces a characteristic diversity in density profiles and circular velocity curves, which is one of the central observational signatures tested in this thesis.

Unlike WDM, self-interacting dark matter does not alter the primordial perturbations: the initial power is unchanged in SIDM, and therefore any difference relative to CDM is generated by self-interactions during the subsequent evolution of structures Despali et al. (2025). Equivalently, the SIDM simulations start from the same initial conditions as the CDM ones, with self-interactions becoming relevant only during the late non-linear evolution Despali et al. (2025). This distinction is crucial for interpretation: if CDM and vSIDM differ in the inner density or $V_{\text{circ}}(r)$

at fixed host and subhalo mass, the origin must lie in gravothermal evolution and/or baryonic processing, not in a different initial power spectrum.

1.5.4 Warm Dark Matter (WDM3)

WDM3 assumes a thermal relic dark matter particle with a mass of 3 keV. Unlike SIDM models, WDM affects structure not through interactions between particles but through their non-negligible thermal velocities at decoupling, which erase density perturbations below a characteristic free-streaming length. For a 3 keV thermal relic, the free-streaming length is approximately $\lambda_{\text{fs}} \approx 20.7 \text{ kpc}$, and the power spectrum is suppressed by 50% at the half-mode mass $M_{\text{hm}} \approx 3.26 \times 10^8 h^{-1} M_{\odot}$ (Despali et al. (2025)).

Warm dark matter modifies structure formation already at the level of the initial conditions by suppressing power below a characteristic scale. In the AIDA–TNG implementation, this is quantified directly from the ratio of WDM and CDM power spectra: one can define the half-mode mass M_{hm} as the halo mass corresponding to the scale where the power spectrum is suppressed by one half, $P_{\text{WDM}}(k)/P_{\text{CDM}}(k) = |T(k)|^2 = 0.5$ (Despali et al. (2025)). As illustrated by the input spectra at $z = 127$ (Fig. 2, left panel), warmer models exhibit a cutoff at increasingly larger scales (smaller k), implying a stronger suppression of low-mass structure (Despali et al. (2025)). Because the central densities of dark-matter haloes reflect the density of the Universe at the time of their formation, reduced small-scale power delays assembly and therefore leads to lower concentrations and shallower central potentials at fixed mass in the mass range relevant for subhaloes near the free-streaming scale (Bullock & Boylan-Kolchin (2017)). Operationally, this motivates using internal-structure diagnostics such as stacked $V_{\text{circ}}(r)$ and the $R_{\text{max}}-V_{\text{max}}$ relation, where reduced concentration manifests as larger R_{max} at fixed V_{max} (or lower V_{max} at fixed mass).

Physically, WDM modifies structure formation through free streaming: residual thermal velocities erase primordial perturbations below a characteristic scale, which can be expressed via a transfer function $P_{\text{WDM}}(k) = T^2(k)P_{\text{CDM}}(k)$. As a result, WDM primarily affects the low-mass end: halo abundances are suppressed and surviving haloes are less concentrated, producing lower central densities. In practice, strong core-like flattening is expected only near/below the suppression scale; at Milky-Way and higher masses, profiles remain close to cuspy CDM-like shapes. This suppression reduces the total number of low-mass subhaloes relative to CDM and lowers the concentration of surviving haloes, resulting in shallower central potentials and lower central densities. Unlike SIDM, WDM3 does not produce large constant-density cores at Milky Way mass scales; its structural effects are most pronounced near and below the half-mode suppression mass.

Warm dark matter and expected internal-structure trends. In warm dark matter (WDM), residual thermal velocities produce free streaming, which erases primordial perturbations below a characteristic scale and suppresses structure formation at the low-mass end. This is commonly expressed via a transfer-function suppression of the linear power spectrum,

$$P_{\text{WDM}}(k) = T^2(k) P_{\text{CDM}}(k), \quad T(k) = (1 + (\alpha k)^{2\nu})^{-5/\nu}, \quad (1.2)$$

where α sets the cutoff scale (larger α corresponds to a warmer model), and $\nu \simeq 1$. Because haloes near the cutoff form later (when the mean cosmic density is lower), WDM haloes are expected to be less concentrated than in CDM, with the strongest deviations confined to the lowest masses.

For WDM models in the $m_{\text{WDM}} \sim \text{keV}$ range, literature results indicate that core-like flattening is primarily a low-mass phenomenon: strong cores can appear below $\sim 10^{11} M_{\odot}$ in very warm cases (e.g. WDM1), while in intermediate models such as WDM3 cores are expected only at the smallest resolved masses. At the high-mass end ($\gtrsim 10^{12} M_{\odot}$), density profiles remain close to cuspy CDM-like shapes. Quantitatively, reported trends include lower central densities by factors of a few in the lowest-mass subhaloes, lower V_{max} at fixed mass, and larger r_{max} at fixed V_{max} (reflecting reduced concentration).

In this thesis, WDM3 is examined in the context of the L100 box $R_{\text{max}}-V_{\text{max}}$ analysis as a structural reference point. However, it is not carried forward into the main profile and demographic comparisons, as recent work suggests that vSIDM provides a more physically complete and observationally consistent description of small-scale structure across mass scales Despali et al. (2025). WDM3 therefore appears in this thesis primarily to illustrate the contrast between a free-streaming suppression mechanism and a self-interaction mechanism, rather than as a primary model under investigation.

A practical caveat in WDM runs is artificial fragmentation, which can generate spurious low-mass haloes in filaments. Where WDM is used in this thesis (mainly as a structural reference in the $R_{\text{max}}-V_{\text{max}}$ plane), we therefore interpret the lowest-mass objects conservatively and rely on resolution/completeness cuts rather than treating the extreme low-mass tail as physically robust.

Figure 1.2 illustrates the structural differences between models directly. In DMO runs (top row), vSIDM and SIDM1 both produce a systematic suppression of the inner density relative to CDM, with vSIDM showing a clear mass dependence — the effect is strongest in low-mass haloes where internal velocities are low and the effective cross-section is largest. WDM3 and WDM1 show a qualitatively different pattern: suppression that is most pronounced at low masses and near the free-streaming scale, but with less diversity across the population. Crucially, the full-physics rows (bottom) show that baryonic processes substantially reduce these model-to-model differences, particularly at intermediate and large radii, underscoring the need for a controlled DMO–HYDRO comparison strategy to disentangle microphysics from baryonic effects.




Figure 1.2: Ratio of the dark matter density profile in each alternative model to the CDM baseline (ρ/ρ_{CDM}) as a function of physical radius r , from Despali et al. (2025). Top and bottom rows show dark-matter-only (DMO) and full-physics (FP) runs respectively; columns correspond to vSIDM, SIDM1, WDM3, and WDM1 (AIDA simulations). Line colours indicate host halo mass, ranging from $\log_{10}(M/M_{\odot}) = 9.5$ (dark purple) to 14.0 (cyan). The vertical dotted line marks the numerical convergence radius. The mass dependence of vSIDM and the baryonic suppression of model differences in the FP rows motivate the analysis strategy of this thesis.

Chapter 2

Cosmological simulations

Understanding the formation and evolution of cosmic structure requires more than observations alone. The gravitational collapse of dark matter, the hydrodynamics of baryonic gas, and the feedback from stars and black holes are coupled non-linear processes that cannot be followed analytically beyond the linear regime. Cosmological simulations provide the essential connection: by numerically evolving a set of initial conditions forward in time under well-defined physical assumptions, they allow direct and controlled comparison between theoretical predictions and observed structures across a wide range of scales.

Their importance becomes even clearer when different dark-matter models are considered. In this case, one is not only interested in the global growth of structure, but also in how changes in dark-matter microphysics propagate into the internal structure of haloes and into the baryonic components of galaxies. Figure 2.1 illustrates this point for a Milky-Way-mass halo, comparing the projected dark-matter, gas, and stellar distributions in CDM, SIDM1, vSIDM, and WDM3. Even at fixed halo mass, the models produce visible differences in the abundance of substructure, the gas morphology, and the shape of the stellar component, highlighting the need for cosmological simulations that consistently evolve both dark matter and baryons.

2.1 From analog models to cosmological N-body simulations

The history of numerical simulation in astrophysics begins with a remarkably creative analog experiment. Holmberg (1941) modelled the tidal encounter of two stellar systems using light bulbs as mass elements, exploiting the fact that light intensity follows the same inverse-square law as Newtonian gravity. Each galaxy was represented by 37 bulbs, and the gravitational force was measured with a photocell rather than computed mathematically. Despite its simplicity, the experiment produced meaningful results: Holmberg found that tidal deformations during the encounter increase the mutual gravitational attraction between the systems, and that the resulting energy loss can lead to capture — with spiral arm formation emerging as a natural by-product. This work established the conceptual template for what would become a central tool of modern astrophysics.

As electronic computers became available, true N-body methods emerged in

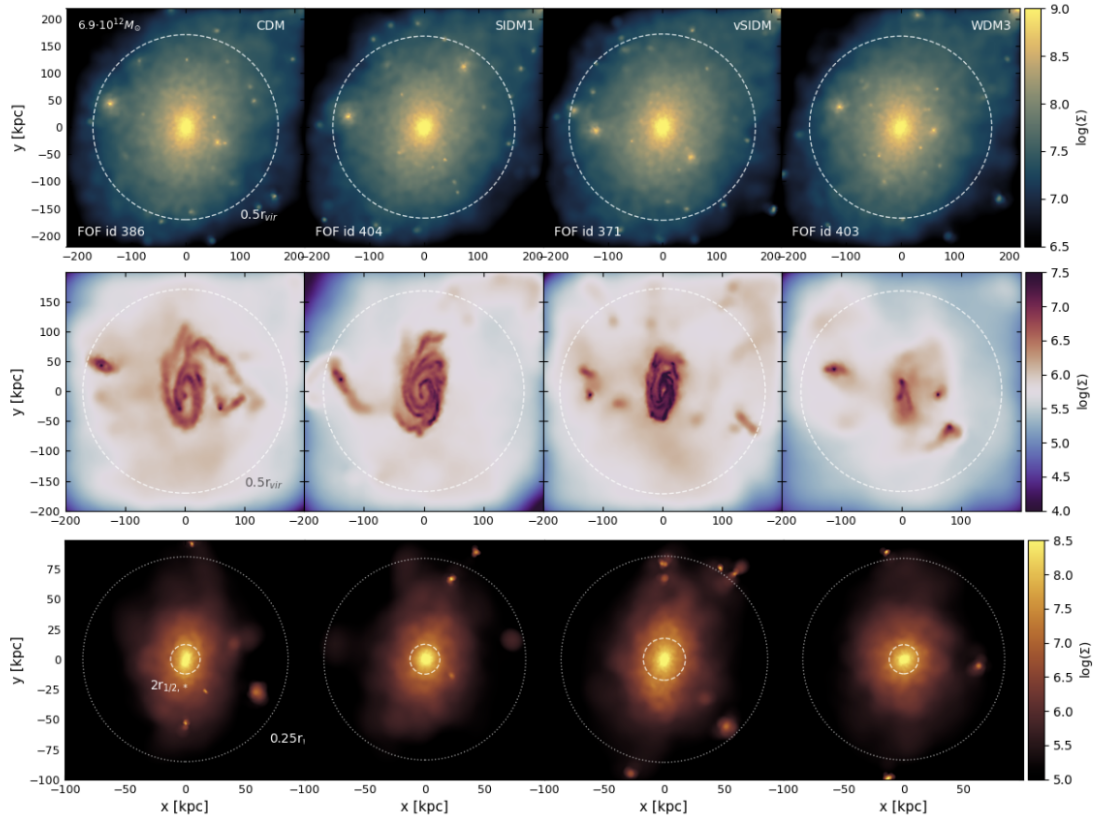


Figure 2.1: Visualisation of the projected mass distributions for a Milky-Way-mass halo ($M_{\text{vir}} = 6.9 \times 10^{12} M_{\odot}$ in CDM at $z = 0$) in the 100/A runs, adapted from Despali et al. (2025). From left to right, the columns show CDM, SIDM1, vSIDM, and WDM3. The top, middle, and bottom rows correspond to dark matter, gas, and stars, respectively. The figure is centred on the inner halo, extending to $\sim 0.5 r_{\text{vir}}$ in the top and middle rows and to $\sim 0.25 r_{\text{vir}}$ in the stellar row; in the latter, the dashed inner circle marks twice the stellar half-mass radius. The comparison shows a reduced abundance of substructure in WDM3 and visible differences in the gas dynamics and stellar morphology across dark-matter models.

which each simulation particle represents a large ensemble of matter and the gravitational force on each particle is computed by summing contributions from all others. The direct summation approach, refined by Aarseth (1985) through the implementation of fourth-order interpolation schemes to maximise the integration timestep, requires of order $N(N-1)/2$ pair evaluations, giving a computational cost that scales as $\mathcal{O}(N^2)$. Simulations of this era were typically limited to a few thousand particles, which was sufficient for studying collisional stellar dynamics but still not enough for cosmological volumes where $N \sim 10^8-10^{10}$.

The critical algorithmic breakthrough came with the hierarchical tree method of Barnes & Hut (1986), which organised particles into a recursive octree structure and approximated the gravitational influence of distant groups via multipole expansions. This approach, which also provided well-defined upper bounds on the force error at each evaluation, reduced the scaling to $\mathcal{O}(N \log N)$ and became the standard for cosmological N-body work, making large-volume simulations computationally

tractable for the first time.

2.2 From gravity-only to hydrodynamic simulations

The early successes of pure gravity simulations in reproducing the large-scale cosmic web came with an important caveat: gravity alone is not sufficient to model the observed Universe at small scales. White & Rees (1978) identified this limitation concretely, showing that early simulations suffered from a cooling catastrophe and failed to reproduce both the long survival times of galaxies in cluster environments and the correct slope of observed matter density profiles. Their key insight was that radiative cooling is the essential physical mechanism allowing gas to condense at the centres of dark matter haloes, increasing the binding energy of the baryonic component and enabling galaxies to survive the tidal forces of their host environment. This work made clear that a realistic description of galaxy formation requires modelling gas physics alongside gravity.

Dark-matter-only (DMO) simulations, in which collisionless particles evolve under gravity alone, nonetheless remained the primary tool for establishing the broad framework of Λ CDM structure formation Springel et al. (2005). They demonstrated how small primordial perturbations grow hierarchically into the cosmic web of filaments, sheets, and voids, while comparisons between their predictions and observations helped bring into focus key small-scale tensions — including the missing satellites and cusp–core problems discussed in Section 1.3 — that motivate the alternative dark matter models considered in this thesis Klypin et al. (1999).

Resolving these tensions, however, requires modelling the complex physics of baryonic matter: gas cooling, star formation, supernova feedback, and black hole growth. None of these processes operate at the resolution scales accessible to cosmological simulations, so modern hydrodynamic (HYDRO) runs rely on *subgrid models* — parametrized prescriptions that approximate the net effect of small-scale physics on the simulated volume.

Distinguishing between dark matter models adds a further requirement: the non-linear interplay between self-interactions, tidal stripping, and baryonic feedback cannot be captured analytically beyond the perturbative regime. Cosmological simulations in which the dark matter model is varied while all other physics is held fixed therefore provide the most direct and controlled route to identifying model-dependent signatures in observable subhalo populations, and they form the methodological foundation of this thesis.

2.3 Identifying structure: FoF and SUBFIND

Extracting physically meaningful objects from a simulation requires dedicated structure-finding algorithms. The standard approach in modern cosmological simulations proceeds in two steps.

The first step applies a Friends-of-Friends (FoF) algorithm, which links particles into groups by connecting any two particles separated by less than a linking length




Figure 2.2: Comparison between observed galaxy redshift surveys and mock galaxy catalogues constructed from the Millennium Simulation. The figure illustrates that the Λ CDM framework reproduces the large-scale cosmic web of filaments, walls, clusters, and voids seen in observations. The observational wedges include the CfA2 Great Wall, the Sloan Great Wall in SDSS, and the 2dFGRS, while the lower and right panels show mock surveys generated from the Millennium dark-matter simulation with matched survey geometry and magnitude limits. From Springel et al. (2006).

b. Davis et al. (1985) standardised the choice of $b = 0.2$ times the mean inter-particle separation, demonstrating that this value identifies groups with an average overdensity comparable to that of a virialized object in a Λ CDM universe. Because FoF is a topological finder, it makes no prior assumptions about halo geometry, making it well suited to the complex non-spherical structures that arise naturally in the cosmic web.

The second step applies the SUBFIND algorithm of Springel et al. (2001) to identify gravitationally self-bound substructures within each FoF group. SUBFIND uses an excursion-set method based on adaptive kernel density estimates: it identifies density peaks as subhalo candidates and locates saddle points in the global density field to separate overlapping regions. Crucially, each candidate then undergoes an iterative unbinding procedure using a tree-based potential calculation, retaining only particles that are genuinely gravitationally bound. The power of this approach was demonstrated in the Aquarius Project Springel et al. (2008), where SUBFIND successfully resolved nearly 300,000 bound subhaloes within a single Milky Way-sized halo and detected up to four generations of sub-substructure. Together, FoF and SUBFIND define the halo and subhalo catalogues on which all subsequent analysis

in this thesis is based.

An important subtlety is that the inclusion of baryons modifies the gravitational potential in ways that affect both subhalo evolution and survival. Gas cooling can deepen central potential wells, while stellar and AGN feedback can redistribute matter and alter the tidal evolution of satellites. As a result, hydrodynamical (HYDRO) simulations do not produce halo and subhalo populations identical to their dark-matter-only (DMO) counterparts even at fixed numerical resolution.

This distinction is physically important and is well established in the literature: baryons can modify both the abundance of haloes and, more importantly for this thesis, the abundance and survival of subhaloes relative to DMO predictions. It is therefore not sufficient to compare different dark-matter models in isolation. Instead, the impact of dark-matter microphysics must be assessed separately in DMO and HYDRO runs, so that changes driven by baryons can be distinguished from those caused by the dark matter model itself.

2.4 The AIDA-TNG suite

The simulations used in this thesis are part of the AIDA-TNG (Alternative Dark Matter in the TNG universe) project [Despali et al. \(2025\)](#), which provides a controlled set of cosmological runs designed to disentangle the effects of alternative dark matter physics from baryonic processes. All runs share the same IllustrisTNG galaxy formation framework, including radiative cooling with metallicity and redshift dependence, stochastic star formation, stellar winds and supernova feedback, and supermassive black hole growth and feedback in both thermal (quasar-mode) and kinetic (low-state) modes [Pillepich et al. \(2018\)](#).

The suite spans three cosmological volumes at different resolutions. All AIDA-TNG volumes follow the standard IllustrisTNG integration from $z = 127$ to $z = 0$. In particular, the L50 box has a comoving side length of 51.7 Mpc (matching TNG50-2 initial conditions), and we use the $z = 0$ snapshot as the primary analysis epoch throughout this thesis. The key numerical parameters for the flagship high-resolution runs are summarised in Table 2.1.

Table 2.1: Numerical parameters for the AIDA-TNG high-resolution runs used in this thesis. m_{DM} and m_{bar} are the dark matter and baryonic particle masses respectively; ϵ is the gravitational softening length. Values are taken from [Despali et al. \(2025\)](#).

Box	Side length	$m_{\text{DM}} (M_{\odot})$	$m_{\text{bar}} (M_{\odot})$	ϵ (kpc)
L100	110.7 Mpc	6.0×10^7	1.1×10^7	1.48
L50	51.7 Mpc	3.6×10^6	6.8×10^5	0.57
L25	25.1 Mpc	$\sim 1 \times 10^5$	—	—

Each volume is run in both DMO and full-physics (HYDRO) configurations and across multiple dark matter models: CDM, WDM (with particle masses of 1, 3, and 5 keV), SIDM1 (constant cross-section $\sigma/m = 1 \text{ cm}^2 \text{ g}^{-1}$), and vSIDM (velocity-dependent, following [Correa \(2021\)](#)). The initial conditions for the L100 and L50

boxes match those of the original TNG100-2 and TNG50-2 simulations respectively, enabling direct comparison with the broader TNG literature.

This thesis focuses primarily on the L50 and L25 boxes, which provide the best balance between sample size and mass resolution for subhalo internal structure studies. The L100 box offers the largest statistical sample and is used where high-mass host statistics are needed, while L50 and L25 allow reliable profiling down to smaller subhalo masses. Profile measurements are trusted only at radii $r \geq 3\epsilon_{\text{DM}}$ to avoid artificially softened cores, and subhaloes are required to contain at least 1000 dark matter particles for reliable structural measurements. This corresponds to approximate minimum resolved subhalo masses of $4.3 \times 10^9 M_{\odot}$ in the L50 box and $7.1 \times 10^{10} M_{\odot}$ in the L100 box. The primary analysis snapshot is $z = 0$ throughout, except where noted for the SIDM1 L50 run where the closest available snapshot at $z \approx 0.2$ is used; based on the stability of model-dependent signatures at low redshift Despali et al. (2025), this is not expected to introduce significant systematic differences. The central dark matter model comparison in this thesis is CDM versus ν SIDM, with SIDM1 and WDM3 included where available as additional reference points.

AIDA-TNG therefore provides a controlled testbed in which the same galaxy-formation model is applied across multiple dark-matter microphysics variants. A key result of the suite is that the baseline TNG model produces broadly reasonable galaxy populations across the considered dark-matter scenarios, and that the overall baryonic suppression of substructure is similar across models, which motivates using matched DMO–HYDRO pairs as a practical strategy for separating baryonic effects from dark-matter microphysics.

2.4.1 Effect of baryonic physics on subhalo populations

Baryonic processes alter subhalo populations through four main channels. Reionization suppresses star formation in low-mass haloes ($M \lesssim 10^9 M_{\odot}$), leaving them as effectively dark, star-free objects. Gas cooling and adiabatic contraction concentrate baryons — and, through the deepened potential well, dark matter — in halo centres, increasing the central concentration of massive subhaloes that host stars and making them more resistant to tidal disruption. Stellar and AGN feedback partially counteracts this by heating and expelling gas, reducing halo masses and shifting the overall mass function; the magnitude of this effect is strongly model-dependent. Finally, the presence of a baryonic disc in the host enhances tidal stripping of infalling subhaloes, particularly those on low-pericentre orbits, further depleting the inner subhalo population.

The net outcome is a systematic suppression of the subhalo mass function (SHMF) relative to dark-matter-only (DMO) runs. Comparing hydrodynamic simulations to their DMO counterparts, Despali & Vegetti (2017) find a reduction of order 20% in EAGLE and up to 40% in Illustris at $m_{\text{sub}} \sim 10^8\text{--}10^{10} M_{\odot}$, with the depletion most severe at the low-mass end. This suppression is accompanied by a shallowing of the SHMF slope α (defined via $dN/dM_{\text{sub}} \propto M_{\text{sub}}^{-\alpha}$): from $\alpha \simeq 0.90$ in DMO to $\alpha \simeq 0.85$ in EAGLE and $\alpha \simeq 0.76$ in Illustris, reflecting a relative reduction of low-mass objects compared to high-mass ones. The amplitude of the effect thus encodes both

the efficiency of baryonic disruption and the strength of the feedback prescription adopted.

These baryonic signatures must be carefully separated from any microphysics-driven differences between CDM and vSIDM. Throughout this thesis, the DMO–HYDRO comparison within each dark-matter model serves precisely this purpose: differences that are present in DMO and persist with the same sign in HYDRO are more plausibly associated with dark-matter microphysics, while differences that appear only in HYDRO, or change sign between the two, are more plausibly linked to baryonic processing.

2.5 Thesis scope and structure

The central aim of this thesis is to determine how the internal structure and demographic properties of dark-matter subhaloes differ between collisionless cold dark matter (CDM) and alternative dark-matter models, and to assess how robust those differences remain once baryonic physics is included. The analysis uses the AIDA–TNG simulation suite [Despali et al. \(2025\)](#) as its primary numerical framework, with main emphasis on the comparison between CDM and velocity-dependent self-interacting dark matter (vSIDM). Among the available non-standard models in the suite, vSIDM provides the most relevant test case for the present work because it allows structurally non-trivial departures from CDM while remaining more phenomenologically viable than more extreme self-interaction scenarios. SIDM1 is included as a useful reference where available, while WDM3 is used primarily as a contrasting example of how suppression of small-scale structure by free streaming differs from the effects of self-interactions; neither is treated as a primary focus of the thesis.

The comparison is carried out in both dark-matter-only (DMO) and full-physics hydrodynamical (HYDRO) simulations. The DMO runs provide the cleanest setting in which to identify signatures of dark-matter microphysics, whereas the HYDRO runs test whether the same trends remain detectable once gas cooling, star formation, and feedback are included. This distinction is central to the interpretation of the results. A structural difference that appears in DMO and remains visible with the same sign in HYDRO is more plausibly linked to the dark-matter model itself, whereas a trend that weakens strongly, disappears, or reverses once baryons are included is more likely to reflect the combined effects of baryonic processing, altered survival, and sample re-selection. The DMO–HYDRO comparison is therefore not a secondary consistency check, but a central part of the strategy used to disentangle dark-matter microphysics from baryonic and environmental effects.

The thesis is organised as follows. Chapter 2 introduces the simulations, definitions, and diagnostic quantities used throughout the analysis. Chapter 3 presents the main structural and demographic results, beginning with catalogue-based structural proxies and stacked dark-matter density and circular-velocity profiles, and then turning to the subhalo mass function as a complementary population-level diagnostic. Chapter 4 examines the role of orbital environment, first through the radial distribution of subhaloes within their hosts and then through distance-dependent structural trends in the HYDRO runs. The final chapter synthesises these results, discusses their main limitations, and summarises the implications for using subhalo

structure as a probe of dark-matter physics.




Figure 2.3: Halo and global subhalo mass functions in the Illustris and EAGLE simulations, comparing dark-matter-only and hydrodynamical runs across several redshifts. The inclusion of baryons changes the predicted abundance of structures, particularly at the subhalo level, illustrating that DMO and HYDRO catalogues cannot be interpreted interchangeably. This motivates the separate analysis of DMO and HYDRO simulations throughout this thesis. From Despali & Vegetti (2017).

Chapter 3

Internal Structure and Abundance of Subhaloes in AIDA–TNG

Roadmap of the chapter. This chapter compares the internal structure and demographic properties of subhaloes across the AIDA–TNG dark-matter models, with primary emphasis on CDM and velocity-dependent SIDM (vSIDM), and with matched comparisons between dark-matter-only (DMO) and full-physics (HYDRO) runs. Throughout, host haloes are characterised by M_{200c} and R_{200c} , subhalo mass refers to the present-day bound mass `SubhaloMass` unless stated otherwise, and the central subhalo of each FoF group is excluded from the subhalo samples. Density and circular-velocity profiles are analysed in physical kpc.

The chapter is organised from compact structural diagnostics to broader population statistics. We first examine the catalogue-based (V_{\max}, R_{\max}) relation and then turn to stacked dark-matter density and circular-velocity profiles, which allow any structural differences to be localised in radius. We then ask whether these structural trends are accompanied by changes in the abundance of surviving subhaloes through the subhalo mass function (SHMF). Unless noted otherwise, comparisons are made at fixed present-day bound subhalo mass, so the results should be interpreted as differences among surviving subhaloes rather than as direct progenitor-matched comparisons.

3.1 Data Products, Definitions, and Analysis Strategy

The analysis uses the AIDA–TNG simulation suite of Despali et al. (2025), which provides matched DMO and HYDRO realizations for four dark matter models: CDM, SIDM1, vSIDM, and WDM3 (a 3 keV thermal relic). This section summarises only the quantities and analysis choices that are directly relevant for interpreting the results shown below.

3.1.1 Halo and subhalo catalogues

We analyse snapshot 99 of the L25, L50, and L100 simulation boxes (comoving side lengths 25.1, 51.7, and 110.7 cMpc), corresponding to redshift $z = 0$. The adopted

cosmology is Planck-like, with $h = 0.6774$, $\Omega_m = 0.3089$, and $\Omega_\Lambda = 0.6911$. The three volumes play complementary roles in the analysis. L100 provides the largest statistical sample for the SHMF; L50 is the main box used for the host-binned structural comparisons; and L25, which has the highest mass and force resolution, is used to assess how far the profile-based CDM–vSIDM trends remain visible when the accessible radial range is extended.

The dark-matter gravitational softening lengths at $z = 0$ are $\epsilon_{\text{DM}} = 0.295$ kpc (L25), 0.57 kpc (L50), and 1.48 kpc (L100), with corresponding dark-matter particle masses $m_{\text{DM}} = 9.5 \times 10^5 M_\odot$, $4.3 \times 10^6 M_\odot$, and $7.1 \times 10^7 M_\odot$. Unless stated otherwise, profile measurements are restricted to $r \geq 3\epsilon_{\text{DM}}$, corresponding to inner cuts of 0.89, 1.71, and 4.44 kpc in L25, L50, and L100, respectively. This cut avoids obvious softening-dominated artefacts, but it should still be regarded as a conservative practical limit rather than a full convergence proof for inner slopes in tidally processed subhaloes. For the SIDM1 L50 run, snapshot 99 is unavailable; we therefore use the nearest available low-redshift output at $z \approx 0.2$. Given the weak evolution of the relevant model differences over this interval reported by Despali et al. (2025), this substitution is not expected to dominate the systematic uncertainty, although it remains an additional source of model-specific uncertainty when SIDM1 is compared directly to the $z = 0$ runs.

Host haloes are identified with the Friends-of-Friends (FoF) algorithm and are characterised by their virial mass M_{200c} and virial radius R_{200c} , defined with respect to $200 \rho_c(z)$:

$$M_{200c} = \frac{4}{3}\pi R_{200c}^3 \cdot 200 \rho_c(z), \quad \rho_c(z) = \frac{3H^2(z)}{8\pi G}. \quad (3.1)$$

These quantities are taken directly from the group catalogues as `Group_M_Crit200` and `Group_R_Crit200`. Host-centric distances are computed from the three-dimensional separation between `SubhaloPos` and `GroupPos`.

Subhaloes are identified within each FoF group using SUBFIND. The central object of each FoF group (`GroupFirstSub`) is excluded from all subhalo samples. Unless otherwise stated, “subhalo mass” refers to `SubhaloMass`, i.e. the total bound mass assigned by SUBFIND at the analysed snapshot. We retain only subhaloes containing at least 1000 dark-matter particles, corresponding to approximate lower mass limits of $9.5 \times 10^8 M_\odot$ (L25), $4.3 \times 10^9 M_\odot$ (L50), and $7.1 \times 10^{10} M_\odot$ (L100). Because this chapter primarily uses present-day bound mass, comparisons should be interpreted as comparisons at fixed surviving subhalo mass rather than, for example, at fixed M_{peak} . This choice is important for interpretation: any model-dependent remapping induced by differential stripping is part of the measured signal in this chapter, rather than something removed by construction.

3.1.2 Diagnostic quantities

The internal structure of subhaloes is characterised using stacked dark-matter density profiles $\rho_{\text{dm}}(r)$ and circular-velocity profiles $V_{\text{circ}}(r)$:

$$V_{\text{circ}}(r) = \sqrt{\frac{G M(< r)}{r}}, \quad (3.2)$$

where $M(< r)$ is the enclosed dark-matter mass within physical radius r . Using only the dark-matter component in both DMO and HYDRO runs ensures that the same quantity is compared across physics types. Density profiles are measured in logarithmically spaced spherical shells centred on the subhalo potential minimum and are quoted in physical kpc throughout. Because $V_{\text{circ}}(r)$ depends on the enclosed mass, it is generally smoother and less sensitive to local shell-to-shell fluctuations than $\rho_{\text{dm}}(r)$; correspondingly, model differences are often easier to identify in density than in circular velocity.

A compact structural summary is provided by the pair $(V_{\text{max}}, R_{\text{max}})$, where V_{max} is taken directly from the catalogue field `SubhaloVmax` and R_{max} from the corresponding catalogue field `SubhaloVmaxRad`. In this chapter, the $(R_{\text{max}}, V_{\text{max}})$ analysis is therefore catalogue-based and should not be confused with any radius inferred separately from reconstructed $V_{\text{circ}}(r)$ profiles. At fixed V_{max} , larger R_{max} generally corresponds to a more extended characteristic scale, while smaller R_{max} corresponds to a more compact one.

At the population level, we consider the SHMF, written as the number of subhaloes per unit logarithmic mass interval.

3.1.3 DMO versus HYDRO comparison strategy

Each diagnostic is analysed in matched DMO and HYDRO runs. Microphysics effects are assessed by comparing vSIDM to CDM at fixed physics type, while baryonic effects are assessed by comparing HYDRO to DMO within the same dark-matter model. In practice, however, HYDRO does not simply “add baryons” to a fixed subhalo sample: it also changes subhalo survival, stripping, and detectability. Accordingly, differences between DMO and HYDRO are interpreted here as the effect of baryons on the measured signal, rather than as a perfectly clean causal decomposition.

This matched-comparison framework is motivated by previous DMO–HYDRO studies such as Despali & Vegetti (2017), who showed that baryons can suppress inner-halo substructure counts substantially relative to matched DMO runs. In the present analysis, the most important practical question is whether a CDM–vSIDM separation seen in DMO persists with the same sign in HYDRO, weakens in HYDRO, or disappears entirely. The structural diagnostics and the abundance diagnostics need not respond in the same way: a structural difference among individual surviving subhaloes can remain visible even when baryonic processing largely erases a corresponding demographic signal.

Because baryonic stripping is strongest toward the host centre, and because configuration-space subhalo finding becomes less complete against the dense host background in the inner halo, any differences seen at the smallest radii in HYDRO are treated as suggestive rather than decisive. The same caution applies to low- N sub-bins and to any interpretation that relies mainly on the most extreme individual profiles.

3.1.4 Sample binning and profile stacking

To control for the strong dependence of subhalo structure and survival on both host scale and subhalo scale, the analysis is organised in bins of host mass and subhalo mass. Host-mass binning limits environmental mixing, while subhalo-mass binning reduces the risk of interpreting ordinary mass-scaling trends as signatures of dark-matter microphysics. This is particularly important for vSIDM, whose effective interaction rate is expected to vary with the characteristic internal velocity of the system.

For the stacked profile analysis, we focus primarily on L50 and adopt two subhalo-mass bins,

- $9.5 < \log_{10}(M/M_{\odot}) < 10.5$,
- $10.5 < \log_{10}(M/M_{\odot}) < 11.5$,

where M denotes the present-day catalogue mass `SubhaloMass`, converted to physical units via $M = \text{SubhaloMass} \times 10^{10} h^{-1} M_{\odot}$. The relevant L50 DMO sample sizes are listed in Table 3.1. We distinguish between the number of mass-selected objects, N_{mass} , and the number that enter the final stacked profiles, N_{stack} , after requiring valid resolved profiles and interpolation onto a common radial grid. Since the profile stacks are therefore drawn from a subset of the mass-selected population, they should be interpreted as properties of the resolved-profile sample rather than of all subhaloes satisfying the mass cut.

A non-negligible selection effect enters the profile stacks: the fraction $N_{\text{stack}}/N_{\text{mass}}$ is substantially lower for vSIDM than for CDM in both subhalo-mass bins. The stacked profile comparison should therefore be interpreted as a comparison of resolved-profile subsamples rather than of the full mass-selected populations. This asymmetry may reflect numerical profile-validity requirements, genuinely more diffuse or irregular vSIDM structures, or both, but it is not disentangled further here. At the host-bin level, the retained samples remain sizeable in both DMO and HYDRO, so the increased difficulty of isolating a clean CDM–vSIDM separation in HYDRO cannot be attributed to sample size alone.

Table 3.1: Sample sizes for the L50 DMO stacked profile analysis. For each subhalo-mass bin and dark-matter model, N_{mass} is the number of mass-selected subhaloes and N_{stack} is the number entering the final stacked-profile sample after requiring valid resolved density profiles and interpolation onto the common radial grid.

Model	Mass bin $\log_{10}(M/M_{\odot})$	N_{mass}	N_{stack}
CDM	$9.5 < \log_{10} M < 10.5$	34219	4704
CDM	$10.5 < \log_{10} M < 11.5$	4668	827
vSIDM	$9.5 < \log_{10} M < 10.5$	44443	2000
vSIDM	$10.5 < \log_{10} M < 11.5$	5479	267

Within each host-mass and subhalo-mass bin, we construct stacked profiles using the arithmetic mean and quantify halo-to-halo variation through the 16th–84th

DMO: catalogue-based R_{\max} - V_{\max} relation for CDM and vSIDM
 Lower panels: $\Delta \log_{10} R_{\max} = \log_{10} R_{\max, \text{vSIDM}} - \log_{10} R_{\max, \text{CDM}}$

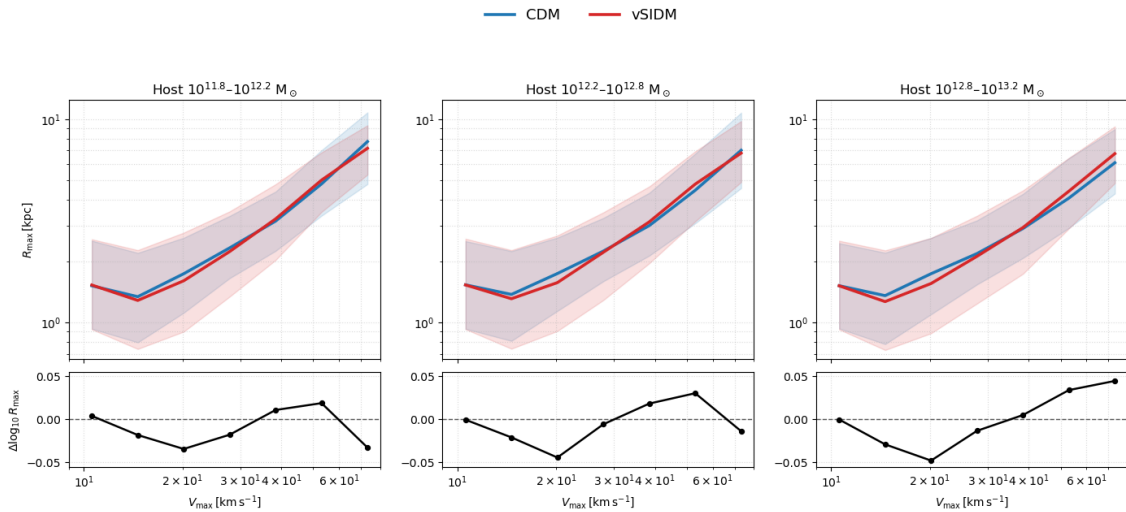


Figure 3.1: Catalogue-based R_{\max} - V_{\max} relation for CDM and vSIDM subhaloes in the DMO runs, split by host-mass bin. The upper panels show the median relations and 16th–84th percentile bands in common V_{\max} bins; the lower panels show the matched-bin offset $\Delta \log_{10} R_{\max} = \log_{10} R_{\max, \text{vSIDM}} - \log_{10} R_{\max, \text{CDM}}$. The offset is small and sign-changing, rather than a single global displacement between the two models.

percentile range. To illustrate the tails of the distribution, we additionally show the three most extreme individual profiles in each bin, identified by their largest deviation from the stacked mean at the innermost resolved radius. Because the arithmetic mean can be affected by skewed tails, especially in sparsely populated bins, differences between CDM and vSIDM are interpreted in the context of the full distribution width and treated cautiously whenever the separation of the mean profiles is comparable to the object-to-object scatter.

3.2 Internal Structure: Density Profiles, Circular Velocity, and the (R_{\max}, V_{\max}) Plane

We begin with the catalogue-based (V_{\max}, R_{\max}) plane, which provides a compact structural summary before turning to the radial profile diagnostics.

Figures 3.1 and 3.2 show the catalogue-based R_{\max} - V_{\max} relations for CDM and vSIDM in common V_{\max} bins, split by host mass, together with the offset

$$\Delta \log_{10} R_{\max} \equiv \log_{10} R_{\max, \text{vSIDM}} - \log_{10} R_{\max, \text{CDM}}. \quad (3.3)$$

This matched-bin construction is important: it allows the median relations to be compared at the same V_{\max} , avoiding the ambiguity that arises when the two models are binned separately.

In the DMO case, the catalogue-based relation does not show a single monotonic shift between CDM and vSIDM. Instead, the matched-bin offset remains

HYDRO: catalogue-based R_{\max} - V_{\max} relation for CDM and vSIDM

Lower panels: $\Delta \log_{10} R_{\max} = \log_{10} R_{\max, \text{vSIDM}} - \log_{10} R_{\max, \text{CDM}}$

— CDM — vSIDM




Figure 3.2: Same as Fig. 3.1, but for the HYDRO runs. The catalogue-based offset remains small in amplitude. In the middle and upper host-mass bins, the vSIDM median lies modestly above CDM at intermediate and high V_{\max} , while the low- V_{\max} offset is near zero or slightly negative.

small, of order a few $\times 10^{-2}$ dex, and changes sign with V_{\max} . In all three host-mass bins, $\Delta \log_{10} R_{\max}$ is mildly negative in part of the low- V_{\max} regime, typically around $V_{\max} \sim 15\text{--}25 \text{ km s}^{-1}$, and then rises toward zero or positive values at higher V_{\max} . A broadly similar sign-changing pattern is visible across the three host-mass bins, although the exact amplitude varies from panel to panel. The catalogue-based (R_{\max}, V_{\max}) plane therefore does not by itself support a strong claim that vSIDM subhaloes are uniformly more extended than CDM at fixed V_{\max} . Instead, it indicates only a weak, scale-dependent structural modulation.

The HYDRO catalogue-based relation is similarly weak in amplitude, but the pattern differs somewhat from DMO. In the middle and upper host-mass bins, the vSIDM median lies modestly above CDM at intermediate and high V_{\max} , while the low- V_{\max} offset is close to zero or slightly negative. This is not the simple low- V_{\max} positive displacement that might be expected if the dominant effect were a clean catalogue-level core-formation signature. In this dataset, the catalogue-based (R_{\max}, V_{\max}) plane therefore suggests only a mild scale-dependent modulation of structure, not a clean model-separated locus.

The radial profile comparison then makes it possible to localise where any structural separation arises. Figures 3.5, 3.6, and 3.7 show stacked DMO dark-matter density profiles $\rho_{\text{dm}}(r)$ and circular-velocity profiles $V_{\text{circ}}(r)$ for CDM and vSIDM, split by host-mass bin and subhalo-mass bin. In each panel, the solid curves show the stacked means, the shaded regions show the 16th–84th percentile range, and the dashed curves highlight the most extreme individual objects. Because the mean can be influenced by tails, especially in low- N bins, the interpretation below focuses on trends that persist over a resolved radial range and are not driven entirely by a few outliers.

A useful resolution cross-check is provided by Figs. 3.3 and 3.4, which compare

the same profile diagnostics between L50 and the higher-resolution L25 DMO runs. The role of this comparison is not to establish strict numerical convergence of the inner slopes, but to test whether the sign of the CDM–vSIDM difference remains visible when the accessible radial range is extended. In L25, the profiles reach substantially farther beyond the adopted inner cut than in L50, and the low-mass-bin separation between CDM and vSIDM is correspondingly easier to identify, particularly in $\rho_{\text{dm}}(r)$. In L50, by contrast, several panels probe only a relatively narrow radial range beyond the inner limit, so the box provides stronger support for the sign of the trend than for detailed statements about the exact inner profile shape.

In the DMO stacks, the clearest separation appears in the lower subhalo-mass bins, where the CDM mean density profile often lies above the vSIDM mean across the resolved inner radial range. The corresponding $V_{\text{circ}}(r)$ curves show the same ordering with smaller amplitude, as expected for an enclosed-mass diagnostic. At fixed present-day bound mass, this points to systematically lower central dark-matter densities in part of the surviving vSIDM population.

At higher subhalo masses the separation becomes less uniform. In some host bins the ordering remains visible, but in others the model offset is comparable to the intrinsic profile-to-profile scatter. The evidence for a robust population-level difference is therefore strongest in the lower-mass subhalo sample.

The tails of the distribution are wide in both models. In particular, some individual vSIDM objects reach inner densities comparable to or above the upper end of the CDM range in specific panels. This behaviour is physically suggestive, but by itself it is not sufficient to establish gravothermal core collapse. In the absence of time evolution, tail statistics, or a robust population-level reversal relative to CDM, these cases are interpreted here only as *collapse candidates* or examples of *collapse-like behaviour*. This caution is especially important in the sparsest bins, where the dashed curves can be dominated by one or two systems and where the innermost radii are also the most sensitive to convergence and subhalo-finder incompleteness.

Figures 3.8 and 3.9 show the corresponding density and circular-velocity profiles in the HYDRO runs. The broad pattern is similar to DMO, but with substantially broader scatter and less uniform panel-to-panel separation. In some host-mass and subhalo-mass bins, the CDM mean remains above vSIDM over part of the resolved radial range in both $\rho_{\text{dm}}(r)$ and $V_{\text{circ}}(r)$; in others, the two models overlap within the intrinsic distribution width.

The most defensible conclusion is therefore that the DMO structural trend is not cleanly erased in HYDRO, but it becomes much harder to isolate once baryonic physics is included. This reflects both broader structural diversity among surviving satellites and the fact that HYDRO re-selects the surviving population rather than simply perturbing a fixed DMO sample.

3.3 Subhalo Mass Function

We next examine the overall demographic abundance of subhaloes through the SHMF. The subhalo mass function (SHMF) provides a complementary population-level diagnostic to the structural measures discussed above and is used here to ask whether the different dark-matter models produce measurable changes in the abundance of surviving subhaloes as a function of present-day bound mass.

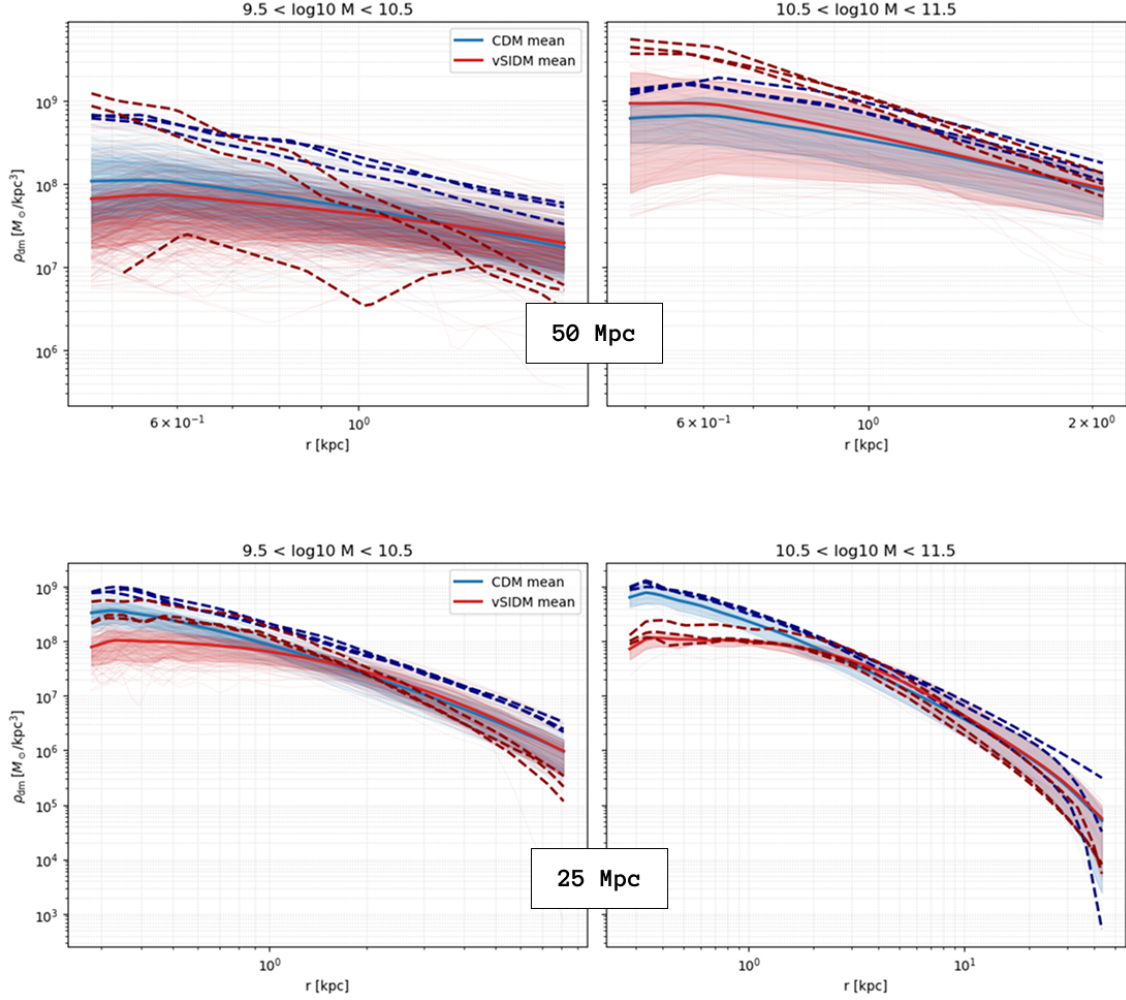


Figure 3.3: Stacked dark-matter density profiles $\rho_{\text{dm}}(r)$ for CDM (blue) and vSIDM (red) in the L50 (top row) and L25 (bottom row) DMO runs, for two subhalo-mass bins. Solid lines show the mean stacked profile; shaded bands indicate the 16th–84th percentile range; thin lines show individual profiles; dashed curves highlight the most extreme objects. The higher-resolution L25 runs extend to larger radii beyond the inner cut and show the low-mass CDM–vSIDM separation more clearly than L50.

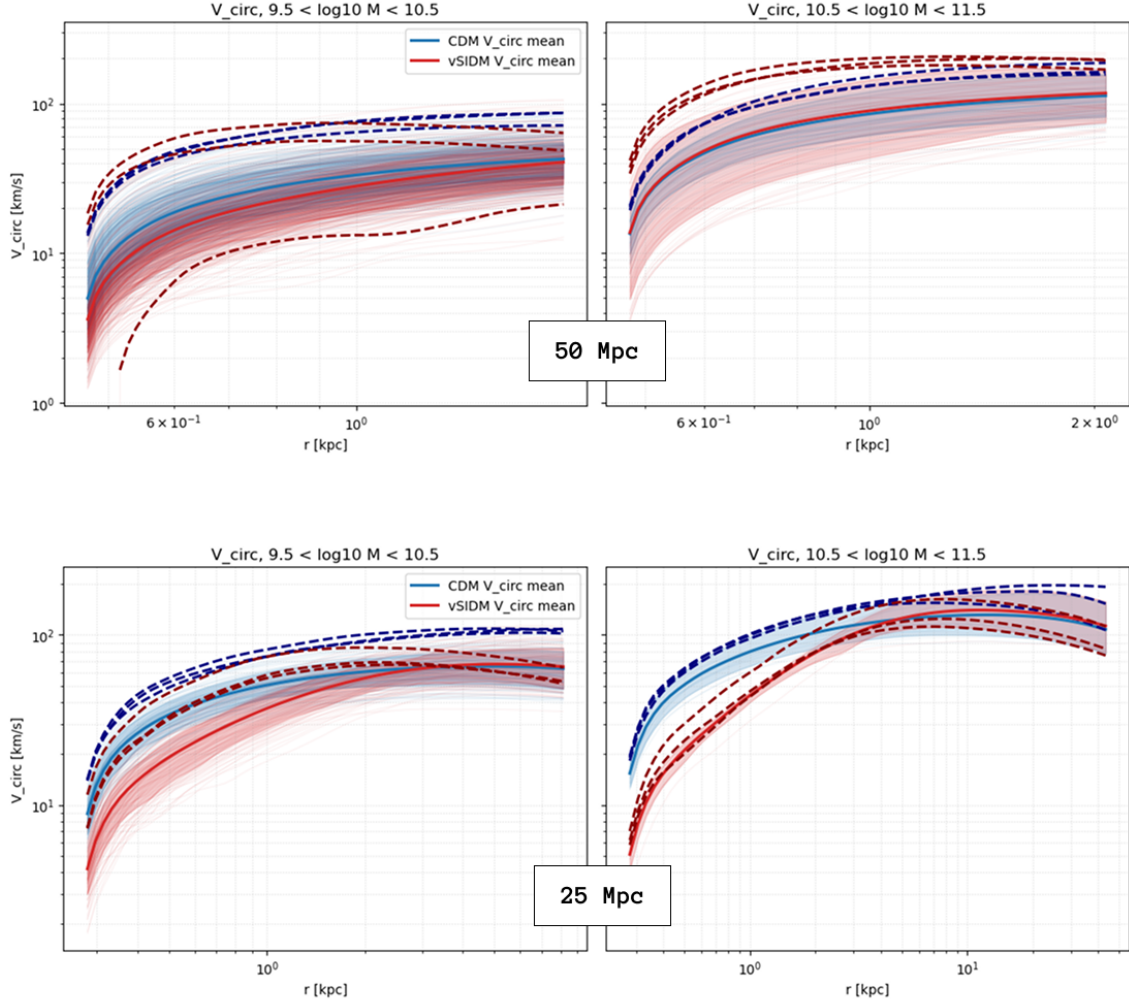


Figure 3.4: Stacked circular-velocity profiles $V_{\text{circ}}(r)$ for CDM (blue) and vSIDM (red) in the L50 (top row) and L25 (bottom row) DMO runs, for the same subhalo-mass bins as Fig. 3.3. The sign of the CDM–vSIDM offset is broadly consistent with the density-based comparison, but the separation is visibly more muted, as expected for an enclosed-mass diagnostic.

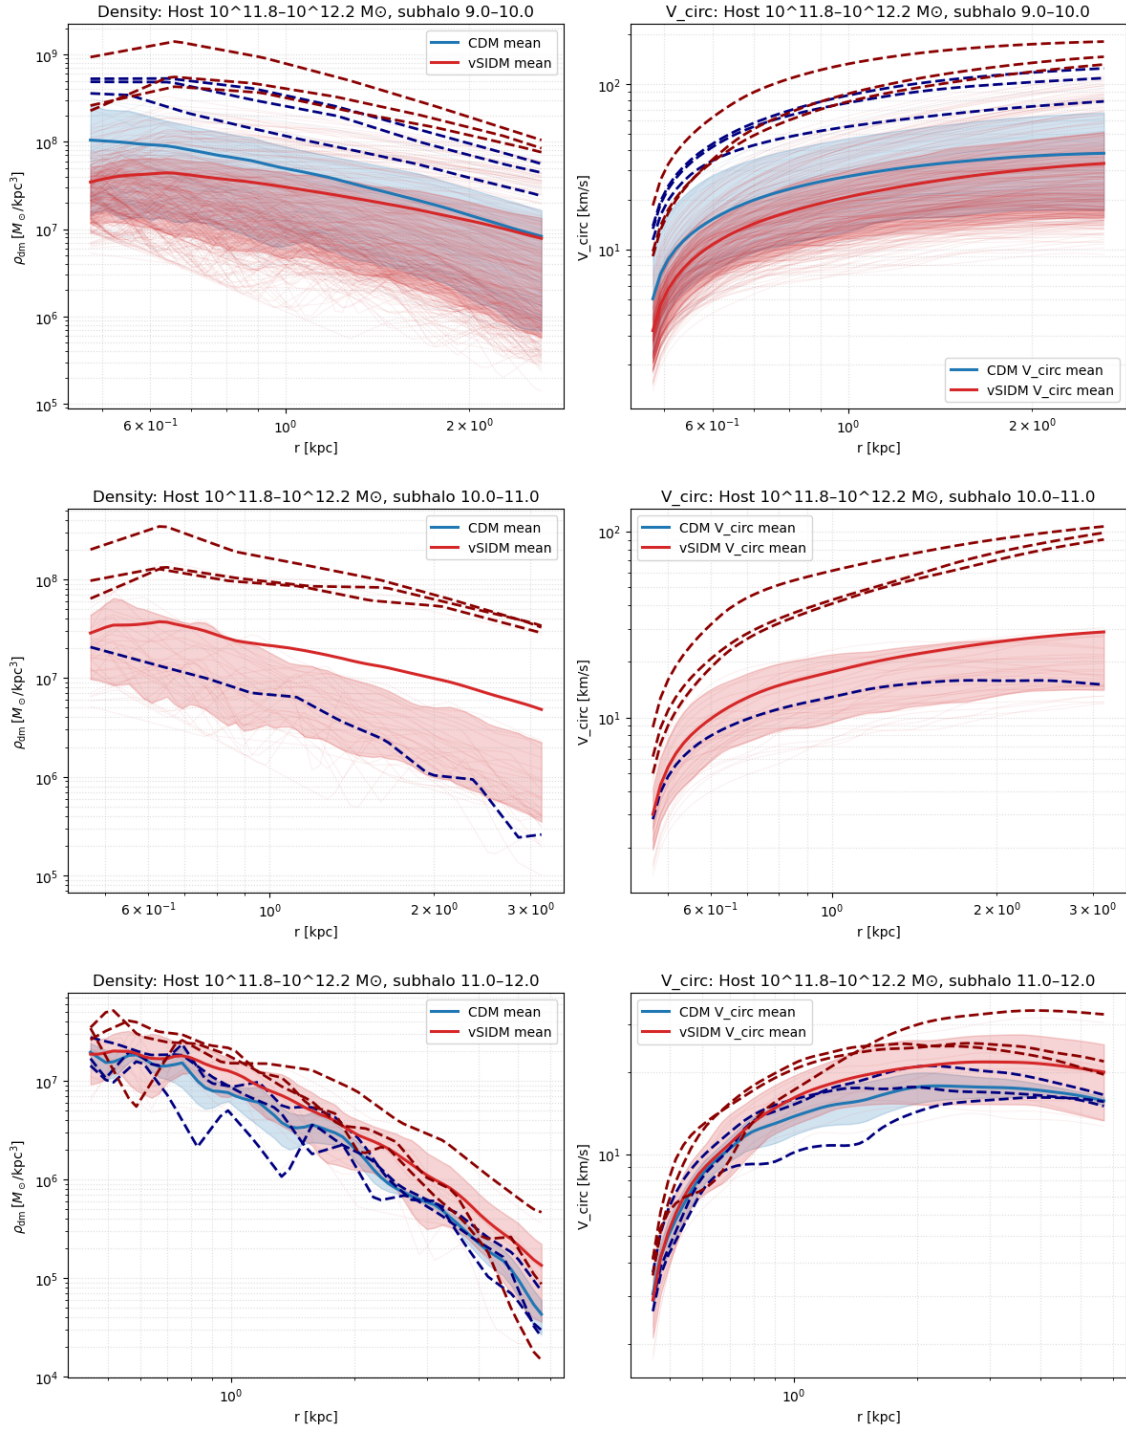


Figure 3.5: DMO density and circular-velocity profiles for host-mass bin $10^{11.8}$ – $10^{12.2} M_{\odot}$. Each row corresponds to a different subhalo-mass bin; left columns show $\rho_{\text{dm}}(r)$ and right columns show $V_{\text{circ}}(r)$ for CDM (blue) and vSIDM (red). Solid lines are mean stacked profiles, shaded bands show the 16th–84th percentile range, and dashed curves highlight the most extreme objects.

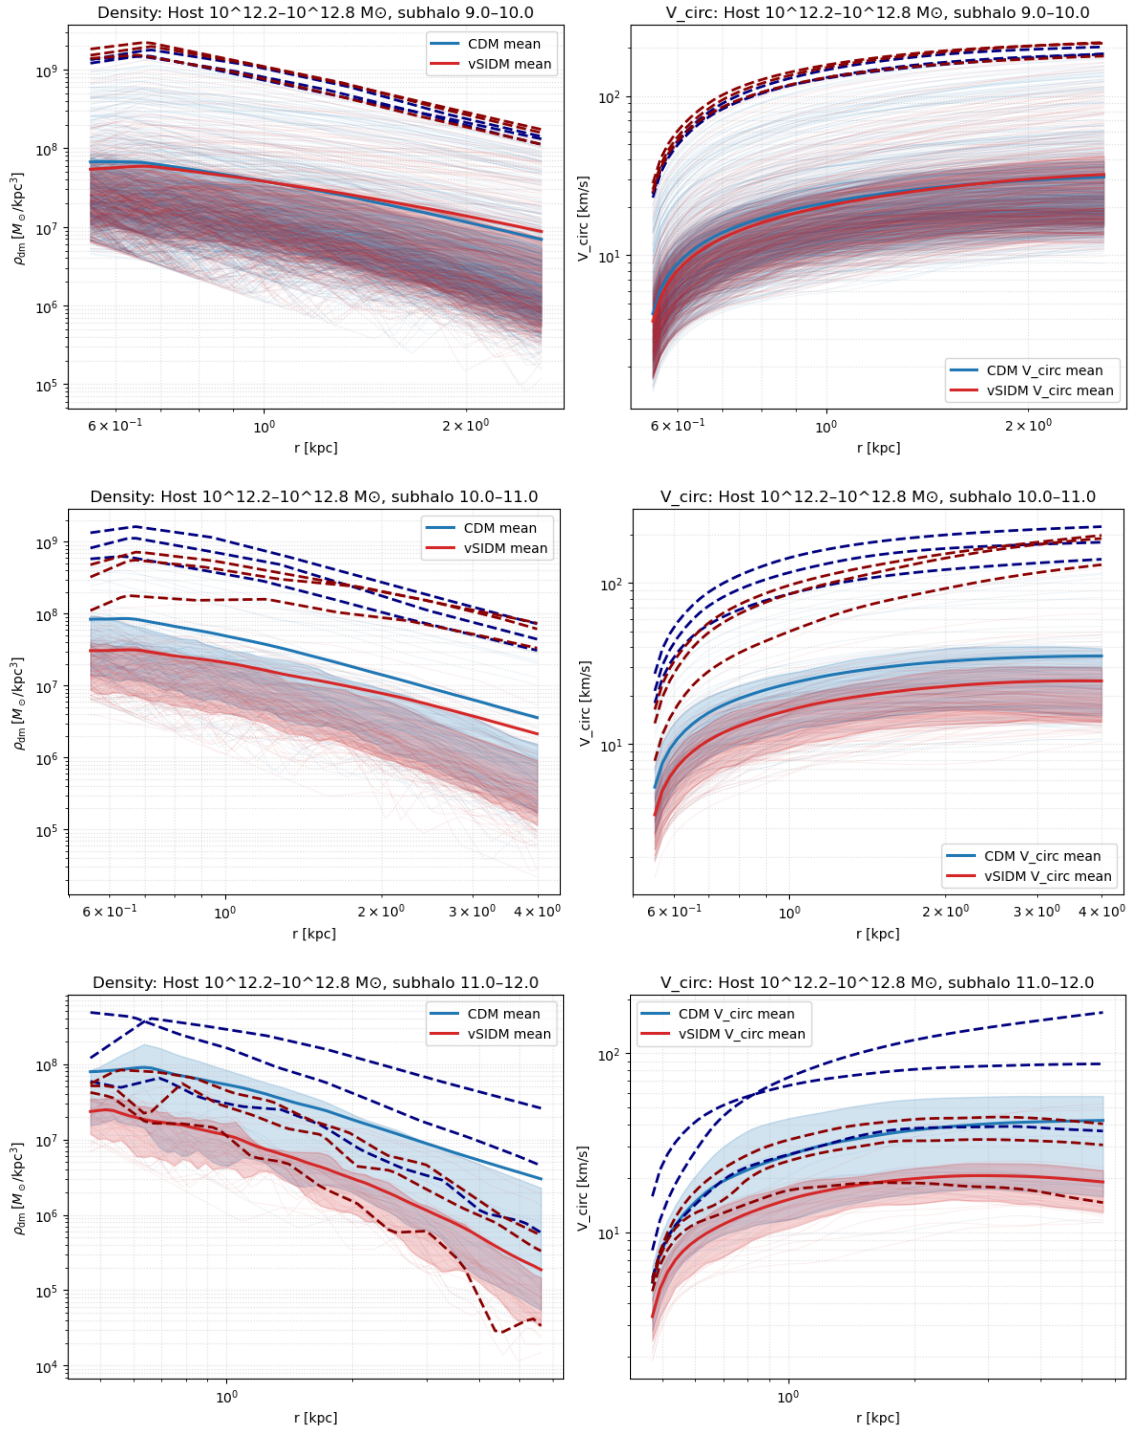


Figure 3.6: Same as Fig. 3.5, but for host-mass bin $10^{12.2}-10^{12.8} M_{\odot}$ (DMO runs).

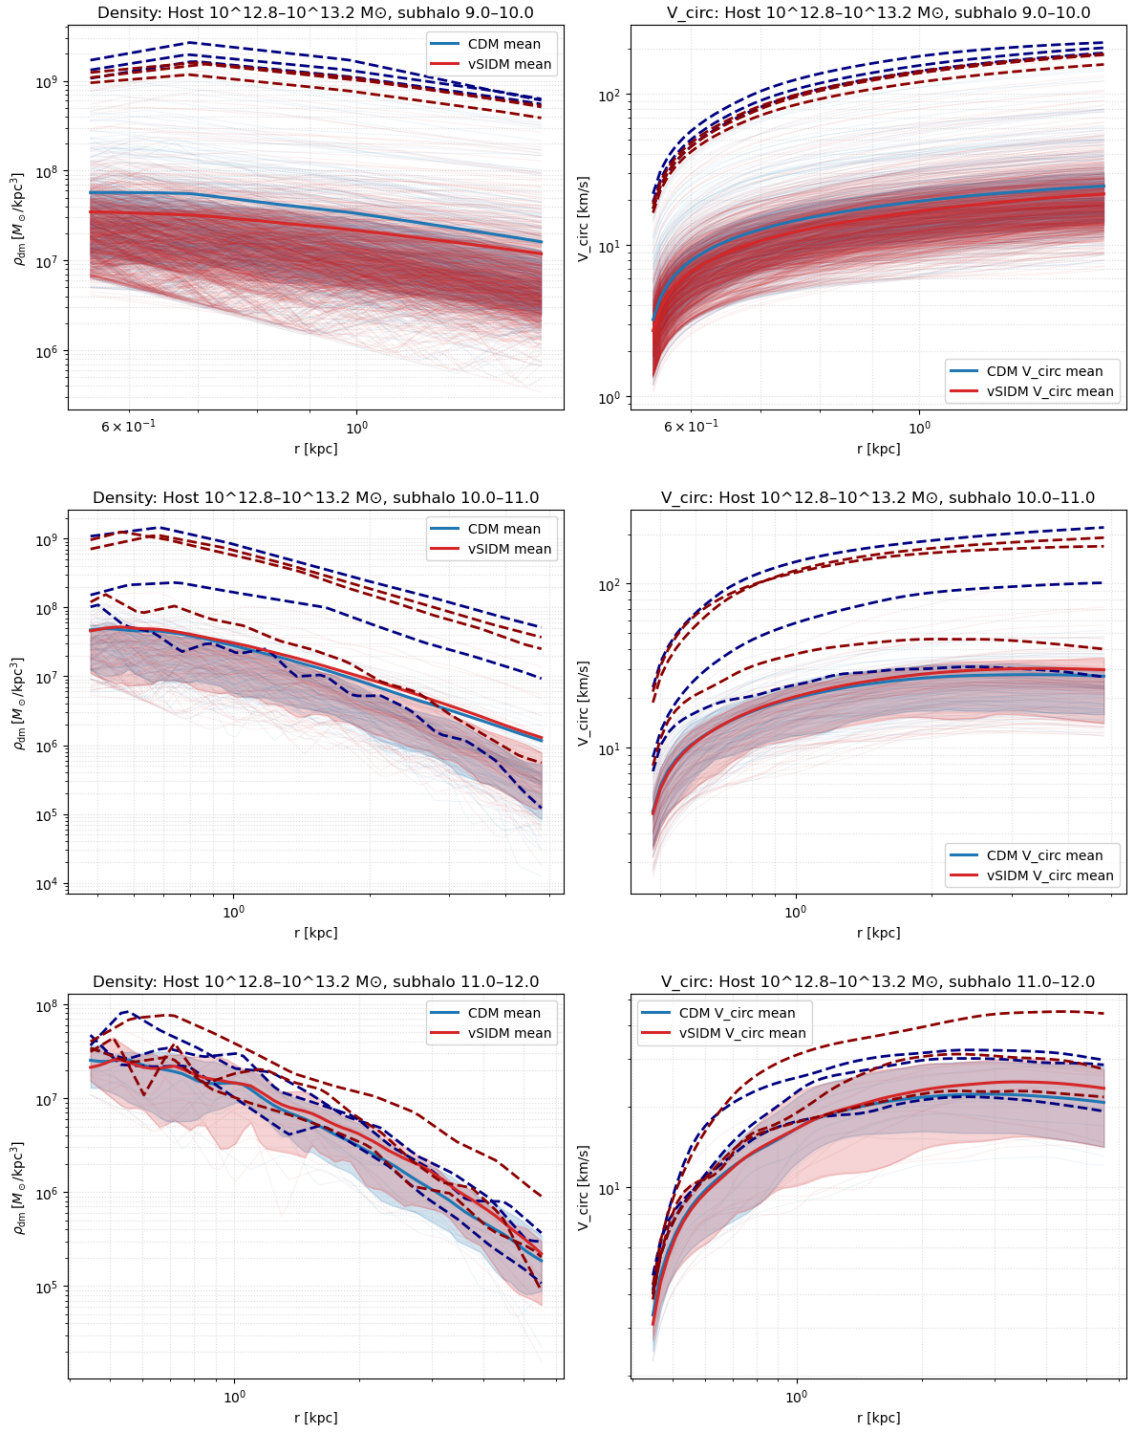


Figure 3.7: Same as Fig. 3.5, but for host-mass bin $10^{12.8}-10^{13.2} M_{\odot}$ (DMO runs).

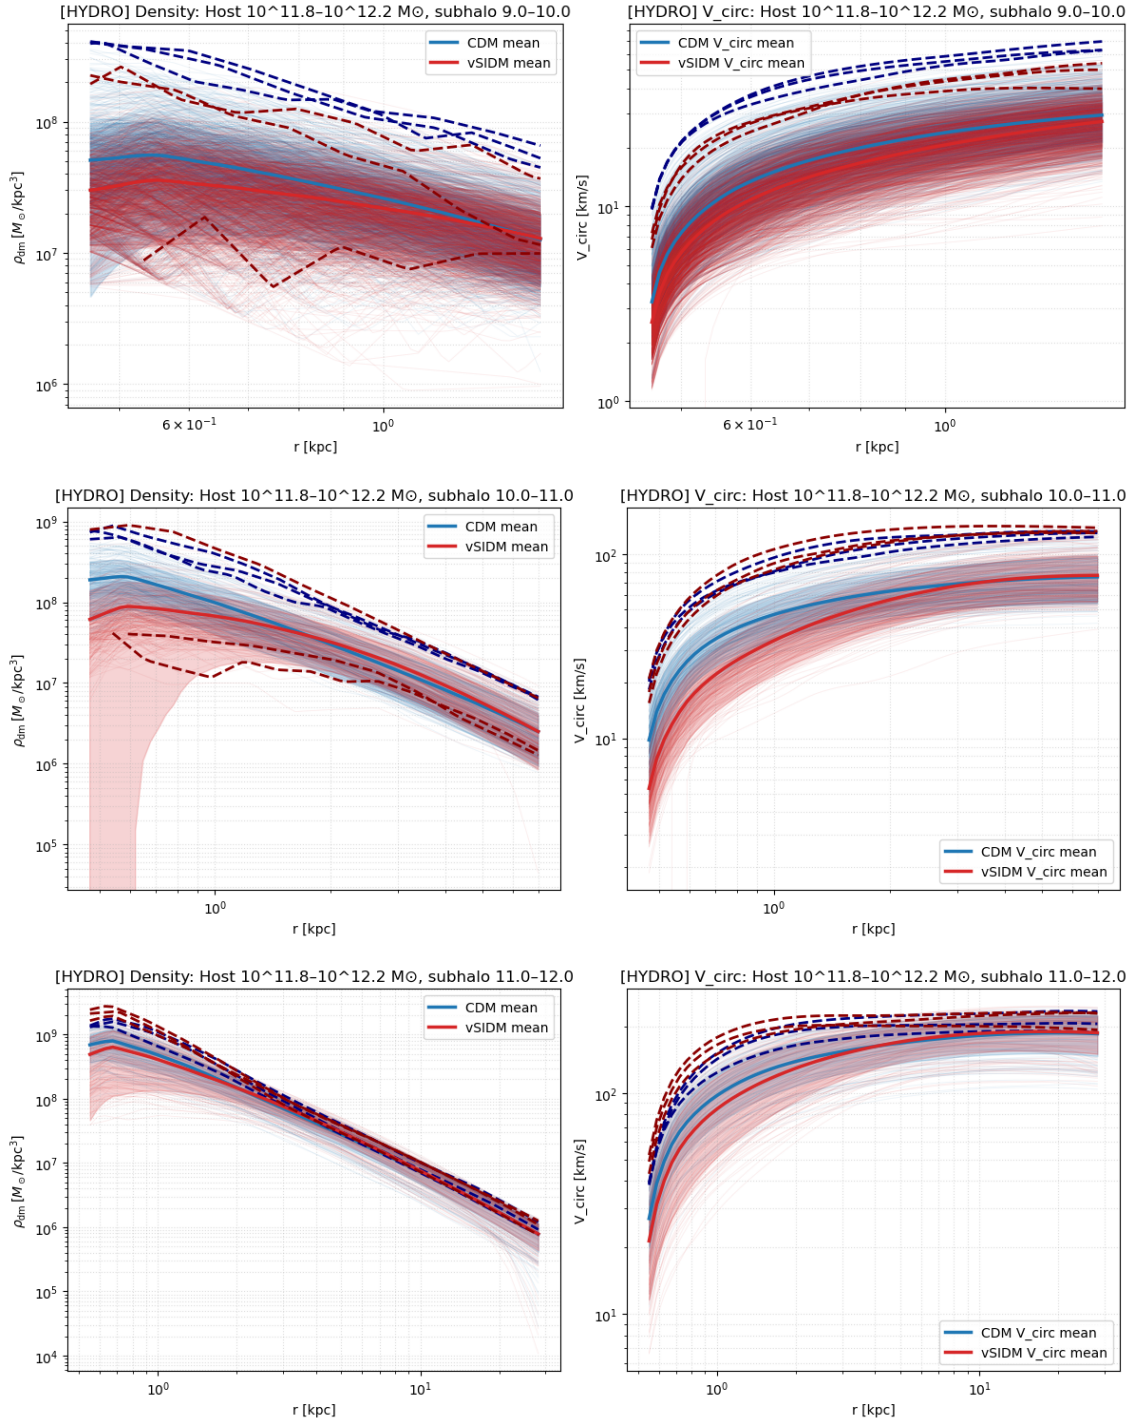


Figure 3.8: Stacked dark-matter density profiles $\rho_{\text{dm}}(r)$ (left columns) and circular-velocity profiles $V_{\text{circ}}(r)$ (right columns) for CDM (blue) and vSIDM (red) in the HYDRO runs, for three subhalo-mass bins and host-mass bin $10^{11.8}-10^{12.2} M_{\odot}$. Line styles follow Fig. 3.3. Compared with DMO, the HYDRO distributions are broader and the CDM–vSIDM separation is correspondingly harder to isolate, especially in the lost-mass panels.

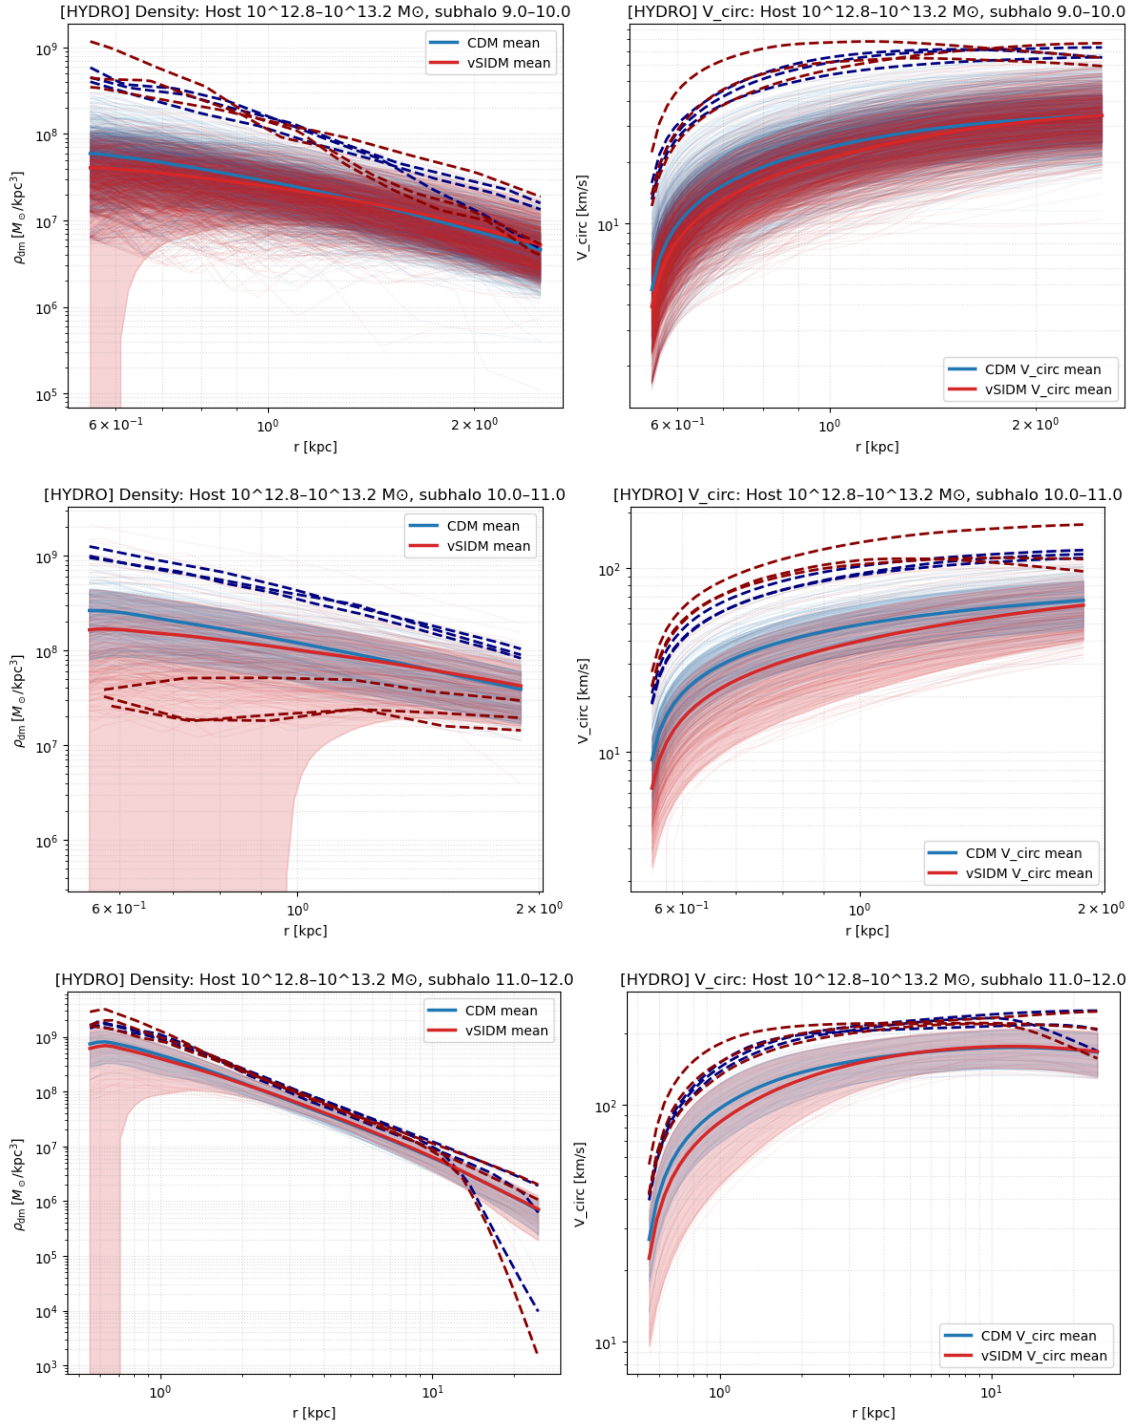


Figure 3.9: Same as Fig. 3.8, but for host-mass bin $10^{12.8}-10^{13.2} M_{\odot}$ (HYDRO runs).

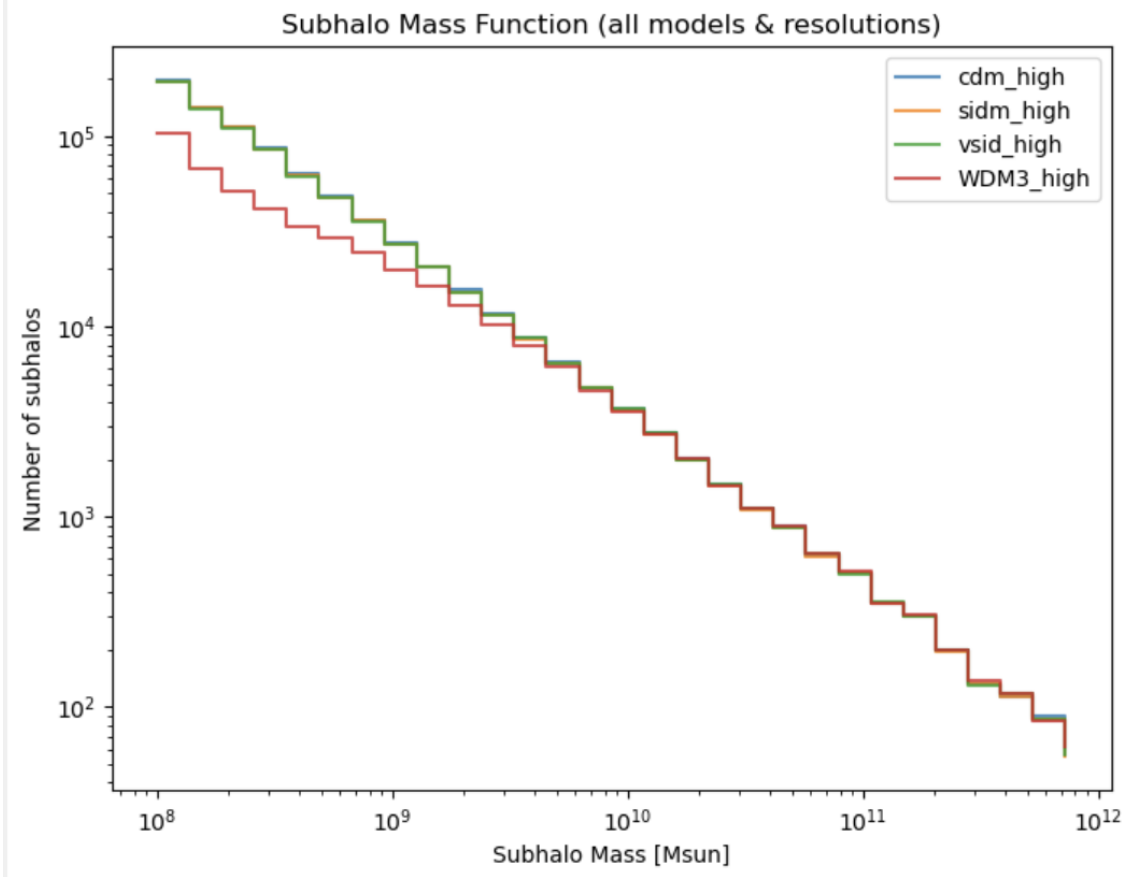


Figure 3.10: Subhalo mass function measured from the high-resolution DMO runs. The total number of subhaloes per logarithmic mass bin is shown for CDM (blue), SIDM1 (orange), vSIDM (green), and WDM3 (red), spanning approximately 10^8 – $10^{12} M_{\odot}$. CDM and SIDM1 are nearly indistinguishable across the plotted range. WDM3 is suppressed relative to CDM at low masses, while vSIDM lies modestly above CDM in the low-mass regime.

In practice, the SHMF declines steeply with increasing subhalo mass, so low-mass subhaloes are intrinsically much more numerous than high-mass ones. Demographic comparisons are therefore most informative at the low-mass end, while the highest-mass bins are always more sensitive to small-number statistics.

Figure 3.10 shows the SHMF measured from the high-resolution DMO runs for all four dark-matter models considered here: CDM, SIDM1, vSIDM, and WDM3. The clearest feature in the figure is the suppression of WDM3 relative to CDM at low subhalo masses. Below roughly a few $\times 10^9 M_{\odot}$, the WDM3 histogram falls progressively below the CDM curve, with the deficit becoming substantial toward $10^8 M_{\odot}$. This is the expected signature of warm dark matter: the free-streaming cutoff in the primordial power spectrum suppresses the formation of low-mass haloes and therefore reduces the abundance of low-mass subhaloes at $z = 0$.

By contrast, the SIDM models show much weaker differences from CDM in this statistic. SIDM1 is visually indistinguishable from CDM across almost the full resolved mass range, indicating that, for the constant-cross-section model considered here, the SHMF is not measurably altered relative to CDM. The vSIDM curve,

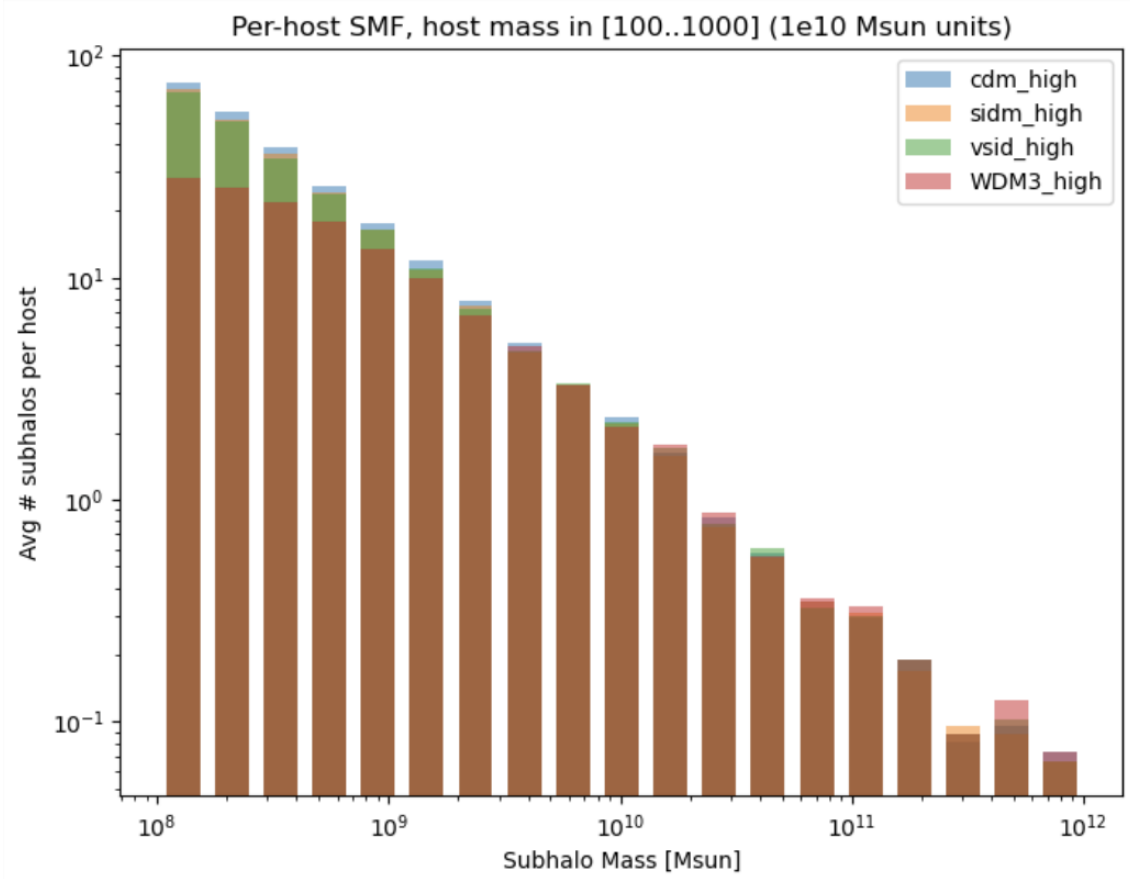


Figure 3.11: Average number of subhaloes per host as a function of subhalo mass for host haloes in the 10^{12} – $10^{13} M_{\odot}$ range, measured from the high-resolution DMO runs. Colours are the same as in Fig. 3.10. The per-host SHMF supports the same broad ordering as the total SHMF: SIDM1 \approx CDM, vSIDM slightly above CDM at low masses, and WDM3 suppressed relative to CDM.

however, lies modestly above CDM at the low-mass end, mainly below $\sim 10^{10} M_{\odot}$, while converging with CDM at higher masses. The vSIDM excess is modest and confined mainly to the low-mass end, indicating that abundance differences are subtler than the structural offsets discussed above.

A per-host view of the same statistic is shown in Fig. 3.11, where we plot the average number of subhaloes per host for host haloes in the 10^{12} – $10^{13} M_{\odot}$ range. The same broad ordering is suggested in this representation: SIDM1 remains close to CDM, vSIDM lies slightly above CDM at low masses, and WDM3 is suppressed relative to both. This per-host view indicates that the low-mass vSIDM excess seen in Fig. 3.10 is not obviously driven only by a very small number of exceptionally subhalo-rich hosts, although it remains a modest effect. Within the present statistic, the sign of the low-mass vSIDM excess is therefore at least qualitatively stable between the total and per-host representations, even if its physical origin remains ambiguous.

The physical origin of the low-mass vSIDM excess cannot be determined from the SHMF alone. Because the selection variable used here is present-day bound subhalo mass (`SubhaloMass`), differences in the SHMF may reflect not only complete

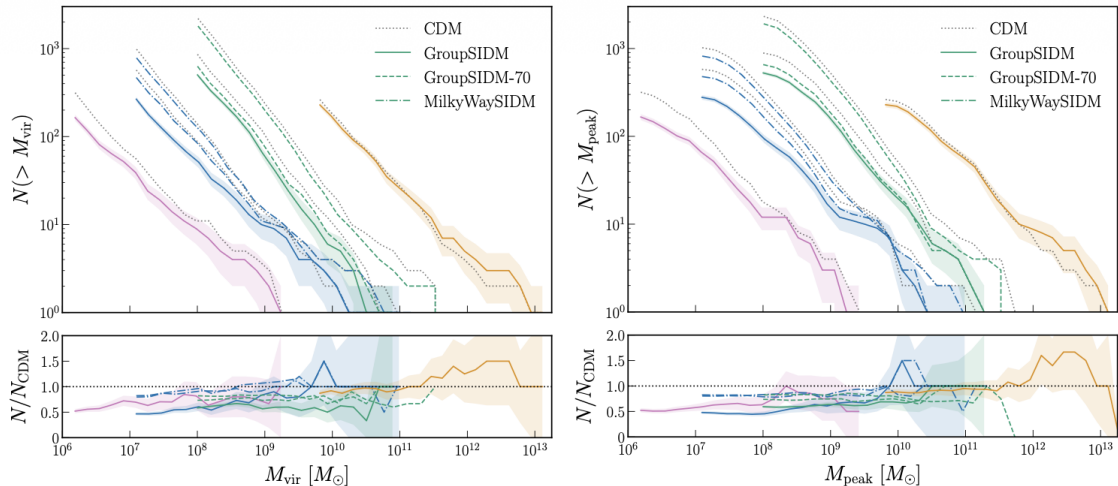


Figure 3.12: Cumulative subhalo mass functions for GroupSIDM, GroupSIDM-70, and MilkyWaySIDM compared to CDM, reproduced from Nadler et al. (2025). Left and right panels show the SHMF as a function of present-day virial mass M_{vir} and peak mass M_{peak} , respectively; lower panels show the ratio N/N_{CDM} . For the more strongly interacting GroupSIDM model, the SHMF is substantially suppressed relative to CDM in several host environments, illustrating that SIDM-induced changes in subhalo abundance are highly model-dependent.

disruption, but also differences in stripping histories and in how structurally distinct subhaloes are redistributed across present-day mass bins. The SHMF therefore does not by itself distinguish between enhanced survival, modified stripping, or remapping between adjacent mass bins. The safest conclusion is simply that, in the present DMO analysis, vSIDM produces a modest excess of low-mass surviving subhaloes relative to CDM when subhaloes are selected by present-day bound mass.

This result differs in sign from the strong suppression found in more aggressively interacting SIDM models. A useful comparison is provided by Nadler et al. (2025), reproduced in Fig. 3.12, where the GroupSIDM model yields cumulative SHMFs suppressed by roughly 50% relative to CDM for LMC-, MW-, and Group-mass hosts. The contrast with our result is not necessarily a contradiction. The contrast instead emphasises that SHMF responses are highly model-specific: the present vSIDM implementation need not reproduce the strong suppression seen in more aggressively interacting SIDM scenarios.

A matched HYDRO SHMF comparison is not included in the present analysis, because the corresponding high-resolution HYDRO runs were not available for all models. Previous work suggests that baryons can reduce subhalo abundances relative to DMO, particularly at low masses and in the inner halo Despali & Vegetti (2017), but whether the low-mass ordering seen here would survive in the high-resolution HYDRO SHMF remains an open question.

In the present dataset, the SHMF shows only modest SIDM-related variation compared with the stronger low-mass suppression seen in WDM3. For CDM versus vSIDM specifically, abundance differences are weaker and less diagnostic than the profile-based structural trends.

Chapter 4

Environmental Dependence of Subhalo Structure

Overview Chapter 3 established that detectable CDM–vSIDM differences are present in the subhalo population, especially in internal-structure diagnostics. The remaining question is how strongly those differences are entangled with orbital environment inside the host halo. Satellites near the host centre have typically experienced very different tidal histories from those near the virial radius, so environmental processing can broaden, mimic, or dilute model-dependent structural trends.

This chapter therefore focuses on orbital environment as an additional axis of interpretation. We first establish the environmental baseline by comparing the radial distribution of subhaloes with the smooth host dark-matter profile, thereby illustrating the familiar antibias of substructure relative to the host halo. We then examine cumulative radial distributions $N(< r)$ in both DMO and HYDRO runs to ask whether the surviving satellite population shows any detectable model-dependent spatial segregation once subhaloes are stacked by host mass and subhalo mass. Finally, we analyse distance-binned density profiles in the HYDRO runs to test more directly how strongly the internal structure of surviving satellites depends on present-day host-centric distance, and whether that environmental variation is comparable to or larger than the CDM–vSIDM offset at fixed distance.

4.1 Subhalo Antibias and Orbital Environment

Before comparing cumulative radial distributions across dark-matter models, it is useful to place the subhalo population in the broader environmental context of the host halo. In CDM, subhaloes are known to be spatially *antibiased* relative to the smooth dark-matter distribution: the number density of surviving substructure increases toward the host centre, but much more slowly than the host dark-matter density itself. As a result, substructure contributes a progressively smaller fraction of the total mass budget in the inner halo. High-resolution studies such as Aquarius (Springel et al. (2008)) interpret this central antibias as a consequence of tidal processing: subhaloes found at small radii are typically earlier-accreted systems that have undergone stronger stripping and, in some cases, disruption.

Figure 4.1 compares, for two host-mass bins, the differential radial distribution of subhaloes and the mean host dark-matter density profile as a function of physical

Subhalo radial distribution vs host DM density profile
(scaled radius r/R_{200c} ; host-mass bins; CDM / SIDM1 / vSIDM)




Figure 4.1: Normalised differential subhalo radial distribution (solid lines) compared with the normalised host dark-matter density profile (dotted lines) as a function of physical radius r for CDM (blue), SIDM1 (green), and vSIDM (orange), in two host-mass bins. In both panels, the smooth host dark-matter profile is substantially more centrally concentrated than the subhalo distribution, illustrating the antibias of substructure relative to the host halo at the level of profile shape. Differences between CDM, SIDM1, and vSIDM are small compared with the overall host-versus-subhalo separation. Because the central subhalo is included in the radial counts, the apparent antibias is conservative relative to a satellite-only measurement.

radius r in kpc. The subhalo curve is effectively $dN/d\ln r$, while the host curve is the mean $\rho_{\text{dm}}(r)$ profile. In each model, both curves are normalised by their own maxima, so the figure isolates differences in *shape* rather than in absolute amplitude. In addition, this version of the diagnostic includes the central subhalo in the radial counts. Because the two curves represent different physical quantities and are independently normalised, the comparison is intended only to contrast relative radial concentration.

The main role of Fig. 4.1 is therefore to show that host environment already imprints a strong baseline radial structure on the surviving subhalo population. Differences between dark-matter models are small compared with the host-versus-subhalo contrast in this diagnostic. Present host-centric distance is thus a useful, albeit imperfect, proxy for cumulative tidal environment and provides a natural coordinate for the more quantitative comparisons that follow.

The current host-centric distance of a subhalo is therefore a useful, albeit imperfect, proxy for its cumulative tidal environment. This motivates the more quantitative model-comparison analysis below and later provides a natural framework for interpreting how structural diagnostics vary with orbital environment.

4.2 Cumulative Radial Distributions

We next test whether the structural differences discussed earlier are accompanied by measurable changes in the radial distribution of surviving satellites. For this purpose we measure cumulative counts $N(< r)$ in bins of host mass and subhalo

mass, expressing radius in units of r/R_{200c} . This statistic does not probe inner structure directly; instead, it asks whether any model-dependent structural differences translate into detectable shifts in the spatial distribution of the surviving subhalo population.

Figures 4.2–4.4 show the DMO cumulative radial distributions for three subhalo-mass bins. In the two lower subhalo-mass ranges, the CDM, vSIDM, and SIDM1 curves are often broadly similar in overall shape across the better-populated lower and intermediate host-mass bins, but the ratio panels show that the agreement is not exact. In particular, at the low-mass end and toward small r/R_{200c} , some of the better-populated bins show a noticeable depletion of surviving vSIDM subhaloes relative to CDM, with the inner ratio reaching values of order ~ 0.5 in the strongest cases. If taken at face value, this corresponds to a substantial reduction in the cumulative number of vSIDM satellites in the inner host halo. SIDM1, by contrast, generally remains closer to CDM within the precision allowed by the statistic.

By contrast, the highest subhalo-mass bin, shown in Fig. 4.4, is dominated by step-like cumulative curves and noisy ratio panels, indicating that the number of objects per panel is too small to support a meaningful systematic comparison. The same caution applies to the most weakly populated host bins in Figs. 4.2 and 4.3, where visually large ratio excursions can be amplified by small- N behaviour rather than reflecting a robust model-dependent trend.

The DMO result is therefore best described as suggestive evidence for enhanced inner depletion of low-mass vSIDM subhaloes relative to CDM, rather than as a uniformly strong effect across all bins. This does not imply that the models are structurally identical. Rather, it indicates that the profile-level differences seen in Chapter 3 may be accompanied by a deficit of surviving vSIDM satellites in the inner halo, although the cumulative radial statistic remains less clean and less direct to interpret than the internal-structure diagnostics.

The HYDRO cumulative distributions, shown in Figs. 4.5–4.7, do not preserve the DMO behaviour in a comparably clear way. In the lowest subhalo-mass bin, some of the better-populated lower host-mass panels show a mild tendency for the vSIDM curve to lie slightly above CDM at small r/R_{200c} , corresponding to ratio values modestly above unity in the inner halo. However, this effect is small, not present in all panels, and not robust enough on its own to support a strong physical interpretation. In the intermediate subhalo-mass bin, CDM and vSIDM again track each other broadly, with ratio panels fluctuating around unity without a coherent sign across host bins. The highest subhalo-mass HYDRO bin is, as in DMO, too sparse to interpret systematically.

Taken together, the HYDRO figures do not reveal a strong or clean model-dependent radial segregation either. Instead, they indicate that once baryonic physics is included, any residual CDM–vSIDM difference in cumulative $N(< r)$ remains small compared with the host-to-host and panel-to-panel variation visible in the sample. The cumulative spatial statistic is therefore consistent with the broader picture established above: baryons and environmental processing dominate the selection of the surviving inner-halo population, while any microphysics-driven signal in radial counts remains subdominant.

At the smallest radii, interpretation is particularly delicate. This is where baryonic stripping is strongest, where configuration-space subhalo finders such as SUB-

DMO: cumulative subhalo distribution $10^9\text{--}10^{10} M_\odot$

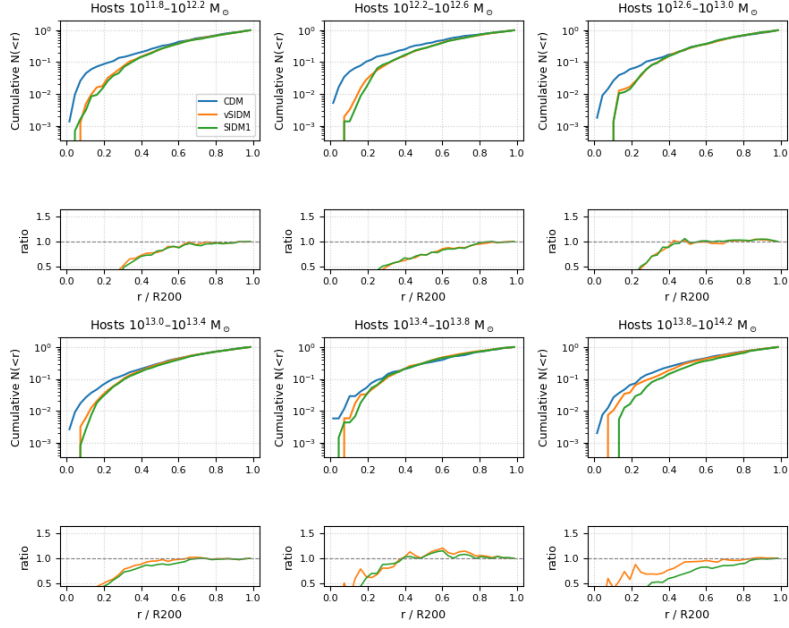


Figure 4.2: DMO cumulative radial distributions $N(< r)$ as a function of r/R_{200c} for subhalo mass bin $10^9\text{--}10^{10} M_\odot$, split by host-mass bin. Curves compare CDM, vSIDM, and SIDM1; ratio panels show model-to-CDM. The central subhalo is excluded. The highest host-mass panels are visibly step-like and should be interpreted with caution.

DMO: cumulative subhalo distribution $10^{10}\text{--}10^{11} M_\odot$

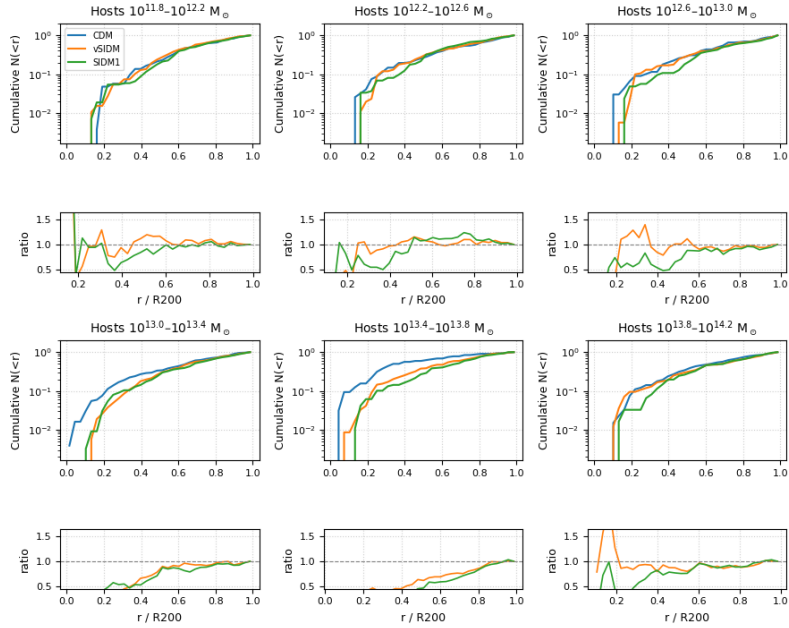


Figure 4.3: Same as Fig. 4.2, but for subhalo mass bin $10^{10}\text{--}10^{11} M_\odot$ (DMO runs).

DMO: cumulative subhalo distribution $10^{11}\text{--}10^{12} M_{\odot}$

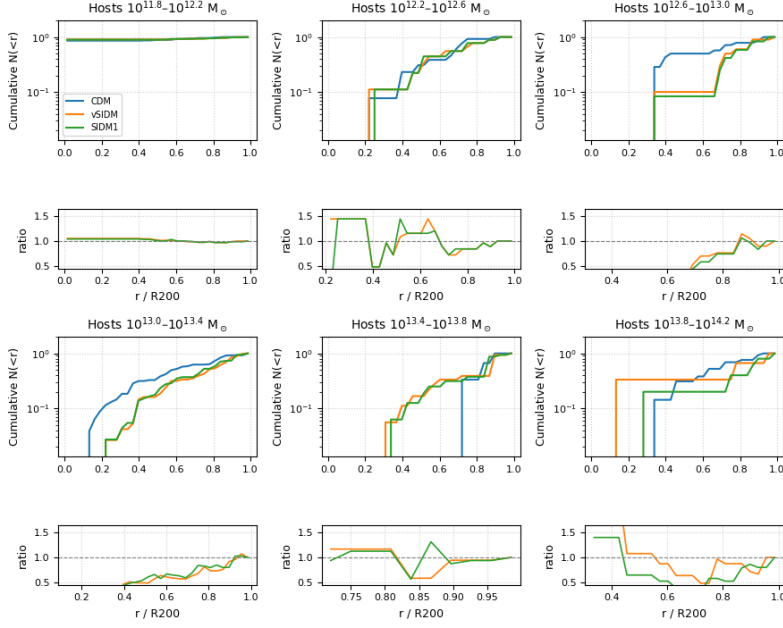


Figure 4.4: Same as Fig. 4.2, but for subhalo mass bin $10^{11}\text{--}10^{12} M_{\odot}$ (DMO runs).

HYDRO: cumulative subhalo distribution $10^9\text{--}10^{10} M_{\odot}$

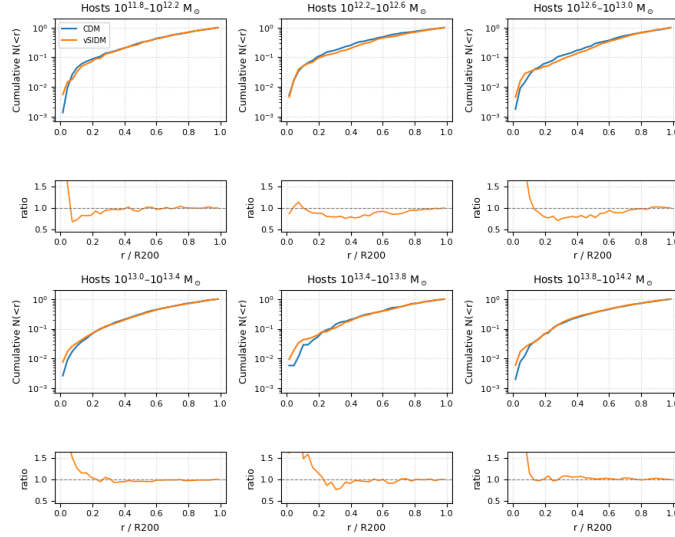


Figure 4.5: HYDRO cumulative radial distributions $N(< r)$ as a function of r/R_{200c} for subhalo mass bin $10^9\text{--}10^{10} M_{\odot}$, split by host-mass bin. Curves compare CDM and vSIDM; ratio panels show model-to-CDM. The central subhalo is excluded. In some of the better-populated lower host-mass bins, vSIDM lies slightly above CDM at small scaled radii, but the effect is weak and should be interpreted cautiously.

HYDRO: cumulative subhalo distribution $10^{10}\text{--}10^{11} M_{\odot}$

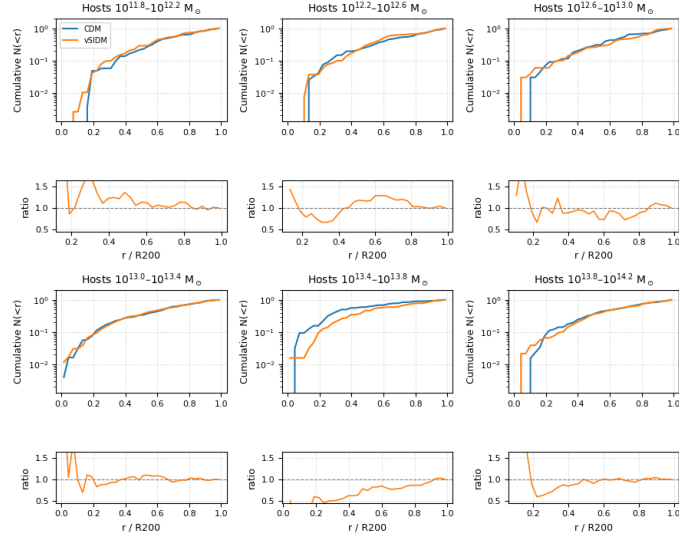


Figure 4.6: Same as Fig. 4.5, but for subhalo mass bin $10^{10}\text{--}10^{11} M_{\odot}$ (HYDRO runs).

HYDRO: cumulative subhalo distribution $10^{11}\text{--}10^{12} M_{\odot}$

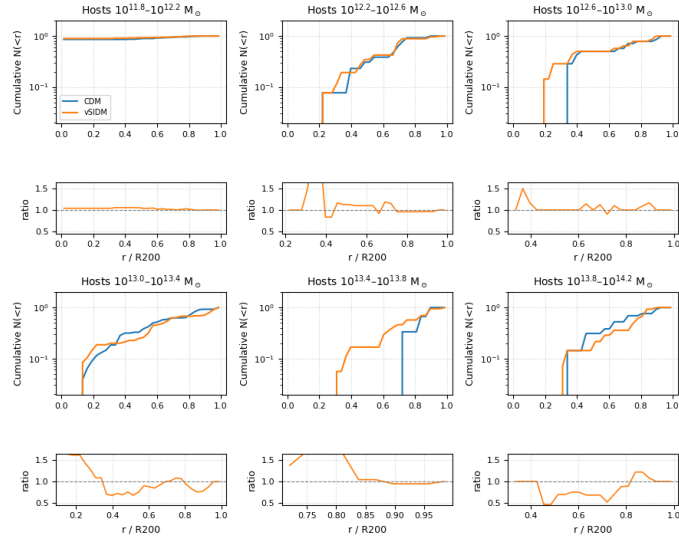


Figure 4.7: Same as Fig. 4.5, but for subhalo mass bin $10^{11}\text{--}10^{12} M_{\odot}$ (HYDRO runs).

FIND are least complete against the dense host background, and where cumulative counts can be affected strongly by only a few objects. For this reason, any inner-halo excess or deficit visible in the ratio panels is treated here as suggestive rather than decisive, especially in the most massive host bins and in the highest subhalo-mass bins.

A useful qualitative DMO benchmark is again provided by Aquarius (Springel et al. (2008)), which showed that subhaloes are centrally antibiased relative to the smooth halo and that inner depletion primarily reflects prolonged tidal processing of early-accreted systems. Their analysis uses a different radial normalisation from ours, so the comparison is qualitative rather than one-to-one, but the overall picture is consistent: the radial distribution of surviving subhaloes is shaped strongly by environmental processing, and any additional CDM–vSIDM signal in the cumulative counts is comparatively subtle.

Relative to the internal-structure diagnostics of Chapter 3, cumulative radial distributions are therefore a less sensitive discriminator of CDM versus vSIDM in the present dataset. In DMO, some of the better-populated low-mass bins show a noticeable inner deficit of surviving vSIDM subhaloes relative to CDM, but this trend is not uniform across all host-mass bins and becomes difficult to assess in the sparse high-mass panels. In HYDRO, the corresponding cumulative radial distributions do not preserve the same behaviour cleanly: some low-mass panels instead show a mild vSIDM excess at small r/R_{200c} , while others remain close to unity. The cumulative radial statistic therefore provides useful environmental context, but it remains a markedly less robust probe of the CDM–vSIDM difference than the direct profile-based structural diagnostics.

4.3 Distance-Dependent Density Profiles in the HYDRO Runs

To probe environmental dependence more directly, we bin HYDRO subhaloes by present-day host-centric physical distance d and stack their dark-matter density profiles within each host-mass bin. Satellites from all hosts in the same host-mass bin are pooled, the central subhalo is excluded, and the distance bins are fixed logarithmic bins in physical kpc.

Profiles are interpolated onto a common logarithmic radial grid, and the stacked profile is computed from the mean in $\log_{10} \rho_{\text{dm}}(r)$ before converting back to linear density for plotting. These stacks therefore behave more like typical-profile estimators than the arithmetic-mean profiles used in Chapter 3, so the comparison between the two chapters should be read qualitatively rather than as a strict one-to-one amplitude match.

A crucial point is that the distance bins are defined in physical kpc rather than in r/R_{200c} , so the same bin does not correspond to the same relative orbital phase across host-mass panels. In addition, no narrow subhalo-mass cut is imposed within each host-mass panel, so the observed distance trend may reflect both genuine environmental processing and radial changes in the mass mix of the surviving satellite population.

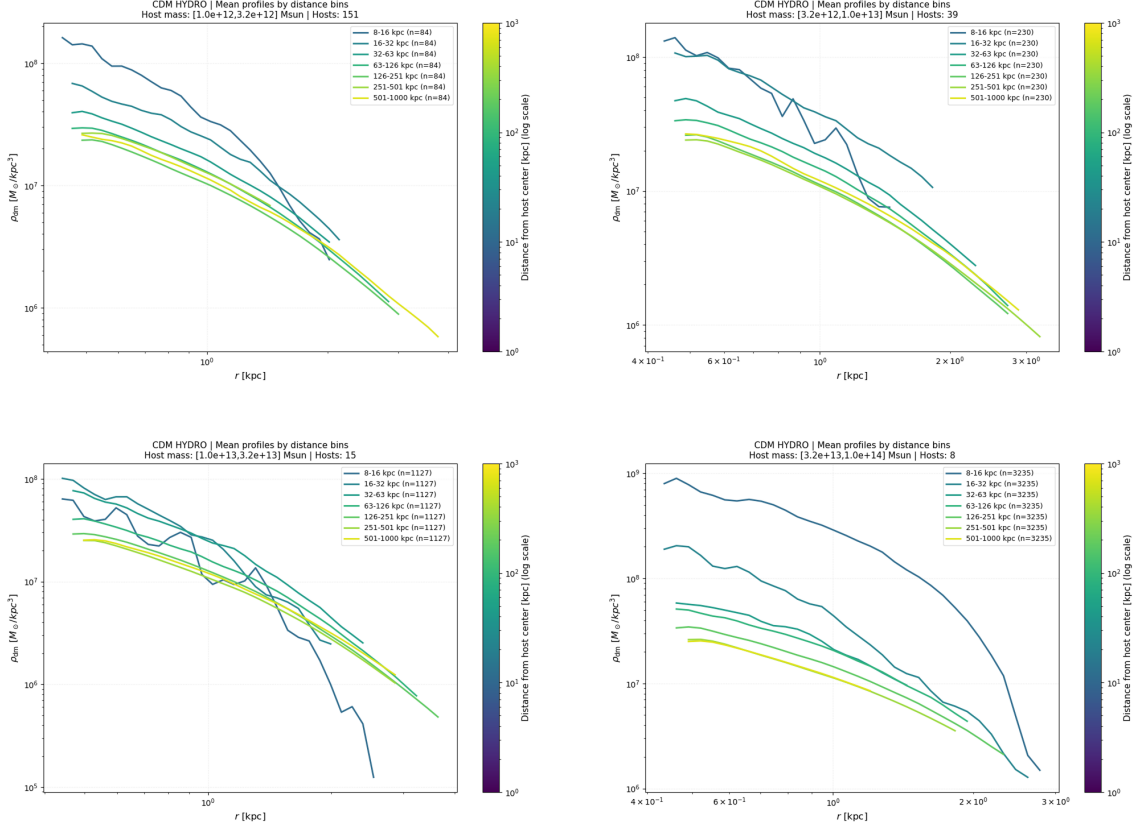


Figure 4.8: Mean dark-matter density profiles $\langle \rho_{\text{dm}}(r) \rangle$ for CDM satellite subhaloes in the HYDRO runs, split into fixed bins of current host-centric physical distance for four host-mass bins. Curves are colour-coded from dark blue (innermost occupied distance bin) to yellow–green (outermost distance bin). In every panel, inner-halo satellites are systematically denser than outer-halo satellites at fixed plotted radius, demonstrating a strong environmental ordering of the mean profiles.

4.3.1 CDM HYDRO

Figure 4.8 shows the distance-binned mean dark-matter density profiles for CDM satellite subhaloes in the HYDRO runs, split by host-mass bin and coloured by host-centric distance. In all four panels, the mean profiles are strongly ordered by distance: satellites in the innermost occupied distance bins have the highest mean densities over the full resolved radial range, while the outermost bins have the lowest mean densities. This ordering is visible across all host-mass bins and is especially smooth in the better-populated high-mass panels.

The most robust conclusion from Fig. 4.8 is therefore that current host-centric distance is a strong predictor of mean subhalo density in the HYDRO satellite population. The separation between the inner and outer distance bins spans a large dynamic range in density, and the ordering is preserved across the resolved radial range rather than being confined to a single point. This trend is qualitatively consistent with stronger tidal processing and stronger survival selection in the inner halo: satellites found at small host-centric distances are, on average, a more heavily processed population than those at large distances. The figure alone, however, does not identify a unique mechanism. Both environmental processing of the profile




Figure 4.9: Average dark-matter density profiles for CDM subhaloes with masses 3.5×10^9 – $6.7 \times 10^9 M_\odot$, binned by current host-centric distance R/R_{200} , reproduced from Heinze et al. (2024) Heinze et al. (2024). Curves are colour-coded from the innermost bin (blue) to the outermost (red); the shaded region marks the softening-limited inner boundary. The same qualitative ordering seen in Fig. 4.8 is present here: inner-distance bins are systematically denser than outer-distance bins at fixed radius.

shape and survival bias—where only the densest satellites survive to small radii—can contribute to the observed ordering.

The apparent separation between the innermost and outermost distance bins also tends to increase toward the more massive host bins. This suggests that environmental processing spans a wider range of outcomes in more massive hosts, although the exact strength of this increase should not be over-interpreted because the distance bins are defined in physical kpc rather than in r/R_{200c} . For this reason, the host-mass trend is interpreted qualitatively rather than as a directly comparable amplitude across panels. Over the resolved radial range, the dominant visible effect is a change in profile normalisation with distance bin; any systematic change in inner slope is less clearly isolated in these stacked means.

A useful literature comparison is provided by Heinze et al. (2024), reproduced in Fig. 4.9. Their CDM subhaloes, binned by current host-centric distance expressed in R/R_{200} , show the same qualitative ordering: inner bins are denser at fixed radius than outer bins over the full resolved radial range. Although the distance definition and simulation setup differ from ours, the qualitative agreement supports the interpretation that host-centric distance is a robust environmental driver of subhalo density.




Figure 4.10: Same as Fig. 4.8, but for vSIDM in the HYDRO runs. The same broad ordering of mean density profiles with current host-centric distance is present as in CDM: inner-distance bins lie above outer-distance bins across the resolved radial range in all host-mass panels.

4.3.2 vSIDM HYDRO

Figure 4.10 shows the corresponding distance-binned mean profiles for vSIDM in the HYDRO runs. The same broad environmental ordering is present as in CDM: satellites in the innermost occupied distance bins are systematically denser than those in the outermost bins over the resolved radial range, and the ordering is clearest in the better-populated high-mass panels.

Figure 4.10 shows the same basic environmental ordering in the vSIDM HYDRO runs: inner-distance bins lie above outer-distance bins across the resolved radial range in every host-mass panel. The key point is that this ordering remains strong in vSIDM rather than being washed out by self-interactions.

Some of the innermost vSIDM bins appear more irregular than their CDM counterparts, particularly in lower- N panels and at the smallest resolved radii, but this is not quantified directly and should not be over-interpreted. Because the CDM and vSIDM distance-binned profiles are shown separately rather than as matched residuals, the relative strength of the distance gradient in the two models cannot be measured robustly from these figures alone.

One potentially interesting feature is that, in several host-mass panels, the vSIDM sample populates an additional innermost distance bin that is not present in the corresponding CDM figure. This indicates a difference in the sampled inner-halo

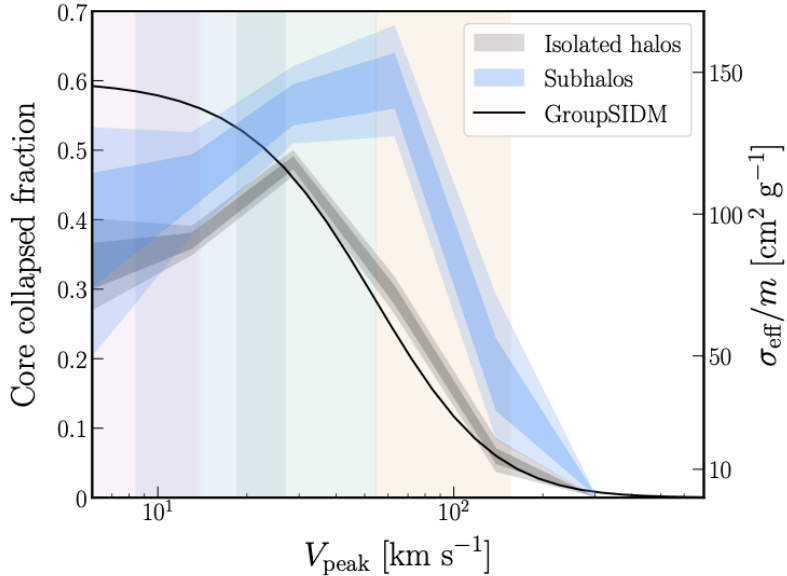


Figure 4.11: Core-collapsed fraction as a function of peak circular velocity V_{peak} for isolated haloes and subhaloes in the GroupSIDM simulations, reproduced from Nadler et al. (2025). Subhaloes exhibit a higher collapsed fraction than isolated haloes at low V_{peak} , illustrating that tidal environment can accelerate gravothermal evolution in strongly interacting SIDM models. This figure is used here as qualitative context rather than as a direct one-to-one comparison with the AIDA–TNG vSIDM model.

population, but by itself it does not identify the cause. It could reflect enhanced survival, different radial occupancy, or another selection effect associated with the fixed-bin procedure. We therefore refrain from assigning a unique physical interpretation to this feature on the basis of the present plots alone. It may be related to the mild inner-halo vSIDM excess seen in some of the HYDRO cumulative $N(< r)$ panels, but the present analysis does not demonstrate that the same objects are responsible for both features.

A useful literature context is provided by Nadler et al. (2025), reproduced in Fig. 4.11, who show that subhaloes in their more aggressively interacting GroupSIDM model have a higher core-collapsed fraction than isolated haloes at low V_{peak} . This supports the general idea that tidal environment can accelerate gravothermal evolution in some SIDM models. However, the Nadler model is substantially more strongly interacting than the vSIDM model considered here, so the comparison is suggestive rather than one-to-one. In the present analysis, the most that can safely be said is that the innermost vSIDM distance bins may sample a more structurally diverse population than the corresponding CDM bins.

4.3.3 Environmental variation versus model variation

The comparison between Figs. 4.8 and 4.10 makes it possible to compare two sources of variation directly: environmental spread within a model, and residual model-to-

model offset at fixed distance. In the robust HYDRO distance bins, the profile variation associated with present-day host-centric distance is generally of the same order as, and usually larger than, the CDM–vSIDM offset at fixed distance.

This does not imply that the model difference vanishes. Rather, it shows that environmental mixing can easily mimic or dilute part of the apparent structural separation unless host-centric distance is controlled explicitly. The distance-binned analysis therefore demonstrates directly that orbital environment is a major confounding variable for model discrimination in subhalo structure.

From an observational point of view, this means that environment-driven profile differences in the inner halo can be comparable to, or larger than, the direct CDM–vSIDM structural offset discussed in Chapter 3. This does not render microphysical signatures undetectable, but it does mean that attempts to infer dark-matter physics from subhalo density structure must account for the strong dependence of apparent inner density on current orbital environment. Otherwise, tidal evolution and survival selection within the host halo can be misidentified as evidence for, or against, a particular dark-matter model.

Chapter summary This chapter shows that orbital environment is a major organising variable of subhalo structure. The antibias of subhaloes relative to the smooth host halo is clear but only weakly model-dependent in the present qualitative diagnostic, and cumulative radial distributions show little robust evidence for model-dependent spatial segregation in either DMO or HYDRO. By contrast, distance-binned HYDRO profiles display a strong environmental ordering in both CDM and vSIDM: satellites closer to the host centre are systematically denser than those farther out. In the well-populated bins, this environmental variation is typically comparable to or larger than the residual CDM–vSIDM offset, showing that orbital environment must be controlled explicitly when interpreting subhalo structure.

Chapter 5

Discussion and Conclusions

This thesis examined how the internal structure and demographic properties of dark-matter subhaloes differ across dark-matter models in the AIDA–TNG simulations, with primary emphasis on the comparison between collisionless cold dark matter (CDM) and velocity-dependent self-interacting dark matter (vSIDM). The central goal was not only to test whether the two models produce distinguishable subhalo populations, but also to determine how strongly those differences are altered or obscured once baryonic physics and environmental processing inside the host halo are taken into account.

The analysis was organised around two complementary questions. First, do CDM and vSIDM produce measurable differences in the structure and abundance of surviving subhaloes? Second, if such differences exist, are they robust once one compares dark-matter-only (DMO) and hydrodynamical (HYDRO) runs and controls, as far as possible, for the effects of host environment? Taken together, the results show that CDM and vSIDM do produce measurable differences in the surviving subhalo population, but that the strength and clarity of those differences depend strongly on the diagnostic considered. In the present dataset, the clearest separation appears in profile-based structural measures, especially in DMO, while abundance statistics and cumulative radial counts provide complementary but less directly interpretable evidence. Once baryonic physics and host environment are included, the structural diversity of the surviving population broadens substantially, making model discrimination in HYDRO more difficult but not eliminating the underlying trends.

5.1 Summary of the main results

The catalogue-based (R_{\max}, V_{\max}) plane shows only a modest CDM–vSIDM separation. In both DMO and HYDRO, the matched-bin offset in R_{\max} at fixed V_{\max} remains small and changes sign across the sampled velocity range, rather than revealing a single monotonic displacement between the two models. This indicates that the (R_{\max}, V_{\max}) plane captures some model dependence, but not in a form that provides a clean standalone discriminator of vSIDM in the present simulations.

Because R_{\max} and V_{\max} compress the profile into two values, they are sensitive not only to inner density structure but also to tidal stripping, truncation, and the outer shape of the circular-velocity curve. A weak sign reversal with V_{\max} is therefore more

naturally interpreted as evidence that these two quantities are influenced by several structural effects, rather than being able to provide a direct monotonic measure of self-interaction-driven expansion or contraction.

A clearer distinction emerges in the stacked dark-matter density and circular-velocity profiles. In the DMO runs, the strongest and most reproducible separation appears in the lower subhalo-mass bins, where CDM subhaloes typically have higher inner dark-matter densities than vSIDM subhaloes over the resolved radial range. The corresponding $V_{\text{circ}}(r)$ profiles show the same ordering with smaller amplitude, consistent with the fact that circular velocity is an enclosed-mass diagnostic and therefore responds more weakly to local structural changes than the density profile itself. At higher subhalo masses the separation becomes less uniform, and in several bins the mean offset is comparable to the intrinsic object-to-object scatter. The DMO result is therefore best described as a modest and non-uniform model dependence in cumulative radial distribution.

The comparison between the L50 and higher-resolution L25 DMO runs strengthens this interpretation and helps define its robust scope. By extending the resolved radial range inward, the L25 run makes the low-mass CDM–vSIDM density difference easier to identify and confirms that its sign is stable across resolution levels. This is physically consistent with the expected vSIDM behaviour, namely reduced central densities in low-mass subhaloes relative to CDM. The remaining caution concerns only the detailed innermost structure: the present analysis does not establish strict numerical convergence of the central slope or a sufficiently robust collapse signature to support stronger claims about the exact inner profile shape.

The HYDRO profile comparison leads to a more mixed outcome. The DMO trend is not cleanly reversed or erased, but it becomes substantially harder to isolate once baryons are included. The HYDRO distributions are broader, panel-to-panel variation is larger, and in several bins the CDM and vSIDM means overlap within the intrinsic distribution width. This indicates that baryonic physics does not simply shift a fixed subhalo population relative to the DMO case. Instead, HYDRO changes the diversity, stripping history, and survival of the satellites that remain in the sample. In this sense, the structural CDM–vSIDM signal appears to persist in HYDRO, but in a form that is no longer cleanly separable from baryon-driven and selection-driven broadening.

The subhalo mass function (SHMF) shows a different pattern. The strongest demographic effect is the expected suppression of low-mass subhaloes in WDM3 relative to CDM, while SIDM1 remains close to CDM over the resolved mass range. The vSIDM SHMF lies modestly above CDM at the low-mass end in the DMO analysis, but this excess is considerably weaker and more ambiguous in interpretation than the profile-based structural differences. Because subhaloes are selected by present-day bound mass rather than by a progenitor quantity such as M_{peak} , the SHMF cannot distinguish cleanly between enhanced survival, altered stripping, and redistribution between present-day mass bins. Physically, the modest low-mass vSIDM excess could therefore arise in several ways: the effective self-interaction rate at the relevant subhalo velocities may be too small to enhance disruption strongly; altered stripping histories may shift subhaloes into lower present-day mass bins rather than destroying them; or survival and selection effects may populate the low-mass surviving sample differently from CDM at fixed `SubhaloMass`. The present

work does not distinguish between these possibilities. The SHMF therefore serves mainly as a secondary demographic diagnostic rather than as the primary evidence for CDM–vSIDM differences.

Although WDM3 is not the primary focus of this thesis, it plays an important comparative role in interpreting the results. In the present dataset, WDM3 produces the clearest demographic signal, namely a low-mass suppression in the SHMF relative to CDM, whereas the vSIDM signal is more visible in internal structure than in abundance. This contrast is itself informative: it suggests that different classes of dark-matter models need not reveal themselves most clearly in the same observable channel. In this sense, the AIDA–TNG results support a useful distinction between WDM-like behaviour, which is most apparent here in the abundance of low-mass subhaloes, and SIDM-like behaviour, which is more apparent in the internal structure of the surviving population. This contrast also suggests a broader observational lesson: different dark-matter scenarios may be more robustly distinguished by combining channels that probe abundance and internal structure simultaneously, rather than by relying on either class of observable alone.

The spatial statistics of Chapter 4 provide only weak evidence for model-dependent differences in the surviving radial distribution of subhaloes. In both DMO and HYDRO, the cumulative radial distributions $N(< r)$ as a function of r/R_{200c} show at most modest and non-uniform deviations between CDM and vSIDM across the better-populated bins, and no strong or reproducibly signed model separation emerges across host masses. In this sense, the cumulative counts are less informative than the profile-based diagnostics. Their limitation is not only interpretive degeneracy, but also the fact that, in the present sample, the model-dependent signal remains weak compared with the bin-to-bin variation.

By contrast, the distance-binned HYDRO profile analysis reveals a strong environmental trend. In both CDM and vSIDM, satellites at smaller present-day host-centric distances are systematically denser than satellites at larger distances across the resolved radial range. This ordering is much stronger than any model-dependent trend visible in the cumulative radial counts. Because the CDM and vSIDM distance-binned profiles are shown in separate figures rather than as a matched residual comparison, the relative size of the environmental spread and the model-to-model offset is not quantified directly in this thesis. Nevertheless, by visual comparison of the two sets of distance-binned profiles, the structural variation associated with host-centric distance appears to be at least of the same order as the residual CDM–vSIDM separation in the better-populated HYDRO bins, and in some panels it appears larger. This conclusion should therefore be understood as qualitative rather than as a measured amplitude ratio.

5.2 Interpretation

Taken together, the results support a coherent hierarchy among the diagnostics used in this thesis. In the present AIDA–TNG dataset, the clearest and most directly interpretable CDM–vSIDM differences appear in the dark-matter density profiles and, more weakly, in the circular-velocity profiles of surviving subhaloes. Catalogue-level structural summaries such as the (R_{\max}, V_{\max}) plane show milder trends, while abundance and cumulative spatial statistics are less informative and, in the present

sample, do not provide a comparably robust or reproducible model separation.

At the same time, the thesis also shows that the detectability of these structural differences depends strongly on what is being held fixed. Throughout most of the analysis, subhaloes are compared at fixed present-day bound mass. This choice is operationally natural, but it means that the measured differences combine at least three effects: genuine differences in inner structure, model-dependent stripping histories, and changes in which subhaloes survive or satisfy the profile-selection criteria. The interpretation is therefore necessarily population-based: the results describe how the *surviving selected subhalo population* differs between CDM and vSIDM, not how matched progenitors would differ under otherwise identical histories.

This distinction is especially important for the DMO–HYDRO comparison. In principle one would like to separate the effect of dark-matter microphysics from the effect of baryonic physics. In practice, however, HYDRO does not act on a fixed DMO subhalo catalogue. It changes the orbits, masses, central densities, stripping histories, and in some cases even the detectability of satellites against the host background. For that reason, the thesis does not claim a clean additive decomposition of “microphysics plus baryons”. The more defensible conclusion is narrower: the CDM–vSIDM structural trend that is relatively visible in DMO becomes substantially harder to isolate in HYDRO because baryons broaden the structural diversity of the surviving sample and enhance the importance of environmental processing.

The environmental analysis sharpens this point further. The strong dependence of HYDRO profile shape on present-day host-centric distance demonstrates that subhalo structure cannot be interpreted in isolation from orbital context. Even when the underlying microphysical signal is present, environmental mixing can mimic, wash out, or re-weight the apparent model separation in stacked samples. In this sense, host-centric distance acts as a confounding variable for model discrimination. This does not make dark-matter microphysics unobservable, but it does imply that any robust attempt to infer microphysics from subhalo structure must control for environment at the same time.

An additional point of interest is the presence of individual vSIDM subhaloes whose inner densities reach, and in some panels exceed, the upper end of the CDM distribution. In the results chapters these were described conservatively as collapse candidates or examples of collapse-like behaviour rather than as confirmed cases of gravothermal core collapse. That caution remains necessary here. The present analysis does not provide the combination of evidence required for a stronger claim: the innermost resolved radii remain subject to convergence limitations, the candidate objects are identified from extreme tails rather than from a robust population-level reversal, and no time evolution is used to test whether these systems underwent the expected transition from core formation to renewed central steepening. Moreover, these objects are not representative of the vSIDM population as a whole, but are confined to the extreme high-density tail and appear only in a limited subset of the plotted DMO bins. This reinforces the interpretation that they are physically interesting outliers rather than, in the present data, evidence for a dominant population-level regime. The current results therefore support the existence of structurally unusual vSIDM outliers, but not a definitive identification of core collapse.

5.3 Limitations of the present analysis

Several limitations should be kept in mind when interpreting the results of this thesis.

The profile-based comparisons adopt an inner cut at $r \geq 3\epsilon_{\text{DM}}$, which is widely used to avoid the regime most directly affected by gravitational softening. This is a conservative and standard choice for comparative profile analysis, but it is not a full demonstration of numerical convergence of the innermost profile shape. The resulting profiles are therefore interpreted as adequate for comparative structural trends across the resolved radial range, while any conclusions that depend strongly on the innermost bins — especially possible collapse-like upturns — are treated with additional caution.

Second, the stacked profile samples represent only a subset of the mass-selected subhalo catalogues. The retained fraction is lower for vSIDM than for CDM, particularly in the profile-based analysis. This asymmetry does not merely reduce sample size. If the rejected vSIDM objects are preferentially the most diffuse or least well-resolved systems — which is physically plausible, since extended low-density profiles are more likely to fail common-grid or inner-range validity requirements than more centrally concentrated profiles — then the retained vSIDM stacks are biased toward the more compact tail of the surviving distribution. In that case, the stacked CDM–vSIDM separation shown in this thesis would be conservative, in the sense that it would underestimate the true structural difference in the full mass-selected population. This direction of the bias is not demonstrated directly here, but it is the most plausible interpretation of the observed selection asymmetry.

Third, most comparisons are performed at fixed present-day `SubhaloMass`. This choice is suitable for describing the properties of the surviving population at $z = 0$, but it mixes together structural evolution, stripping, and survival selection. A progenitor-matched analysis using quantities such as M_{peak} or V_{peak} would be better suited to isolating the effect of dark-matter microphysics on the evolution of individual systems.

Fourth, the cumulative radial statistics become noisy in the highest subhalo-mass bins and in the least-populated host bins. The absence of a strong model-dependent signal in those panels should therefore be interpreted as an absence of evidence within the precision of the present sample, not as a proof of exact equality between the models.

Finally, a matched high-resolution HYDRO SHMF comparison was not available for all dark-matter models. This limits how far the demographic comparison can be extended beyond the DMO case and leaves open the question of whether the low-mass vSIDM excess seen in the DMO SHMF would survive once baryonic suppression is included consistently across all models.

5.4 Future directions

The present results point naturally toward several extensions.

The most important next step would be to complement the present-day-mass analysis with progenitor-based selections, for example using M_{peak} or V_{peak} . This

would help disentangle genuine structural differences from remapping between present-day mass bins caused by differential stripping or disruption.

A second priority would be to incorporate explicit time evolution. Following the same subhalo populations across multiple snapshots would allow one to test whether the low-mass DMO CDM–vSIDM separation is stable in time, whether collapse candidates remain exceptional or become more common at late times, and whether environmental signatures strengthen after infall.

Third, the environmental analysis could be tightened by repeating the distance-binned profile comparison in bins of r/R_{200c} rather than in fixed physical kpc, and by imposing narrower subhalo-mass cuts within each host-mass panel. This would reduce environmental and mass-mixing effects and make it easier to compare the strength of the distance trend across host masses. A direct matched comparison between CDM and vSIDM at fixed distance bin would be particularly valuable, since it would allow the relative magnitude of environmental and model-to-model variation to be quantified explicitly rather than inferred visually.

A particularly important extension concerns the collapse-candidate population. Instead of relying on a few extreme profiles in stacked panels, one could quantify the fraction of systems occupying the high-density tail, follow those objects across multiple snapshots, and test explicitly for the expected temporal transition from core growth to renewed central steepening. This would provide a much firmer basis for distinguishing genuine gravothermal evolution from ordinary tail scatter in the surviving population.

Finally, extending the matched HYDRO analysis across all available dark-matter models, especially for the SHMF, would help determine whether the qualitative hierarchy found here remains true more generally: namely, that structure-based diagnostics are more sensitive to CDM–vSIDM differences than abundance or radial-count statistics once baryonic physics is included.

5.5 Final remarks

The main conclusion of this thesis is that, in the AIDA–TNG simulations studied here, CDM and vSIDM do produce measurably different surviving subhalo populations, and that the clearest signatures of this difference lie in internal dark-matter structure rather than in abundance alone. In DMO, the strongest population-level separation appears in the lower-mass profile stacks, where vSIDM subhaloes are typically less centrally dense than their CDM counterparts over the resolved inner radial range. By contrast, the cumulative radial distributions show only modest and non-uniform model dependence in the present sample, and do not provide a discriminator between CDM and vSIDM that is as robust as the direct profile-based structural diagnostics. In HYDRO, these model-dependent trends become more difficult to isolate, not because they necessarily disappear, but because baryonic physics and orbital environment broaden the structural diversity of the surviving population and complicate any direct mapping between present-day structure and dark-matter microphysics.

The thesis therefore supports a qualified but positive conclusion. Dark-matter microphysics can leave detectable and astrophysically meaningful signatures in subhalo structure and survival, but the visibility of those signatures depends strongly

on the diagnostic used and on whether baryonic effects are included. In particular, baryonic and environmental processes are strong enough that any attempt to use subhalo structure as a discriminator of dark-matter physics must treat them not as secondary complications, but as part of the problem itself. At the same time, the fact that some diagnostics do not provide a clean separation across the full subhalo mass range does not make them uninformative: individual mass regimes that show stronger or more systematic differences may still provide useful points of contact with observational data, especially where those regimes correspond to the scales probed by satellite kinematics, strong lensing, or stellar-stream perturbations.

Bibliography

- Aarseth S. J., 1985, in Goodman J., Hut P., eds, IAU Symposium Vol. 113, Dynamics of Star Clusters. pp 251–258
- Barnes J., Hut P., 1986, *A hierarchical $O(N \log N)$ force-calculation algorithm*, , 324, 446
- Bosma A., 1981, *21-cm line studies of spiral galaxies. I. Observations of the galaxies NGC 5033, 3198, 5055, 2841, and 7331.*, , 86, 1791
- Boylan-Kolchin M., Bullock J. S., Kaplinghat M., 2011, *Too big to fail? The puzzling darkness of massive Milky Way subhaloes*, , 415, L40
- Bullock J. S., Boylan-Kolchin M., 2017, *Small-Scale Challenges to the Λ CDM Paradigm*, , 55, 343
- Correa C. A., 2021, *Constraining velocity-dependent self-interacting dark matter with the Milky Way’s dwarf spheroidal galaxies*, , 503, 920
- Cruz A., et al., 2025, *Dwarf diversity in Λ CDM with baryons* ([arXiv:2510.11800](https://arxiv.org/abs/2510.11800)), <https://arxiv.org/abs/2510.11800>
- Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *The evolution of large-scale structure in a universe dominated by cold dark matter*, , 292, 371
- Despali G., Vegetti S., 2017, *The impact of baryonic physics on the subhalo mass function and implications for gravitational lensing*, , 469, 1997
- Despali G., Moscardini L., Nelson D., Pillepich A., Springel V., Vogelsberger M., 2025, *Introducing the AIDA-TNG project: Galaxy formation in alternative dark matter models*, , 697, A213
- Heinze F. M., Despali G., Klessen R. S., 2024, *Not all subhaloes are created equal: modelling the diversity of subhalo density profiles in TNG50*, , 527, 11996
- Kaplinghat M., Keeley R. E., Linden T., Yu H.-B., 2014, *Tying Dark Matter to Baryons with Self-Interactions*, , 113, 021302
- Klypin A., Kravtsov A. V., Valenzuela O., Prada F., 1999, *Where Are the Missing Galactic Satellites?*, , 522, 82
- Lovell M. R., Frenk C. S., Eke V. R., Jenkins A., Gao L., Theuns T., 2014, *The properties of warm dark matter haloes*, , 439, 300

- McCarthy M. E., Lux-Zeplin Team 2024, in APS April Meeting Abstracts. p. C14.001
- Nadler E. O., et al., 2021, *Constraints on Dark Matter Properties from Observations of Milky Way Satellite Galaxies*, , [126](#), [091101](#)
- Nadler E. O., Kong D., Yang D., Yu H.-B., 2025, *SIDM Concerto: Compilation and Data Release of Self-interacting Dark Matter Zoom-in Simulations*, , [991](#), [69](#)
- Navarro J. F., Frenk C. S., White S. D. M., 1997, *A Universal Density Profile from Hierarchical Clustering*, , [490](#), [493](#)
- Oman K. A., et al., 2015, *The unexpected diversity of dwarf galaxy rotation curves*, , [452](#), [3650](#)
- Pillepich A., et al., 2018, *Simulating galaxy formation with the IllustrisTNG model*, , [473](#), [4077](#)
- Planck Collaboration et al., 2020, *Planck 2018 results. I. Overview and the cosmological legacy of Planck*, , [641](#), [A1](#)
- Read J. I., Walker M. G., Steger P., 2018, *The case for a cold dark matter cusp in Draco*, , [481](#), [860](#)
- Rubin V. C., Ford Jr. W. K., Thonnard N., 1980, *Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 (R=4kpc) to UGC 2885 (R=122kpc).*, , [238](#), [471](#)
- Spergel D. N., Steinhardt P. J., 2000, *Observational Evidence for Self-Interacting Cold Dark Matter*, , [84](#), [3760](#)
- Springel V., Yoshida N., White S. D. M., 2001, *GADGET: a code for collisionless and gasdynamical cosmological simulations*, , [6](#), [79](#)
- Springel V., et al., 2005, *Simulations of the formation, evolution and clustering of galaxies and quasars*, , [435](#), [629](#)
- Springel V., Frenk C. S., White S. D. M., 2006, *The large-scale structure of the Universe*, , [440](#), [1137](#)
- Springel V., et al., 2008, *The Aquarius Project: the subhaloes of galactic haloes*, , [391](#), [1685](#)
- Wang L., Dutton A. A., Stinson G. S., Macciò A. V., Penzo C., Kang X., Keller B. W., Wadsley J., 2015, *NIHAO project - I. Reproducing the inefficiency of galaxy formation across cosmic time with a large sample of cosmological hydrodynamical simulations*, , [454](#), [83](#)
- Wetzell A. R., Hopkins P. F., Kim J.-h., Faucher-Giguère C.-A., Kereš D., Quataert E., 2016, *Reconciling Dwarf Galaxies with Λ CDM Cosmology: Simulating a Realistic Population of Satellites around a Milky Way-mass Galaxy*, , [827](#), [L23](#)
- White S. D. M., Rees M. J., 1978, *Core condensation in heavy halos: a two-stage theory for galaxy formation and clustering.*, , [183](#), [341](#)