



A Physically Motivated Framework to Compare Pair Fractions of Isolated Low- and High-mass Galaxies across Cosmic Time

Katie Chamberlain¹ , Gurtina Besla¹ , Ekta Patel^{2,3,13} , Vicente Rodriguez-Gomez⁴ , Paul Torrey⁵ , Garreth Martin^{1,6,7} , Kelsey Johnson⁸ , Nitya Kallivayalil⁸ , David Patton⁹ , Sarah Pearson^{10,13} , George Privon^{5,8,11} , and Sabrina Stierwalt¹²

¹ University of Arizona, 933 N. Cherry Avenue, Tucson, AZ 85721, USA; katiechambe@arizona.edu

² Department of Astronomy, University of California, Berkeley, 501 Campbell Hall, Berkeley, CA 94720, USA

³ Department of Physics and Astronomy, University of Utah, 115 South 1400 East, Salt Lake City, UT 84112, USA

⁴ Instituto de Radioastronomía y Astrofísica, Universidad Nacional Autónoma de México, Apdo. Postal 72-3, 58089 Morelia, Mexico

⁵ Department of Astronomy, University of Florida, P.O. Box 112055, Gainesville, FL 32611, USA

⁶ School of Physics and Astronomy, University of Nottingham, University Park, Nottingham, NG7 2RD, UK

⁷ Korea Astronomy and Space Science Institute, 776 Daedeokdae-ro, Yuseong-gu, Daejeon 34055, Republic of Korea

⁸ Department of Astronomy, University of Virginia, 530 McCormick Road, Charlottesville, VA 22904, USA

⁹ Department of Physics and Astronomy, Trent University, 1600 West Bank Drive, Peterborough, ON, K9L 0G2, Canada

¹⁰ The Center for Cosmology and Particle Physics, New York University, 726 Broadway, New York, NY 10003, USA

¹¹ National Radio Astronomy Observatory, 520 Edgemont Road, Charlottesville, VA 22903, USA

¹² Physics Department, Occidental College, 1600 Campus Road, Los Angeles, CA 90041, USA

Received 2023 September 22; revised 2023 December 20; accepted 2023 December 30; published 2024 February 16

Abstract

Low-mass galaxy pair fractions are understudied, and it is unclear whether low-mass pair fractions evolve in the same way as more massive systems over cosmic time. In the era of JWST, Roman, and Rubin, selecting galaxy pairs in a self-consistent way will be critical to connect observed pair fractions to cosmological merger rates across all mass scales and redshifts. Utilizing the Illustris TNG100 simulation, we create a sample of physically associated low-mass ($10^8 < M_* < 5 \times 10^9 M_\odot$) and high-mass ($5 \times 10^9 < M_* < 10^{11} M_\odot$) pairs between $z = 0$ and 4.2. The low-mass pair fraction increases from $z = 0$ to 2.5, while the high-mass pair fraction peaks at $z = 0$ and is constant or slightly decreasing at $z > 1$. At $z = 0$ the low-mass major (1:4 mass ratio) pair fraction is $4\times$ lower than high-mass pairs, consistent with findings for cosmological merger rates. We show that separation limits that vary with the mass and redshift of the system, such as scaling by the virial radius of the host halo ($r_{\text{sep}} < 1R_{\text{vir}}$), are critical for recovering pair fraction differences between low-mass and high-mass systems. Alternatively, static physical separation limits applied equivalently to all galaxy pairs do not recover the differences between low- and high-mass pair fractions, even up to separations of 300 kpc. Finally, we place isolated mass analogs of Local Group galaxy pairs, i.e., Milky Way (MW)–M31, MW–LMC, LMC–SMC, in a cosmological context, showing that isolated analogs of LMC–SMC-mass pairs and low-separation (< 50 kpc) MW–LMC-mass pairs are $2\text{--}3\times$ more common at $z \gtrsim 2\text{--}3$.

Unified Astronomy Thesaurus concepts: [Galaxy pairs \(610\)](#); [Interacting galaxies \(802\)](#); [Galaxy evolution \(594\)](#); [Local Group \(929\)](#)

1. Introduction

Galaxy mergers have been studied in detail as a mechanism for driving galaxy evolution, and have been identified as a trigger of, for example, active galactic nuclei (e.g., Hopkins et al. 2008; Treister et al. 2010; Ramos Almeida et al. 2011; Satyapal et al. 2014; Comerford et al. 2015; Glikman et al. 2015; Blecha et al. 2018; Ellison et al. 2019), star formation (e.g., Mihos & Hernquist 1996; Di Matteo et al. 2008; Ellison et al. 2011; Hopkins et al. 2013; Patton et al. 2013; Martin et al. 2017; Hani et al. 2020; Patton et al. 2020; Martin et al. 2021), and changes in morphology (e.g., Conselice 2003; Lotz et al. 2008; Casteels et al. 2014; Patton et al. 2016; Bignone et al. 2017; Rodriguez-Gomez et al. 2017; Martin et al. 2018; Jackson et al. 2019; Snyder et al. 2019; Jackson et al. 2022; Guzmán-Ortega et al. 2023). Constraining the merger rate for galaxies is critical

for quantifying the importance of mergers for galaxy evolution and testing predictions for hierarchical assembly in cold dark matter theory (e.g., Stewart et al. 2009; Hopkins et al. 2010; Rodriguez-Gomez et al. 2015).

In practice, merger rates cannot be measured straightforwardly from observations and, rather, are calculated by converting the observed frequency of a galaxy merger signature (i.e., asymmetry, concentration, pair fraction, etc.) to a merger rate using an observability timescale (e.g., Lotz et al. 2011). However, morphological signatures of mergers are often caused by nonmerger phenomena. For example, low-mass galaxies are more commonly morphologically disturbed by flybys and nonmerger interactions than by mergers (Martin et al. 2021), and star-forming galaxies at high redshift tend to be clumpy and asymmetric even in isolation (Wuyts et al. 2013). Additionally, detection of these signatures can be strongly dependent on image depth or galaxy stellar mass, and identifying tidal features that rely on human classifications may be unreliable (Martin et al. 2022).

Pair fractions, on the other hand, can be calculated independently of morphological information, and thus offer a more robust method of observationally quantifying merger rates across cosmic time. Pair fractions of high-mass ($M_* \gtrsim 10^{10} M_\odot$)

¹³ Hubble Fellow.



galaxies have been well studied across cosmic time both observationally (e.g., Patton et al. 2002; Lin et al. 2004, 2008; Lotz et al. 2011; Ferreras et al. 2014; Man et al. 2016; Duncan et al. 2019) and theoretically (e.g., Rodriguez-Gomez et al. 2015; Snyder et al. 2017, 2023), with theoretical studies typically done in projection for comparison to observational campaigns.¹⁴ Low-mass galaxy pairs ($10^8 < M_* < 5 \times 10^9 M_\odot$) have been studied at low redshift (e.g., Stierwalt et al. 2015; Pearson et al. 2016; Besla et al. 2018; Paudel et al. 2018; Lubert et al. 2022), but are less well understood across cosmic time owing to the difficulty in observing faint systems outside of the Local Volume. However, JWST (Gardner et al. 2006), as well as the next generation of deep and wide-field surveys from Rubin Observatory¹⁵ (Ivezić et al. 2019) and Roman Space Telescope¹⁶ (Spergel et al. 2015), will significantly revolutionize our ability to identify such systems at high redshift ($z \gtrsim 10$; Behroozi et al. 2020) and in abundance at lower redshift ($z \lesssim 6$; Robertson et al. 2019a, 2019b).

There are reasons to believe that low-mass and high-mass pairs evolve differently as a function of time. For example, semi-empirical and cosmological studies find that galaxy merger rates vary both with redshift and the mass of the most massive galaxy of the pair (see, e.g., Guo & White 2008; Stewart et al. 2009; Hopkins et al. 2010; Rodriguez-Gomez et al. 2015; Martin et al. 2021). It is thus reasonable to assume that the pair fractions of these two mass scales reflect these evolutionary differences as well.

The redshift evolution of low- and high-mass pair fractions has not yet been studied in simulations in a self-consistent way, where high-mass and low-mass pairs are selected from simulations using otherwise equivalent selection criteria. We aim to characterize the redshift behavior of pair fractions of low-mass and high-mass pairs, independent of environmental and projection effects, and to create a robust framework in which to fairly compare pair fractions of different mass scales across cosmic time.

In particular, we take the approach of consistently selecting physically associated pairs of low-mass ($10^8 < M_* < 5 \times 10^9 M_\odot$) and high-mass ($5 \times 10^9 < M_* < 10^{11} M_\odot$) galaxies. Specifically, we identify major (1–1/4) and minor (1/10–1/4) stellar mass ratio pairs from $z = 0$ to 4.2 in the IllustrisTNG cosmological simulation, TNG100. We require that pairs are part of the same Friends-of-Friends (FoF) group and that no other more massive perturbers are nearby, which allows us to ensure that the recovered pair fractions are inherent to the population of selected pairs rather than a function of environment.

Typically, pair fraction studies via simulations and observations apply physical projected separation cuts that are constant over time and do not vary with the mass of the target system. In this study, we identify how the application of a fixed-separation criterion affects inferred pair fractions, and show that a time- and mass-evolving separation cut is necessary to permit equitable pair fraction comparisons across different mass regimes. Although implementing alternative selection criteria in future observational pair fraction studies may not always be strictly necessary, we recommend that pair selection criteria be reevaluated for observational studies that seek to compare pair fractions as a function of mass or redshift.

Finally, there are a number of galaxy pairs in the Local Group that are mass analogs of the isolated pairs in our sample. For example, the Milky Way (MW)–M31 system is a high-mass major pair, the MW–LMC and M31–M33 systems are high-mass minor pairs, and the LMC–SMC are a low-mass minor pair. Studies have examined the frequency of such configurations, particularly at low redshift or in the form of progenitor systems of the Local Group in cosmological simulations (Boylan-Kolchin et al. 2011; Fattahi et al. 2013; Patel et al. 2017a; Geha et al. 2017; Mao et al. 2021), but the prevalence of such pairs in isolation has not yet been quantified as a function of redshift. We utilize our pair data set to quantify the likelihood of finding isolated mass analogs of Local Group pairs as a function of redshift, particularly when the present-day separations of these systems are folded in.

This paper is structured as follows. In Section 2, we outline our methodology for selecting physically associated high-mass and low-mass pairs from the TNG simulation. In Section 3, we provide an overview of the properties of the selected sample, including the number of primaries, pairs, and their stellar mass ratios as a function of time. We present the time-evolving pair fraction of high-mass and low-mass pairs in Section 4, and show how they change for different separation criteria. In Section 5, we give context to our results by drawing comparisons to Local Group pairs and other Illustris-based pair fraction studies, and discuss implications for observational campaigns. Finally, we summarize our results and conclusions in Section 6.

2. Methodology

We aim to quantify and characterize the frequency, or pair fraction, of low-mass ($M_* = 10^8\text{--}5 \times 10^9 M_\odot$) and high-mass ($M_* = 5 \times 10^9\text{--}10^{11} M_\odot$) galaxy pairs in cosmological simulations as a function of cosmic time. To this end, we utilize the IllustrisTNG suite of simulations to select subhalo pairs as a function of redshift, according to the selection criteria outlined in the remainder of this section.

In Section 2.1, we provide motivation for and details of the simulation utilized in this study. Section 2.2 outlines the initial mass cuts used to define our Subhalo Catalog, to which we will add stellar mass information. Section 2.3 outlines the abundance matching prescription used to associate dark matter subhalos with stellar masses, and the creation of the Subhalo + Stellar Mass Catalog, from which we will select pairs. Section 2.4 describes the second set of selection criteria that we use to finally construct the Full Pair Catalog.

2.1. Simulation Details

The IllustrisTNG project (Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018; Pillepich et al. 2018; Springel et al. 2018) consists of a suite of dark-matter-only N -body and full physics cosmological simulations that adopt the Planck2015 Lambda cold dark matter (Λ CDM) cosmology (Planck Collaboration et al. 2016).

In this study, we utilize data from TNG 100-1, the main high-resolution, full physics simulation of a $(110.7 \text{ Mpc})^3$ volume (hereafter TNG100).¹⁷ This simulation follows the evolution of baryons and 1820^3 dark matter particles from $z = 127$ to $z = 0$. The volume of this simulation is sufficiently large, and the resolution is sufficiently high, to conduct a simultaneous

¹⁴ The studies cited here have limits on the projected separation of their sample, such that the projected separations range between 5 kpc and ~ 140 kpc.

¹⁵ <https://rubinobservatory.org/>

¹⁶ <https://roman.gsfc.nasa.gov/>

¹⁷ We will also use data from TNG100-1-Dark, the main high-resolution dark-matter-only run (hereafter TNG100-Dark), as presented in Section 4.3.

analysis of the statistics of both low-mass and high-mass galaxy pairs, as shown by studies of pair statistics in similarly sized volumes (Sales et al. 2013; Patel et al. 2017a, 2017b; Besla et al. 2018). We utilize the group catalogs produced by the SUBFIND algorithm (Springel et al. 2001; Dolag et al. 2009). This catalog was created using the FoF algorithm (Davis et al. 1985), which links nearby dark matter particles to define large halos of associated particles, from which subhalos are identified as overdense and gravitationally bound dark matter structures. In addition, we utilize the merger trees provided by the SUBLINK algorithm (Rodríguez-Gomez et al. 2015), which tracks subhalos between snapshots, enabling us to trace the mass evolution of our selected subhalos.

2.2. Choosing Low- and High-mass Subhalos in TNG100

In this section, we outline the steps to create the Subhalo Catalog, which we will utilize to assign stellar masses to each of the dark matter subhalos.

To assess whether the pair fractions of low-mass and high-mass pairs evolve in fundamentally different ways as a function of redshift, rather than as a product of their environmental conditions, we will focus on galaxy pairs in low-density environments only.¹⁸ As such, we ensure that our target subhalos are sufficiently isolated by placing limits on the virial mass of the FoF group to which our low-mass or high-mass samples belong.

The following selection process is repeated for each snapshot of the simulation over a redshift range of $z=0-4.2$. We stop our analysis at $z=4.2$ since the population of massive subhalo pairs falls off rapidly at larger redshifts, where the sample size of high-mass primaries falls well below 100 subhalos per snapshot.

At each snapshot, we first apply a cut on the FoF group virial mass, M_{vir} , given by `Group_M_TopHat200` in the TNG100 group catalogs.¹⁹ We define the group mass range for low-mass and high-mass groups as

$$\begin{aligned} \text{low mass: } M_G &= 8 \times 10^{10} - 5 \times 10^{11} M_\odot \\ \text{high mass: } M_G &= 10^{12} - 6.5 \times 10^{12} M_\odot. \end{aligned}$$

The FoF group mass criteria is fixed for all redshifts, which means that some low-mass groups at high z may be the progenitors of high-mass systems at $z=0$.

By requiring that the FoF group virial mass does not exceed the above limits, we ensure that there are no subhalos more massive than these limits that will perturb the dynamical state of identified pairs. For example, selecting low-mass pairs from the low-mass FoF groups ensures that the selected pairs are not satellite systems of high-mass subhalos. Over 99% of subhalos selected from FoF groups in these mass limits are not within 1.5 Mpc of a more massive perturber at $z=0$.

From the set of FoF groups that pass the group mass cut, we create a catalog of all subhalos within each FoF group that pass a current “minimum mass” threshold. At the given snapshot,

we require a minimum subhalo mass, M_h , to be

$$\text{minimum subhalo mass: } M_h > 1 \times 10^9 M_\odot.$$

M_h is given by the `SubhaloMass` field in the group catalogs, and is the total combined mass of all bound dark and baryonic matter. This “minimum mass” threshold ensures that subhalos are resolved into >100 dark matter particles, and thus should be robustly identifiable using the SUBFIND and SUBLINK algorithms (Rodríguez-Gomez et al. 2015).

For each subhalo that passes both criteria, i.e., the subhalo belongs to either a high-mass or low-mass FoF group and also passes the “minimum mass” threshold, we use the SUBLINK merger trees to identify the maximum subhalo mass achieved by the subhalo, M_{peak} . We consider only the given snapshot and all previous snapshots in the determination of M_{peak} .

All subhalos in a given snapshot that pass both the FoF group and “minimum mass” selection criteria form our full sample of subhalos, which we call the Subhalo Catalog. The catalog contains subhalos at each snapshot from $z=0$ to 4.2 (snapshot numbers 20–99) and their associated properties, namely, Subhalo ID, FoF Group Number, current subhalo mass M_h , and peak subhalo mass M_{peak} . At $z=0$, this selection process results in 44,656 subhalos in low-mass groups and 38,350 subhalos in high-mass groups.

From here, we will use additional selection criteria to identify pairs of subhalos in each group.

2.3. Abundance Matching

We utilize a stellar mass to halo mass (SMHM) relationship to assign stellar masses to each of the subhalos in our Subhalo Catalog. There are a few reasons we opt to assign stellar masses to our subhalos, rather than use those computed directly from the stellar particles in TNG100.

First, and primarily, utilizing an abundance matching prescription enables the direct comparison of results between dark-matter-only and full hydrodynamics simulations, since the stellar masses are assigned in an identical and prescriptive way. We make a brief note of results for our equivalent analysis using the dark-matter-only, TNG100-Dark simulation, in Section 4.3.

Second, while the SMHM function of IllustrisTNG at $z=0$ closely reproduces the profile of the SMHM relation from various abundance matching and semi-empirical models (Pillepich et al. 2018; Nelson et al. 2019), using stellar masses as calculated from abundance matching allows us to avoid any simulation-dependent stellar mass effects. In particular, we would like to avoid a dependence between our results and the particular subgrid physics model implemented in IllustrisTNG.

Third, abundance matching allows us to account for the observed spread in the SMHM relationship, as we can sample the relation many times to get a distribution of stellar masses for each subhalo in the simulation. Otherwise, we would only be able to perform this analysis once given the set of stellar masses from the simulation. This is particularly important in the low-mass regime ($M_h \lesssim 10^{11} M_\odot$) where the scatter in the SMHM function is large, between ~ 0.3 for $M_h \sim 10^{10} M_\odot$ up to ~ 1 dex for $M_h \sim 10^8 M_\odot$ (Munshi et al. 2021).

We use the abundance matching relationship presented in Moster et al. (2013). The SMHM relationship therein provides an analytic prescription to assign stellar masses to dark matter halos as a function of subhalo mass and redshift, and includes terms to account for the systematic scatter in the SMHM

¹⁸ Note that our pair fraction calculations are thus specifically for isolated systems, and the global pair fraction (including both isolated and nonisolated pairs) will be different.

¹⁹ The mass `Group_M_TopHat200` is the mass enclosed by a sphere with mean density $\Delta_c^* \rho_c$, where Δ_c is the overdensity constant from Bryan & Norman (1998) and ρ_c is the critical density of the Universe at the time calculated. The corresponding virial radius in TNG100 is given by `Group_R_TopHat200`.

relationship, with a larger scatter at lower halo masses. We were careful to choose the input subhalo mass and redshift that would assign the most accurate stellar masses to each halo given their individual histories.

The abundance matching prescription is calibrated for centrals of FoF groups; thus, we elect to use the maximum subhalo mass M_{peak} in the stellar mass calculation, rather than the subhalo mass at the given snapshot (see Besla et al. 2018). Using M_{peak} mass allows us to remain robust to scenarios in which a secondary has formed most of its stars, then loses a significant portion of its dark matter content through tidal interactions with a primary, but retains the bulk of its stellar content. In fact, Munshi et al. (2021) found that the stellar mass of subhalos at $z=0$ in the ‘‘Marvel-ous Dwarfs’’ and ‘‘DC Justice League’’ zoom simulations are more closely correlated with M_{peak} than the $z=0$ halo mass for halos with peak halo mass $10^8 < M_{\text{peak}} < 10^{11} M_{\odot}$.

We also use the current redshift of the given snapshot, z_{snap} , of each subhalo in the stellar mass calculation. Using z_{snap} means that we account for changes in the stellar mass of both the primary and secondary halo, even after the secondary has entered the primary’s halo. This assumption is consistent with findings from Akins et al. (2021), which found that massive dwarf satellites ($M_* \approx 10^8\text{--}10^9 M_{\odot}$) entering MW-mass host halos are rarely quenched, and with Geha et al. (2013), which found that dwarfs >1.5 Mpc from a MW-type galaxy are often star-forming and rarely quenched. Additionally, the SAGA survey has found that large satellites of MW-type galaxies are often very blue, with infall into the halo spurring high rates of star formation due to the large gas fraction in dwarfs (Mao et al. 2021). Thus, our abundance matching prescription is of the form $M_* = f(M_{\text{peak}}, z_{\text{snap}})$.

We calculate the stellar mass for a given subhalo by Gaussian sampling each of the fitting parameters of the analytic framework from Moster et al. (2013), in order to account for the spread in the SMHM relationship. We generate a single stellar mass ‘‘realization’’ using an independent draw from the SMHM distribution to calculate a stellar mass for each subhalo of the full Subhalo Catalog. For each snapshot, we repeat this process 1000 times to generate 1000 separate realizations of assigned stellar masses for the Subhalo Catalog. The resulting catalog is called the Subhalo + Stellar Mass Catalog, and consists of the set of all subhalos from the Subhalo Catalog, as well as the 1000 stellar mass realizations. Each realization is treated as an independent sample of galaxy stellar masses, which will allow us to report realistic spreads of pair properties.

2.4. Pair Selection

Starting from the Subhalo + Stellar Mass Catalog, we outline the pair-matching process used to generate the Full Pair Catalog below. At each redshift, and for each stellar mass realization, we will identify subhalo pairs consisting of a ‘‘primary’’ and a ‘‘secondary,’’ where primaries are the more massive of the pair by stellar mass.

2.4.1. Selecting Primaries

Primary galaxies (equivalently, ‘‘primary subhalos’’) are the most massive galaxy of their FoF group by stellar mass, such that each group will have a singular primary galaxy. Each stellar mass realization is treated independently, and so the

subhalo identified as the primary of a group may change between stellar mass realizations.

At a given snapshot, and for each stellar mass realization, we rank-order the subhalos of each FoF group by stellar mass. Primaries are defined as the subhalo with the highest stellar mass (M_{*1}) in their FoF group that passes the following criteria:

$$\text{low mass primaries: } 10^8 < M_{*1} < 5 \times 10^9 M_{\odot}$$

$$\text{high mass primaries: } 5 \times 10^9 < M_{*1} < 10^{11} M_{\odot}.$$

The stellar mass criteria is fixed, and does not change as a function of redshift. At $z=0$, the stellar mass range for low-mass primaries corresponds to isolated analogs of the LMC or M33, while the high-mass primaries represent isolated analogs of the MW or M31.

Since our selection is based on stellar mass, and there is a large spread in the SMHM relation, a primary subhalo may not be the most massive subhalo in terms of total subhalo mass.

2.4.2. Selecting Secondary Companions

As before, the selection of secondary companions occurs independently for each stellar mass realization, and for each snapshot. Secondary subhalos are defined as the second most massive subhalo in a FoF group by stellar mass (M_{*2}). Secondary subhalos must also have a stellar mass ratio with respect to primaries of

$$\text{stellar mass ratio: } M_{*2}/M_{*1} > 1/10.$$

We do not include companions with a stellar mass ratio $M_{*2}/M_{*1} < 1/10$ as we will be limiting our pair sample to traditional definitions of major and minor pairs, which are typically defined to have stellar mass ratios $M_{*2}/M_{*1} > 1/10$ (i.e., Lotz et al. 2011; Rodriguez-Gomez et al. 2015; Snyder et al. 2017; Duncan et al. 2019; Wang et al. 2020; Guzmán-Ortega et al. 2023).

Our subhalo ‘‘minimum mass’’ threshold of $M_h = 10^9 M_{\odot}$ (described in Section 2.2) corresponds to a mean stellar mass of approximately $10^6 M_{\odot}$ at $z=0$. This stellar mass is well below the 1/10 criteria for the lowest stellar mass primary considered ($M_{*1} > 10^8 M_{\odot}$), ensuring that our pair sample will be complete even at the lowest stellar masses considered.

2.4.3. Creating the Pair Catalog

Pairs consist of a single primary and secondary in a FoF group as defined via the criteria above, such that only one pair is identified in a single FoF group. Each pair is categorized as either a major or minor pair based on the stellar mass ratio between the primary and secondary. Following, for example, Lotz et al. (2011) and Rodriguez-Gomez et al. (2015), the stellar mass ratio criteria for major and minor pairs are

$$\text{major pair: } \frac{1}{4} \leq \frac{M_{*2}}{M_{*1}} < 1$$

$$\text{minor pair: } \frac{1}{10} \leq \frac{M_{*2}}{M_{*1}} < \frac{1}{4}.$$

For each pair, we calculate the physical separation (proper kiloparsec, not comoving kiloparsec) between the two subhalos using the `SubhaloPos` field from the subhalo catalogs. We require that each pair have a minimum separation of 10 kpc between primary and secondary to limit the impact of very

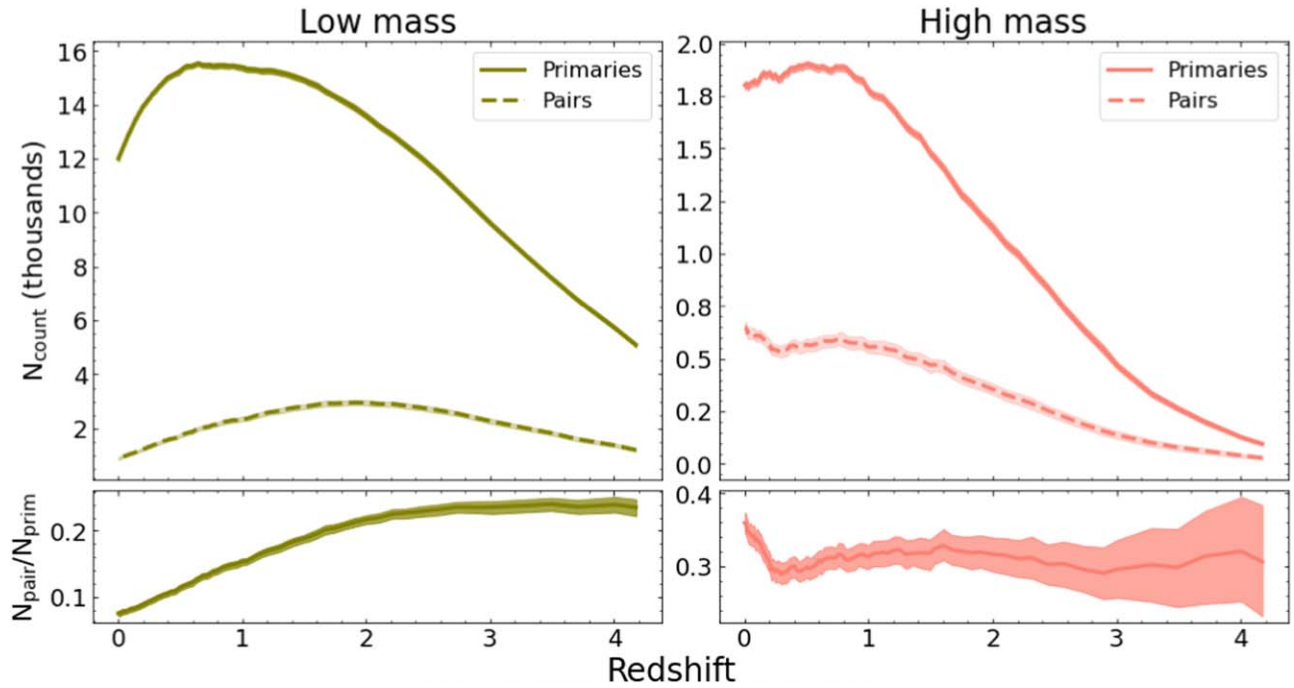


Figure 1. Top: number of isolated low-mass (left) and high-mass (right) primaries and pairs as a function of redshift. The solid and dashed lines represent the median of the set of total counts from each of the 1000 stellar mass realizations in the Full Pair Catalog, while the shaded regions depict the 1st–99th percentile spread of the median. There are approximately 8 times as many low-mass primaries as high-mass primaries. The low-mass primary count (left) peaks at $z = 1$ (with $\sim 15,000$ primary subhalos per realization), while the low-mass pair count peaks at $z = 2$ (with ~ 3000 pairs per realization). The high-mass primary count (right) peaks at $z \sim 1$ (with ~ 1900 primary subhalos per realization), while the high-mass pair count peaks at $z = 0$ (with ~ 700 pairs per realization). Bottom: total pair fraction (fraction of primaries with a major or minor secondary) as a function of redshift (see Section 3 for calculation details). The low-mass total pair fraction is approximately flat between $z = 2.5$ and 4, and decreases from $z = 2.5$ to 0. The high-mass total pair fraction is flat or decreasing from $z = 1$ to 4, but peaks sharply between $z = 0$ and 0.25.

close subhalos becoming indistinguishable in SUBFIND as a result of the resolution limit due to the softening length.

If a primary subhalo does not have a companion that meets the stellar mass ratio criteria and separation criteria, the subhalo will still be considered a primary. We refer to all primaries that do not have a companion with these criteria as “isolated primaries,” including primaries with a companion that meets the stellar mass ratio criteria but has a separation < 10 kpc. The total number of primaries includes both isolated primaries and those with selected companions and is larger than the total number of pairs.

For each snapshot and each of the 1000 realizations of the stellar mass in Subhalo + Stellar Mass Catalog, we identify a set of isolated primaries and pairs. Additionally, within each realization, no single subhalo can be a part of two separate pairs, such that the primary and secondary of every pair are unique to the pair. The final Full Pair Catalog is the collection of all isolated primaries and pairs at each snapshot and includes the following information. For each isolated primary, we store the current subhalo mass (M_h) and stellar mass from the given realization. For each pair, we store the primary and secondary subhalo masses (M_h) and stellar masses, the pair separation, and the virial radius of the FoF group (see Section 2 for virial definitions).

3. Sample: Overview of Pair Properties

Utilizing the Full Pair Catalog, for each snapshot we compute the total number of primaries (isolated and paired) and pairs (including major and minor) in each of the 1000 realizations. We then compute the median and 1st and 99th

percentile spread on the median over all realizations.²⁰ Additionally, we compute the low- and high-mass total pair fraction for each individual realization, defined here as the ratio of the total number of pairs (N_{pairs}) to the total number of isolated and paired primaries ($N_{\text{primaries}}$):

$$f_p = \frac{N_{\text{pairs}}}{N_{\text{primaries}}}.$$

We again compute the median and spread over all 1000 realizations.

Figure 1 shows the median number (solid and dashed lines) of identified low- and high-mass primaries and pairs over the redshift range $z = 0$ –4. The shaded regions show the 1%–99% spread of the set.

The number of identified primaries is lowest at $z = 4$, and rises to a maximum around $z = 1$ for both low- and high-mass primaries. The median count of low-mass primaries (green solid line in top left panel) reaches a maximum of $15,545^{+89}_{-88}$ halos at $z \sim 0.6$, then decreases by $\sim 22\%$ to $12,002^{+85}_{-86}$ halos at $z = 0$. The count of high-mass primaries (pink solid line in top-right panel) reaches a maximum of 1901^{+16}_{-19} at $z = 0.5$, and slightly declines to $z = 0$. Our sample of high-mass primaries represents approximately 20% of all subhalos in TNG100 with the same range of stellar masses at $z = 0$. There are roughly 8 times as many low-mass primaries as high-mass primaries. Note that the comoving volume of IllustrisTNG is the same at all redshifts.

²⁰ The spread on the median of each realization is very small. Thus, we opt to show the 1st and 99th percentile spread rather than those which align with traditional definitions of 1σ or 2σ measurements.




Figure 2. Stellar mass ratio distribution of all low-mass (top) and high-mass (bottom) pairs for all 1000 stellar mass realizations combined at each redshift. Major pairs (solid lines) are defined as pairs with mass ratio $M_{*2}/M_{*1} > 1/4$, while minor pairs (dotted lines) are defined as pairs with stellar mass ratio $1/10 < M_{*2}/M_{*1} < 1/4$. Overall, the stellar mass ratio distribution of major and minor pairs of low- and high-mass galaxies show little evolution from $z = 4$ (right) to $z = 0$ (left). Major pairs make up 51%–55% of the total sample of pairs at every redshift for both low- and high-mass pairs.

Unlike the primary count, the pair counts for low- and high-mass pairs peak at very different redshifts. The count of low-mass pairs (green dashed line) peaks much earlier, at $z = 1.9$ with 2956^{+99}_{-95} pairs, and decreases to 896^{+51}_{-50} pairs at $z = 0$. The pair count for high-mass galaxies (pink dashed line), on the other hand, behaves more similarly to the primary count, increasing from $z = 4$ to $z \sim 1$, then decreasing to $z = 0.25$ before peaking with 647 ± 27 pairs at $z = 0$.

The bottom panel of Figure 1 shows the total pair fraction for low-mass and high-mass pairs, or, equivalently, the fraction of primaries with a major or minor companion. The total pair fractions for both low- and high-mass pairs are roughly flat for $z = 2.5$ – 4 , and display opposite behavior for low redshifts between $z = 0$ and 2.5 . The low-mass total pair fraction decreases from $0.233^{+0.008}_{-0.009}$ to 0.075 ± 0.004 , a decrease of roughly 68%, while the high-mass total pair fraction remains flat between $z = 1$ and 2.5 , ranging between 0.275 and 0.351 . At very low redshifts, from $z = 0$ to 0.25 , the high-mass total pair fraction spikes sharply from $0.288^{+0.016}_{-0.013}$ to $0.359^{+0.015}_{-0.014}$, an increase of 37%.

In Figure 2, we show the combined distribution of stellar mass ratios of every pair from all 1000 realizations in the Full Pair Catalog. Major pairs make up 51%–55% of the full sample of pairs at every redshift for both low- and high-mass pairs. In general, the shape of the distribution remains constant from $z = 0$ (left) to $z = 4$ (right) for both low- and high-mass pairs, and changes weakly as a function of mass scale and of redshift. There are between ~ 3.3 and 4.8 times more pairs with mass ratios $\sim 1/4$ than $\sim 1/1$, and about 1.8 – 2.2 times more pairs have $1/10$ than $1/4$. Roughly 2% of the total pair population is a 1:1 mass ratio encounter; this is true for low- and high-mass galaxies and across all redshifts considered.

4. Results: The Frequency of Low-mass and High-mass Pairs

We have created catalogs of isolated low- and high-mass galaxy pairs from $z = 0$ to 4.2 in the TNG100 simulation. In this section, we will analyze the frequency of major and minor pair types across cosmic time, with the goal of identifying potential differences between high- and low-mass galaxies. In Section 4.1, we examine the redshift evolution of the fraction of primaries with a major or minor companion (the “pair fraction”) and compare the results for low-mass and high-mass pairs. In Section 4.2, we examine the redshift evolution of the pair fraction as a function of pair separation. In Section 4.3, we briefly describe how our equivalent pair fraction analysis utilizing the TNG100-Dark simulation compares to the TNG100 results presented in the previous two sections.

In the following analysis, we will treat each stellar mass realization in the Full Pair Catalog as an independent sample. At each snapshot, we calculate the median and 1st and 99th percentile spread on the median of the pair fractions in each of the 1000 realizations of the target pair sample. Each of the following figures shows the median as the solid or dashed lines, and the shaded regions correspond to the spread on the median.

4.1. Major and Minor Pair Fractions

The major and minor pair fractions are computed in the way as defined in Section 3, Equation (1), where N_{pairs} is the total number of major or minor pairs in the sample. For example, the low-mass major pair fraction is the number of low-mass major pairs divided by the number of low-mass primaries, which can also be interpreted as the likelihood of finding a major companion of an isolated low-mass primary.

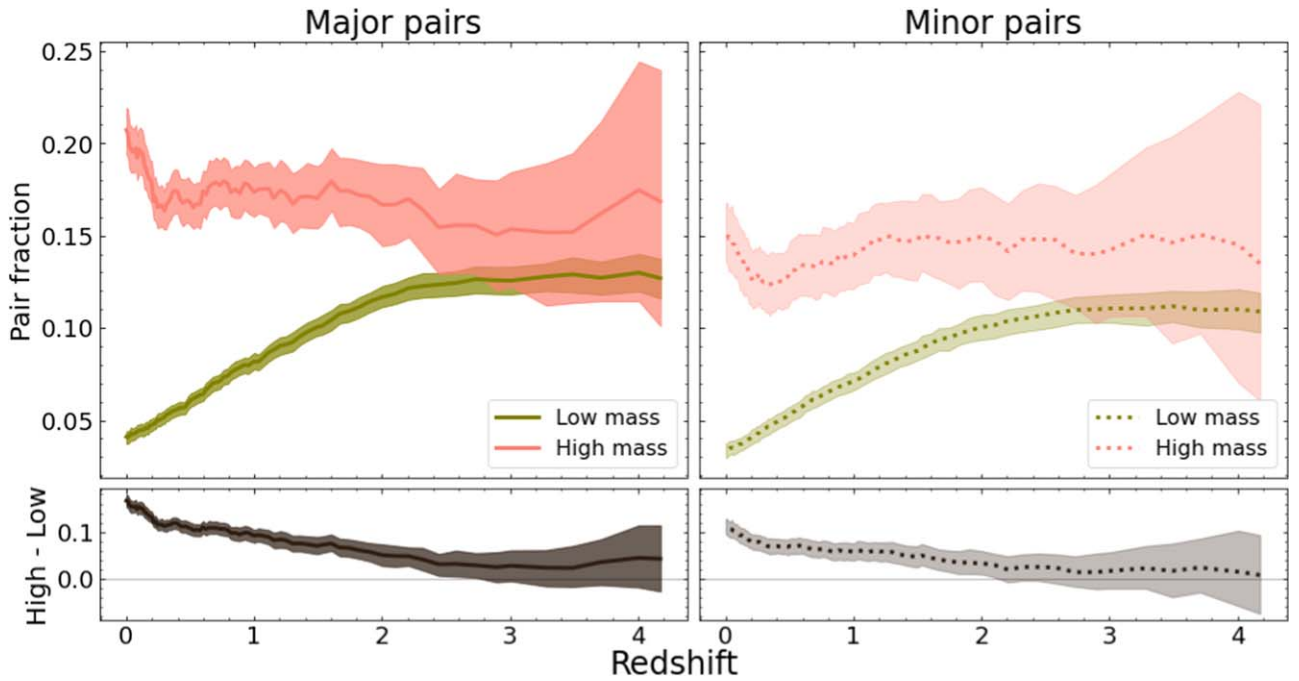


Figure 3. Top: median pair fraction as a function of redshift, defined as the fraction of low-mass or high-mass primaries with a major (solid) or minor (dashed) companion (see Section 2.4). All pairs in the Full Pair Catalog must have a minimum separation of at least >10 kpc, with no constraint on the maximum separation. Shaded areas show the 1st–99th percentile range on the median (solid and dashed lines) from 1000 stellar mass realizations, as discussed at the beginning of Section 4. Low-mass major and minor pair fractions (green) are both at their minima at $z = 0$, and increase by about 200% by $z \sim 2$ – 2.5 , at which point they level off and remain constant from $z = 3$ to 4. On the other hand, high-mass major and minor pair fractions reach their maxima at $z = 0$, then abruptly decline until $z \sim 0.25$ before remaining approximately constant from $z = 1$ to 4. Bottom: the median and 1st–99th percentile range of the subtracted difference between high- and low-mass pair fractions. The difference peaks at $z = 0$ for both major and minor pairs, and declines with increasing redshift. This panel shows that the redshift evolution of the pair fractions of low- and high-mass pairs proceeds differently, particularly at low redshift where pairs are more common for high-mass galaxies than low-mass galaxies.

Figure 3 shows the pair fractions calculated for low- and high-mass pairs (including both major and minor pairs separately) as a function of time for $z = 0$ – 4 . Note that, as discussed in Section 2.4, we adopt a lower stellar mass ratio floor of $1/10$, and consequently the total pair sample is dominated by major pairs.

The low-mass pair fraction evolves distinctly from that of high-mass pairs. Low-mass pair fractions for both major and minor pairs (green) are approximately constant from $z = 3$ to 4, then decline monotonically to $z = 0$. At $z = 0$, the major pair fraction is $0.041^{+0.003}_{-0.004}$, while the minor pair fraction is 0.034 ± 0.004 . At $z = 3$, the pair fractions have increased by 207% to $0.126^{+0.008}_{-0.007}$ for major pairs and 0.111 ± 0.008 for minor pairs.

High-mass pair fractions (pink) remain approximately constant for $z > 1$, where the median of the major pair fraction fluctuates between 0.150 and 0.179, and the median for minor pairs fluctuates between 0.135 and 0.151. Between $z = 0.3$ and $z = 0$, the major pair fraction increases from $0.163^{+0.011}_{-0.010}$ to $0.207^{+0.012}_{-0.014}$, while the minor pair fraction increases from $0.125^{+0.017}_{-0.014}$ to $0.152^{+0.016}_{-0.017}$.

The bottom panels of Figure 3 show the low-mass pair fraction subtracted from the high-mass pair fraction (labeled “High – Low”). The difference between high- and low-mass pair fractions increases with decreasing redshift, peaking at $z = 0$ with a difference of 0.166 ± 0.013 for major pairs and 0.111 ± 0.017 for minor pairs. Thus, at low redshift we expect both major and minor pairs to be more common in high-mass galaxies than in low-mass galaxies. From $z = 2.5$ to 4, the difference is approximately 0, and thus both major and minor pairs of high- and low-mass galaxies are equivalently common at high redshift.

Overall, these results show that low-mass and high-mass pair counts evolve differently over time, particularly at very low redshift, despite the pair fractions being roughly equal at higher redshift. The implications for the difference in the evolution of pair fractions for low-mass and high-mass pairs across time are discussed in detail in Section 5.

4.2. Major Pair Fractions as a Function of Separation

In this section, we analyze subsets of the low- and high-mass major pairs from the previous section that pass additional separation criteria. We study two different sets of separation criteria to compare the resulting pair fractions to the full sample shown in Figure 3. In Section 4.2.1, we apply a separation criterion that requires secondaries to be within a given factor of the virial radius of the FoF group, which is a reasonable proxy for the virial radius of the primary. In Section 4.2.2, we use a range of limits on the 3D pair separation, where limits are consistent with values adopted in the literature.

In all cases, solid lines correspond to the median pair fraction and shaded regions to the 1st–99th percentile, as explained at the beginning of Section 4.

4.2.1. Pair Separation Limits as a Function of the Virial Radius of the FoF Group

We calculate the pair fraction for subsamples of low-mass and high-mass major pairs by selecting only pairs that have separations less than some factor of the virial radius of its FoF group. Note that a minimum separation limit of 10 kpc is always applied to all of our pairs (see Section 2.4). The virial

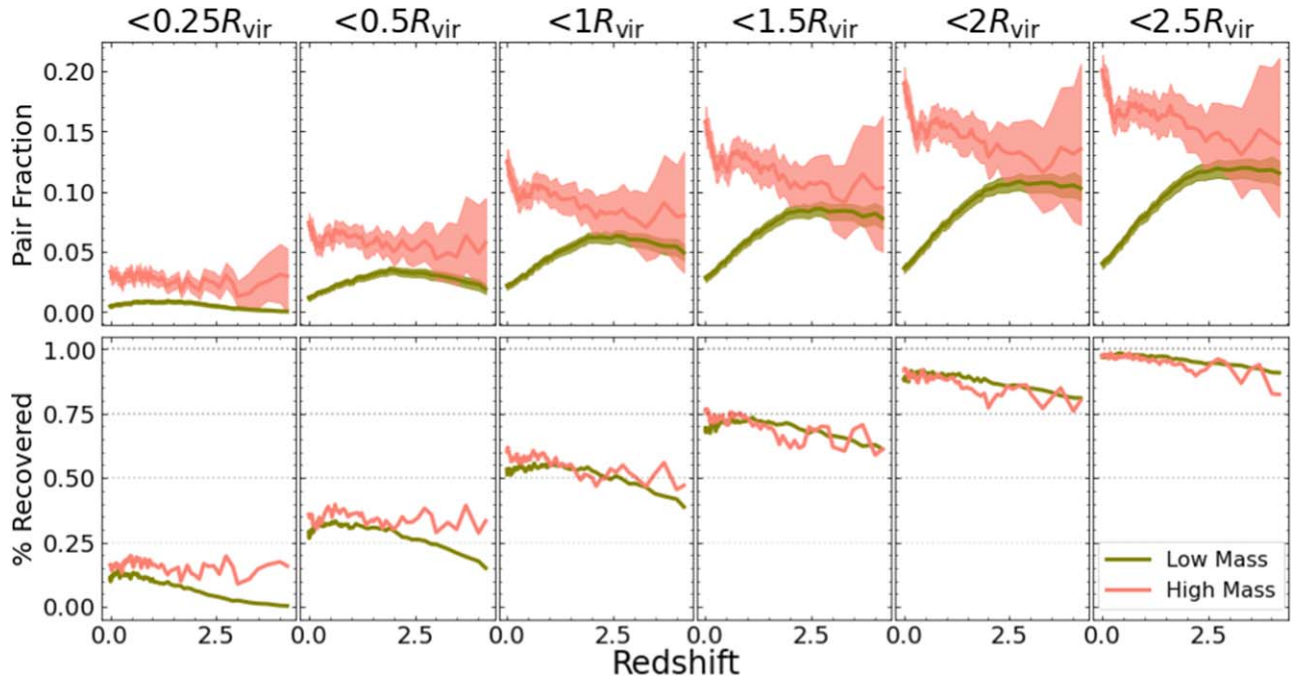


Figure 4. The median and 1st–99th percentile spread are shown by the solid lines and shaded regions, respectively. Top: the median pair fraction of the subset of high-mass (pink) and low-mass (green) pairs with 3D separations within a given factor of the pair’s FoF group virial radius. Such a separation criteria will vary as a function of both the mass of the pair and the redshift. Recovering the redshift evolution seen for the total sample (left panel of Figure 3) requires separation cuts of $r_{\text{sep}} < 1.5R_{\text{vir}}$ for low-mass pairs and $r_{\text{sep}} < 0.5R_{\text{vir}}$ for high-mass pairs. Recovering the differences between high- and low-mass pair fraction trends seen for the total sample at $z > 2$ requires $r_{\text{sep}} < 0.5R_{\text{vir}}$ for both galaxy types. Bottom: the recovery fraction, calculated as the fraction of the total collection of pairs recovered by the subset of pairs at the given separation cut. We find that recovering the redshift evolution of the pair fractions of all galaxy pairs requires a separation cut such that the number of close pairs constitutes more than $\sim 50\%$ of the full sample at all redshifts.

radius is taken from the `Group_R_TopHat200` field in the group catalogs.

Major pairs typically correspond to the two most massive subhalos in their FoF group. The virial mass of their FoF group is reasonably approximated by the combined subhalo mass of both galaxies in the pair. The combined mass of major low-mass galaxy pairs recovers on average 98% of the FoF group mass, and the combined mass of major high-mass galaxy pairs recovers an average of 93%. For example, the mass of the Local Group is dominated by that of the MW and M31 (e.g., Chamberlain et al. 2023). The virial radius of the FoF group is thus reasonably approximated by the virial radius of a halo with a virial mass equal to the combined subhalo mass of the two major pair members.

By focusing our separation criteria on the FoF group virial radius, the separation cut will vary both as a function of redshift and as a function of the combined mass of the pair consistently. For the high-mass FoF groups, the median virial radius at $z = [0, 1, 2, 3, 4]$ is $\sim(348, 206, 134, 97, 76)$ kpc. For the low-mass FoF groups, the median virial radius at $z = [0, 1, 2, 3, 4]$ is $\sim[134, 85, 59, 43, 33]$ kpc.

We choose six subsamples, consisting of pairs with separations less than $0.25, 0.5, 1.0, 1.5, 2.0,$ and $2.5R_{\text{vir}}$. The top row of Figure 4 shows the median pair fraction (solid lines) for high-mass (pink) and low-mass (green) pairs in each of the six subsamples. In the bottom row, we show the recovery fraction, which is the fraction of the total pair sample that is recovered by each subsample, i.e., the number of pairs that pass each separation criterion divided by the total number of pairs in the full sample presented in Section 4.1.

As the pair separation limits increase (left to right), the recovery fraction increases. So too does the pair fraction, as expected since

each consecutive selection cut is less restrictive and will contain more of the full sample. More than 75% of the full pair sample is recovered at all redshifts when the sample contains all pairs within $2R_{\text{vir}}$ (the two rightmost panels of Figure 4).

All subsamples of high-mass pairs display broadly the same pair fraction redshift evolution as the full high-mass sample (pink line in left panel of Figure 3): each of the subsampled high-mass pair fractions peak at $z=0$ and remain roughly constant or decrease at higher redshift. The finer detail trends from the full sample, particularly the upturn at low z , are reasonably captured if $r_{\text{sep}} < 0.5R_{\text{vir}}$, despite the recovery being $\sim 40\%$.

The subsamples of low-mass pairs, however, show a different trend from the full sample (green line in left panel of Figure 3) if the separation cut is too small. For $r_{\text{sep}} < 0.25R_{\text{vir}}$, the pair fraction is flat or decreasing as a function of redshift (especially for $z > 1$), similar to the behavior for high-mass pairs. For $r_{\text{sep}} < 0.5R_{\text{vir}}$ and $r_{\text{sep}} < 1.0R_{\text{vir}}$, the pair fraction rises to a peak at $z \sim 2$ and then decreases. When the separation is limited to $r_{\text{sep}} < 1.5R_{\text{vir}}$, the trends with redshift for the full low-mass sample are recovered, particularly the roughly flat behavior from $z \sim 2-4$, despite a sample recovery of $< 75\%$. We also find that differences between the redshift evolution of high-mass pairs versus low-mass pairs (particularly at $z < 2$) become apparent if $r_{\text{sep}} < 0.5R_{\text{vir}}$ for both galaxy types. Overall, we find that recovering the redshift evolution of the pair fractions of all galaxy pairs requires a separation cut that contains at least the closest $\sim 40\%$ – 50% of pairs from the full sample at all redshifts.

4.2.2. Pair Separation Limits Based on Static 3D Physical Separation

We calculate the pair fraction for subsamples of major low-mass and high-mass pairs by including only pairs with

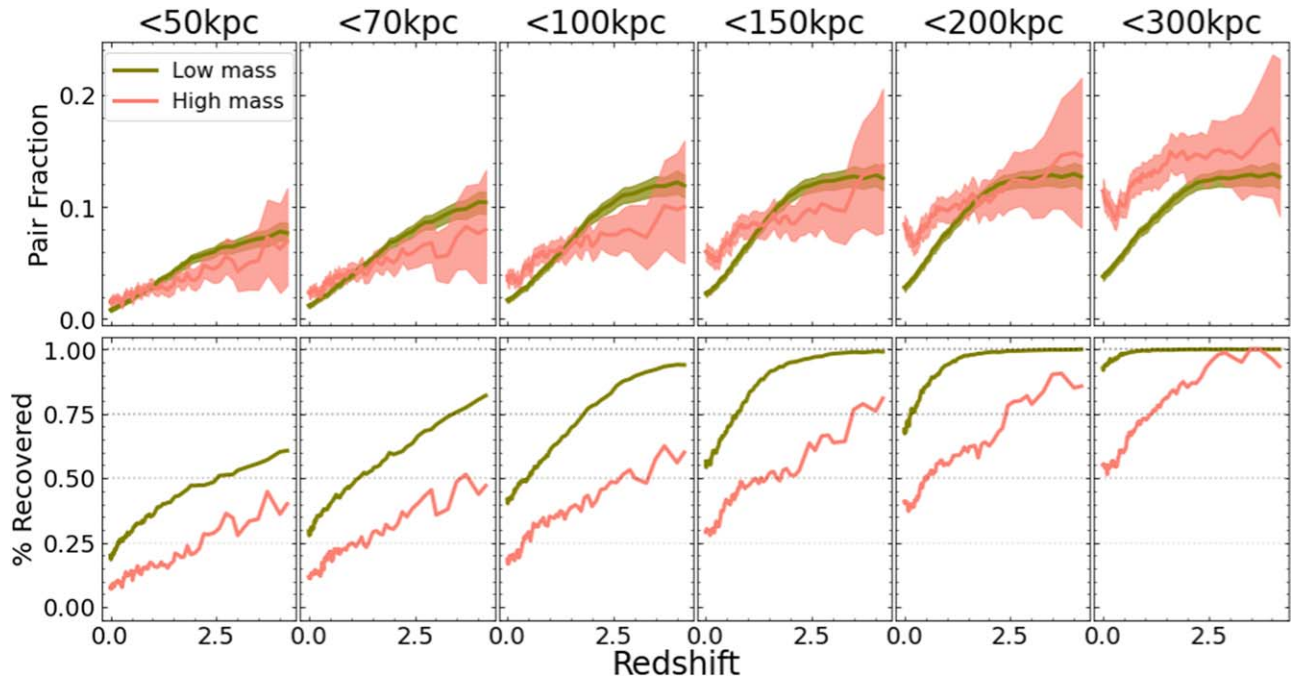


Figure 5. Top: the major pair fraction of the subset of high-mass (pink) and low-mass (green) pairs with 3D separations less than the value given at the top of each column. The median and 1st–99th percentile spread are shown by the solid lines and shaded regions, respectively. The 3D separation cut does not vary as a function of redshift or mass of the system. The general redshift evolution of the full high-mass pair sample (decreasing with redshift for $z > 1$) is not recovered for any physical separation cut, while the behavior of low-mass pair fractions is well recovered by the $r_{\text{sep}} < 150$ kpc subsample. The pair fractions for low-mass and high-mass pairs are virtually indistinguishable for separation cuts less than 70 kpc. Bottom: the fraction of the total collection of pairs in the subset that passes each separation cut. The recovered fraction varies as a function of redshift because a constant separation is chosen. A higher fraction of low-mass pairs is recovered than high-mass pairs at all separation cuts. Less than 50% of all high-mass pairs have separations $r_{\text{sep}} < 70$ kpc at all redshifts, and $>50\%$ of the full high-mass sample is recovered at all redshifts only for $r_{\text{sep}} < 300$ kpc. Between 20% and 60% of low-mass pairs have separations $r_{\text{sep}} < 50$ kpc at all redshifts, while a separation cut of $r_{\text{sep}} < 150$ kpc captures roughly 50%–100% of the low-mass pair population at all redshifts.

separations less than [50, 70, 100, 150, 200, and 300] kpc. Note that these separation criteria do not vary as a function of redshift or mass of the pair, and that a minimum separation limit of 10 kpc is always applied (see Section 2.4). These separation criteria create subsamples of the Full Pair Catalog containing low- and high-mass pairs with separations between 10 and 50 kpc, 10 and 70 kpc, and so on.

Figure 5 shows the median pair fractions for the subsamples of pairs with separations lower than the physical separation listed at the top of each column. The bottom row of the plot shows the recovery fraction, or the number of pairs in each subsample compared to the full sample of major pairs. Again, as the maximum separation increases (left to right), so too does the pair fraction and recovered fraction of the subsample.

The low-mass major pair fraction (green) maintains roughly the same behavior (decreasing with decreasing redshift) for each separation cut, and first recovers the redshift evolution and magnitude of the pair fraction of the full sample (left panel of Figure 3) for a separation cut of $r_{\text{sep}} < 150$ kpc. This is not surprising based on our results from the previous section as this separation is larger than the median virial radius of the sample and the recovery fraction is larger than 0.50 at all redshifts (reaching nearly 100% at $z \sim 3$). The low-mass pair fraction does not change significantly by excluding pairs with separations >150 kpc.

On the other hand, none of the subsamples of the high-mass pairs accurately recover the behavior of the full high-mass sample. In particular, the redshift evolution of the pair fraction between $z = 0.25$ and 2.5 is not readily distinguishable from that of low-mass pairs for any subsample from a physical

separation cut. At higher redshifts, the median pair fraction increases rather than leveling off or decreasing as in the full sample. Even for $r_{\text{sep}} < 300$ kpc, the high-mass pair fraction peaks at $z = 4$, at odds with results using the full pair sample or any of the virial radius cuts in the previous section.

By applying a separation cut that is constant as a function of the redshift and mass of the system, the recovery fraction of the total sample varies markedly as a function of redshift. This is in stark contrast to results from the previous section where the recovery fraction was roughly constant as a function of redshift. Because the same separation cut is applied to low- and high-mass pairs, the recovery fraction for low-mass pairs is universally higher than for high-mass pairs at each separation cut, which erases any differences in redshift evolution between the two galaxy types.

4.3. Comparison between TNG100-1 and TNG100-Dark

We repeated the entirety of our analysis using data from the TNG100-Dark simulation. Each step in our selection criteria was repeated identically, including our selection of group halos, abundance matching process, and selection of primary and secondary halos. We created an equivalent catalog to the Full Pair Catalog containing only subhalo properties from the TNG100-Dark simulation, and calculated the low- and high-mass pair fractions for $z = 0$ –4.2.

We found that the redshift evolution of low-mass and high-mass pair fractions in the TNG100-Dark simulation mimic those of the TNG100 simulation. In particular, low-mass pair fractions remain flat between $z \sim 2$ and 4, and decline from $z = 2$ to a minimum at $z = 0$. High-mass pair fractions likewise

decrease with higher redshift for $z > 1$, and peak at $z = 0$. The numerical value of the pair fraction at all redshifts is $\sim 10\%$ larger in TNG100-Dark for both low- and high-mass pairs and for major and minor pairs, which is likely a result of the suppressed abundance of low-mass halos due to baryonic physics (Chua et al. 2017).

5. Discussion

We performed a pair fraction analysis for low-mass and high-mass pairs in the TNG100 simulation from $z = 0$ to 4.2, utilizing the full spatial information that simulations enable to ensure that the pairs are physically colocated as part of the same FoF group. In this section, we discuss the broader impacts and implications of the behavior of high- and low-mass pair fractions over time.

In Section 5.1, we draw comparisons between our pair sample and various Local Group pairs to quantify the likelihood of finding isolated analogs of such pairs at higher z . In Section 5.2, we compare our results to previous studies of pair fractions that utilize simulations. We discuss the observational implications of the pair fraction behavior of low-mass pairs and high-mass pairs as a function of redshift in Section 5.3. We conclude by discussing some possibilities for the pair fraction difference between low- and high-mass pairs in Section 5.4.

5.1. Implications for Local Group Galaxies

Our study encompasses the stellar mass range and stellar mass ratios of a few well-studied pairs within the Local Group and Local Volume. Thus, we can make predictions for the frequency of finding isolated mass analogs of these systems at $z > 0$ for the first time. For example, isolated mass analogs of the following pairs are included in each specified sample: the LMC–SMC and NGC 4490–4485 (low-mass minor pair),²¹ the MW–LMC and M31–M33 (high-mass minor pair), and the MW–M31 (high-mass major pair). Note that both the NGC 4490–4485 and MW–M31 pairs pass all of our pair selection criteria, while the LMC–SMC, MW–LMC, and M31–M33 pairs are nonisolated pairs within a more massive group environment.

Mass analogs of the MW–M31 pair are most common today (pink solid line in the left panel of Figure 3). Given their large separation of ~ 760 kpc, they would only be included in the $< 2.5R_{\text{vir}}$ panel of Figure 4.

We find that isolated mass analogs of the LMC–SMC and NGC 4490–4485 are roughly 3 times as common at $z > 3$ than at $z = 0$ (green dotted line in the right panel of Figure 3). These high- z pairs may enter a more massive group halo to become satellites of a larger galaxy (like the LMC–SMC; Besla et al. 2007; Patel et al. 2017a), or, if they remain isolated, may continue to merge in a manner similar to NGC 4490–4485 (Pearson et al. 2018).

Isolated mass analogs of the MW–LMC and M31–M33 systems are equally likely at $z > 1$ as they are at present (right panel of Figure 3). If we fold in the present-day separation of the MW–LMC system (~ 50 kpc), this high-mass minor pair becomes rare at low z .²² Specifically, we find that 2% of MWs

would host such a close LMC at $z = 0$, which is consistent with other cosmological studies, such as Patel et al. (2017a), which finds that 3.8% of MW-mass halos host an LMC-mass analog within 50 kpc at $z = 0$. We also find that such systems (high-mass minor pairs with a low separation) are 2–3 times more common at high z . M31–M33 analogs, which have higher separations (~ 200 kpc), are more common than MW–LMC analogs at all redshifts.

5.2. Comparison to Existing Pair Fraction Studies

A direct comparison to pair fractions reported from observations is not straightforward, as observationally selected pairs suffer from contamination due to projected pairs and restrictive separation criteria that exclude more widely separated pairs. Instead, we compare our results to pair fractions and merger rates reported by other studies of the Illustris cosmological simulations to establish the reliability of our findings.

Besla et al. (2018) quantified low-mass galaxy ($2 \times 10^8 < M_* < 5 \times 10^9 M_\odot$) pair fractions at $z \sim 0$ in the Illustris-1 hydrodynamic cosmological simulation, utilizing both projected and 3D pairs. The selection criteria for projected pairs include a projected separation cut, $r_p < 150$ kpc, and a relative line-of-sight velocity difference $\Delta v_{\text{los}} < 150 \text{ km s}^{-1}$.

Searching for pairs in a projected space while having access to the true 3D positions and velocities of the galaxies allowed Besla et al. (2018) to quantify the contamination fraction of false pairs due to projection effects. They found that up to $\sim 40\%$ of identified companions were unrelated but appeared to be close due to projection effects. The projected pair sample also enabled a direct comparison of the cosmologically derived pair fractions with an equivalently selected low-mass galaxy pair fraction from the Sloan Digital Sky Survey (SDSS). They found that the low-mass major pair fraction from the SDSS sample was in good agreement with the simulations. The Besla et al. (2018) study points to the Illustris simulation’s ability to robustly constrain pair fractions of low-mass galaxies, specifically at low z .

We find reasonable agreement with the sample of physical 3D pairs in Besla et al. (2018), which found a major pair fraction between 0.003 and 0.018 over a mass range of $2 \times 10^8 < M_* < 5 \times 10^9 M_\odot$. Our roughly equivalent low-mass major pair sample, when adopting a separation cut of $r_{\text{sep}} < 150$ kpc, results in a major pair fraction of $0.022^{+0.002}_{-0.003}$ at $z = 0$. We believe our pair fractions are somewhat higher than in Besla et al. (2018) as our stellar mass range for primaries extends to a lower value of $10^8 M_\odot$. As such, our results for low-mass pairs are robust.

For high-mass pairs, Snyder et al. (2017) created mock catalogs using light cones in the Illustris-1 simulation to select pairs in the same fashion as done in observations. They consider major pairs (stellar mass ratio $> 1/4$) with $1 \times 10^{11} > M_{*1} > 1 \times 10^{10.5} M_\odot$, projected distances between 14 and 71 kpc, and a redshift separation of $\Delta z < 0.02(1 + z_{\text{pri}})$, corresponding to a velocity separation of $< 1.8 \times 10^4 \text{ km s}^{-1}$ at $z = 2$. They find that the major pair fraction is constant or decreasing for $z > 1$, which is in good agreement with observational studies.

Snyder et al. (2023) extended this work, utilizing the TNG simulation to create mock images of extragalactic survey fields mimicking future planned surveys like JADES. From the mock images, they calculated the pair fraction for major pairs with projected separations between 5 and 70 kpc and with redshift separations of $\Delta z < 0.02(1 + z)$. Again, they found a flat or

²¹ NGC 4490–4485 is a low-mass ($M_* = 7.2$ and $0.82 \times 10^9 M_\odot$, respectively) galaxy binary that is ~ 7 Mpc away from the MW (Theureau et al. 2007; Pearson et al. 2018).

²² Figure 5 displays pair fraction trends for major pairs. The equivalent plot for minor pairs, which we did not include here, shows the same trends as the major pairs, particularly in the < 50 kpc panel where the pair fraction for high-mass minor pairs is < 0.02 at $z = 0$.

decreasing major pair fraction with increasing redshift above $z = 1$.

We cannot directly compare the values of the pair fractions in this work and Snyder et al. (2023), since their work is done in projected space, but we can compare the trends as a function of redshift. We find that the high-mass major pair sample has a pair fraction that decreases with redshift above $z = 1$, in agreement with Snyder et al. (2023).

In addition, Rodriguez-Gomez et al. (2015) examine the merger rates of galaxies in the Illustris simulations as a function of stellar mass and redshift. From their Figures 7 and 10, there is roughly a factor of 4 difference in the major merger rate for high-mass ($M_* \sim 10^{11} M_\odot$) versus low-mass ($M_* \sim 10^8 M_\odot$) galaxies at $z = 0$, and the difference becomes smaller at higher redshift. A strong mass dependence of the galaxy merger rate at $z = 0$ is also seen in Figure 12 of Guzmán-Ortega et al. (2023), as well as in results from semi-empirical models (Stewart et al. 2009; Hopkins et al. 2010). From our Figure 3, there is also a factor of 4 difference in the pair fraction for high- and low-mass galaxies at $z = 0$, which becomes smaller at higher redshift. The consistency between our results and the variation with stellar mass and redshift of the galaxy merger rate suggests that our finding of different pair fractions between low- and high-mass galaxies at $z = 0$, as well as their different redshift evolution trends, is indeed reliable.

5.3. Implications for Observational Pair Fraction Studies

In Section 4.1, we presented the pair fraction of high-mass and low-mass galaxy pairs and found that the relative frequency of the two populations evolves distinctly from $z = 0$ to 4.2. High-mass pair fractions peak at $z = 0$, decrease from $z = 0$ to 0.3, after which they remain roughly constant or mildly decrease with increasing redshift. In contrast, low-mass pair fractions increase with increasing redshift until $z = 2.5$, after which the frequency remains roughly constant with increasing redshift. This behavior is seen for both major and minor pairs.

In Section 4.2, we characterized the behavior of pair fractions for subsamples of the full pair catalog that pass additional separation cuts. We found that physical separation cuts which do not vary as a function of time and mass eliminate the ability to distinguish the different redshift evolution of high- and low-mass pair fractions. Instead, by adopting a separation cut based on the virial radius of the group halo, we can accurately recover the different redshift evolution of low-mass versus high-mass pair fractions, particularly at $z < 2.5$.

These results indicate that future observational studies that seek to compare low-mass and high-mass pair fractions, particularly as a function of time, must take care when determining their pair selection criteria. We advocate for separation criteria that varies as both a function of mass and redshift, such as our choice of $r_{\text{sep}} < 1.0 R_{\text{vir}}$, where R_{vir} can be inferred using an estimate of the combined dark matter mass of the pair.²³ We have shown that utilizing fixed physical separation cuts can lead to significant

²³ Taking the observed stellar mass of each pair member, the halo mass for each pair can be estimated using the SMHM relation (e.g. Moster et al. 2013). The combined halo mass of the pair can then be used as a proxy for the virial mass of the FoF group (M_{vir}) to compute the virial radius, e.g.,

$$R_{\text{vir}} = \sqrt[3]{\frac{3M_{\text{vir}}}{4\pi\Delta_c\rho_c}}, \quad (1)$$

where ρ_c is the critical density of the Universe, and Δ_c is the overdensity constant (see Binney & Tremaine 2008). Both ρ_c and Δ_c are functions of redshift, such that R_{vir} changes as a function of both the mass of the pair and redshift.

deviation in the behavior of galaxy pair fractions for different galaxy masses and between $z = 0$ and 4, and thus it is imperative to carefully consider the selection criteria used in future observational pair fraction studies of low- and high-mass pairs over time.

In addition, our findings are specific to systems in isolated environments, and thus may not be representative of the pair fractions of a more “standard” observational field that contains pairs in isolated and high-density environments. Mitigating this issue in observations would require making certain isolation cuts, such as those employed in Geha et al. (2013) to identify low-mass pairs in isolated (> 1.5 Mpc away from an L^* galaxy) and nonisolated (within 1.5 Mpc of an L^* galaxy) environments.

Given that chance projections will artificially boost the pair fraction as the pair separation is increased, recovering the general evolution of the pair fraction as a function of redshift may be more realistic than recovering the magnitude of the pair fraction itself, since the former requires smaller separation limits. The requirement of $r_{\text{sep}} < 1.0 R_{\text{vir}}$ for low-mass pairs translates to median separation limits that are less than ~ 60 (30) kpc at $z \sim 2(4)$, which are sufficiently small as to avoid major projection effects. For high-mass pairs, this translates to separations less than $\sim 135(75)$ at $z \sim 2(4)$.

Pair fractions are typically translated to galaxy merger rates using an observability window (Lotz et al. 2011). The good agreement between the difference in the low-mass and high-mass pair fractions and the mass-dependent galaxy merger rates from Rodriguez-Gomez et al. (2015) suggests that applying separation cuts which can recover this mass dependency are critical to reliably translating pair fractions to galaxy merger rates. In fact, JWST is expected to identify low-mass galaxies ($M_* > 10^8 M_\odot$) out to at least $z = 10$ (Cowley et al. 2018; Williams et al. 2018; Behroozi et al. 2020). We argue that the adopted pair separation criteria, particularly for low-mass pairs, will greatly affect the measured pair fractions and consequently the inferred merger rates.

5.4. Underlying Physical Behavior of Galaxy Pairs

Though we find a distinct difference in the redshift evolution of low-mass and high-mass pair fractions, our study does not explicitly identify the specific driver of these differences. In a Λ CDM universe, the buildup of structure proceeds hierarchically, with smaller structures merging to form larger and larger structures at later times. This may explain the significant decrease in the pair fraction of isolated low-mass galaxies from $z = 2$ to $z = 0$, since more and more low-mass systems would be accreted into larger group structures and would thus be removed from our isolated pair sample.

On the other hand, the inherent dynamics of the pairs themselves may dictate the redshift evolution of their respective pair fractions. For example, if low-mass pairs tend to have very short merger timescales, the total number of low-mass pairs would decrease much more rapidly than massive pairs with a very long timescale.

Understanding in further detail the origin of the discrepancy in the redshift evolution of low-mass versus high-mass pairs requires tracking the specific pairs across simulation snapshots to trace their orbits and study changes in the mass and environment of the pairs across cosmic time. We leave this as the focus of future work.

6. Summary and Conclusions

In this paper, we construct a sample of low-mass ($10^8 < M_* < 5 \times 10^9 M_\odot$) and high-mass ($5 \times 10^9 < M_* < 10^{11} M_\odot$) pairs from the TNG100 simulation from $z = 0$ to 4.2. Pairs are selected as belonging to the same FoF group such that they are isolated and physically associated, and not contaminated by projection effects. Major and minor pairs are determined using the stellar mass ratio (1–1/4 and 1/4–1/10, respectively) between the primary and secondary halo, where stellar masses are assigned using an abundance matching prescription from Moster et al. (2013) to generate 1000 separate realizations of selected pairs at each redshift in order to sample the error space of the SMHM relation.

From this pair sample, we calculate the pair fraction as a function of redshift for four different pair types: low-mass major pairs, low-mass minor pairs, high-mass major pairs, and high-mass minor pairs. Our goal is to quantify the evolution of low-mass and high-mass galaxy pairs across cosmic time and to identify any differences in the redshift evolution of pair fractions of low-mass and high-mass pairs. We also aim to determine how separation criteria for pair selection can affect the pair fraction evolution of high- and low-mass pairs. To this end, we additionally employ both static and time- and mass-evolving separation cuts to explore the resulting impact on the pair fractions.

Our main findings are as follows:

1. The pair fraction for low-mass and high-mass pairs does not proceed identically throughout cosmic time. In fact, the two mass scales have opposite behaviors at $z < 2.5$ (see Figure 3). The pair fraction of low-mass pairs increases from $z = 0$ to $z \sim 2.5$, while the pair fraction of high-mass pairs peaks at $z = 0$ and is constant or slightly decreasing at $z > 1$.
2. At $z = 0$, we find a low-mass major (minor) pair fraction of $0.041^{+0.003}_{-0.004}$ (0.034 ± 0.004), and a high-mass major (minor) pair fraction of $0.207^{+0.012}_{-0.014}$ ($0.152^{+0.016}_{-0.017}$). These results are consistent with other simulation-based studies that find that the merger rate for high-mass pairs is $\sim 4\times$ higher than for low-mass galaxies (e.g., Rodriguez-Gomez et al. 2015).
3. Low-mass minor pairs evolve similarly to low-mass major pairs as a function of redshift.
4. The differences in redshift evolution of pair fractions for high- and low-mass pairs is well recovered using a subsample of pairs with a maximum separation cut that varies with mass and redshift, specifically for $r_{\text{sep}} < 1R_{\text{vir}}$, where R_{vir} is the virial radius of the FoF group. This separation cut corresponds to separations less than [134 (348), 59(134), 33(76)] kpc at $z = [0, 2, 4]$ for low-mass (high-mass) pairs. In particular, a maximum separation cut of $r_{\text{sep}} < 1R_{\text{vir}}$ encompasses between $\sim 40\%$ and 60% of the full population of pairs at all redshifts, and recovers the distinct differences in the redshift evolution of low-mass versus high-mass pair fractions from $z = 0$ to 2.5.
5. The identified differences in the redshift evolution of high- and low-mass pairs are erased when a static physical separation cut is employed, i.e., a separation that does not evolve with redshift or with mass (e.g., $r_{\text{sep}} < 50$ kpc). This occurs because fixed-separation cuts capture roughly 10%–50% more pairs in the low-mass pair sample than the high-mass sample at all redshifts.

Consequently, selecting the same volume to study the pair fractions for low- and high-mass galaxies via a static separation cut will bias inferred pair fractions.

6. Isolated mass analogs of the MW–M31 are most common at $z = 0$, while analogs of the MW–LMC and M31–M33 are equally common at $z = 0$ as at $z > 1$. However, MW–LMC-type systems with very low separations ($\lesssim 50$ kpc) are 2–3 times more common at higher redshift ($z \sim 2$). Isolated analogs of the LMC–SMC and NGC 4490–4485 are 3 times more common for $z > 3$.

A number of studies have identified a mass and redshift dependence in the galaxy merger rate as a function of time, particularly in cosmological simulations and semi-empirical models (see, e.g., Stewart et al. 2009; Hopkins et al. 2010; Rodriguez-Gomez et al. 2015). Our results show that galaxy pair fractions for physically associated, isolated pairs likewise evolve as a function of mass and redshift. However, we also find that the recovered pair fractions are sensitive to the separation criteria that is used to define pairs. The redshift evolution of pair fractions for physically associated low-mass and high-mass galaxy pairs can only be distinguished when the separation criteria are a function of both mass and redshift.

Observational campaigns that seek to recover low-mass and high-mass pair fractions should consider using mass- and redshift-evolving separation criteria to select pairs. Many observational pair fraction studies use static definitions of the maximum projected separation to determine close pairs, then translate between pair fractions and merger rates via an observability timescale. The observability timescale, however, is not typically calibrated as a function of both mass and redshift, meaning that identical measured pair fractions of high- and low-mass pairs will yield indistinguishable merger rates.

In order to observationally recover the cosmologically expected difference between the merger rates of high- and low-mass galaxies, it is crucial to ensure that differences in the redshift evolution of the pair fraction for these galaxy populations can be adequately captured. This will be of particular importance in the era of JWST, where more low-mass ($M_* \sim 10^8 M_\odot$) galaxies and galaxy pairs will be discovered at higher redshift (up to $z > 10$), and for future wide-field galaxy surveys enabled by the Roman Space Telescope and Rubin Observatory.

Acknowledgments

K.C. would like to thank Greg Snyder for fruitful early conversations during this work. K.C. and G.B. are supported by NSF CAREER award AST-1941096. E.P. acknowledges financial support provided by NASA through grant No. HST-GO-16628. Support for this work was also provided by NASA through the NASA Hubble Fellowship grant No. HST-HF2-51540.001-A, awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA contract NAS5-26555. Support for S.P. was provided by NASA through the NASA Hubble Fellowship grant No. HST-HF2-51466.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA contract NAS5-26555. The work of S.S. is supported by the National Science Foundation Astronomy & Astrophysics Research grant No. 1911107.













The ILLUSTRITNG simulations were undertaken with compute time awarded by the Gauss Centre for Supercomputing (GCS) under GCS Large-Scale Projects GCS-ILLU and GCS-DWAR on the GCS share of the supercomputer Hazel Hen at the High Performance Computing Center Stuttgart (HLRS), as well as on the machines of the Max Planck Computing and Data Facility (MPCDF) in Garching, Germany.

The National Radio Astronomy Observatory is a facility of the National Science Foundation operated under cooperative agreement by Associated Universities, Inc.

This work is based upon High Performance Computing (HPC) resources supported by the University of Arizona TRIF, UITS, and Research, Innovation, and Impact (RII) and maintained by the UArizona Research Technologies department.

We respectfully acknowledge the University of Arizona is on the land and territories of Indigenous peoples. Today, Arizona is home to 22 federally recognized tribes, with Tucson being home to the Oodham and the Yaqui. Committed to diversity and inclusion, the University strives to build sustainable relationships with sovereign Native Nations and Indigenous communities through education offerings, partnerships, and community service.

ORCID iDs

Katie Chamberlain  <https://orcid.org/0000-0001-8765-8670>
 Gurtina Besla  <https://orcid.org/0000-0003-0715-2173>
 Ekta Patel  <https://orcid.org/0000-0002-9820-1219>
 Vicente Rodriguez-Gomez  <https://orcid.org/0000-0002-9495-0079>
 Paul Torrey  <https://orcid.org/0000-0002-5653-0786>
 Garreth Martin  <https://orcid.org/0000-0003-2939-8668>
 Kelsey Johnson  <https://orcid.org/0000-0001-8348-2671>
 Nitya Kallivayalil  <https://orcid.org/0000-0002-3204-1742>
 David Patton  <https://orcid.org/0000-0002-1871-4154>
 Sarah Pearson  <https://orcid.org/0000-0003-0256-5446>
 George Privon  <https://orcid.org/0000-0003-3474-1125>
 Sabrina Stierwalt  <https://orcid.org/0000-0002-2596-8531>

References

Akins, H. B., Christensen, C. R., Brooks, A. M., et al. 2021, *ApJ*, 909, 139
 Behroozi, P., Conroy, C., Wechsler, R. H., et al. 2020, *MNRAS*, 499, 5702
 Besla, G., Kallivayalil, N., Hernquist, L., et al. 2007, *ApJ*, 668, 949
 Besla, G., Patton, D. R., Stierwalt, S., et al. 2018, *MNRAS*, 480, 3376
 Bignone, L. A., Tissera, P. B., Sillero, E., et al. 2017, *MNRAS*, 465, 1106
 Binney, J., & Tremaine, S. 2008, *Galactic Dynamics* (2nd ed.; Princeton, NJ: Princeton Univ. Press)
 Blecha, L., Snyder, G. F., Satyapal, S., & Ellison, S. L. 2018, *MNRAS*, 478, 3056
 Boylan-Kolchin, M., Besla, G., & Hernquist, L. 2011, *MNRAS*, 414, 1560
 Bryan, G. L., & Norman, M. L. 1998, *ApJ*, 495, 80
 Casteels, K. R. V., Conselice, C. J., Bamford, S. P., et al. 2014, *MNRAS*, 445, 1157
 Chamberlain, K., Price-Whelan, A. M., Besla, G., et al. 2023, *ApJ*, 942, 18
 Chua, K. T. E., Pillepich, A., Rodriguez-Gomez, V., et al. 2017, *MNRAS*, 472, 4343
 Comerford, J. M., Pooley, D., Barrows, R. S., et al. 2015, *ApJ*, 806, 219
 Conselice, C. J. 2003, *ApJS*, 147, 1
 Cowley, W. I., Baugh, C. M., Cole, S., Frenk, C. S., & Lacey, C. G. 2018, *MNRAS*, 474, 2352
 Davis, M., Efstathiou, G., Frenk, C. S., & White, S. D. M. 1985, *ApJ*, 292, 371
 Di Matteo, P., Bournaud, F., Martig, M., et al. 2008, *A&A*, 492, 31
 Dolag, K., Borgani, S., Murante, G., & Springel, V. 2009, *MNRAS*, 399, 497
 Duncan, K., Conselice, C. J., Mundy, C., et al. 2019, *ApJ*, 876, 110
 Ellison, S. L., Nair, P., Patton, D. R., et al. 2011, *MNRAS*, 416, 2182
 Ellison, S. L., Viswanathan, A., Patton, D. R., et al. 2019, *MNRAS*, 487, 2491
 Fattahi, A., Navarro, J. F., Starkenburg, E., Barber, C. R., & McConnachie, A. W. 2013, *MNRAS*, 431, L73

Ferreras, I., Trujillo, I., Marmol-Queralto, E., et al. 2014, *MNRAS*, 444, 906
 Gardner, J. P., Mather, J. C., Clampin, M., et al. 2006, *SSRv*, 123, 485
 Geha, M., Brown, T. M., Tumlinson, J., et al. 2013, *ApJ*, 771, 29
 Geha, M., Wechsler, R. H., Mao, Y.-Y., et al. 2017, *ApJ*, 847, 4
 Glikman, E., Simmons, B., Mailly, M., et al. 2015, *ApJ*, 806, 218
 Guo, Q., & White, S. D. M. 2008, *MNRAS*, 384, 2
 Guzman-Ortega, A., Rodriguez-Gomez, V., Snyder, G. F., Chamberlain, K., & Hernquist, L. 2023, *MNRAS*, 519, 4920
 Hani, M. H., Gosain, H., Ellison, S. L., Patton, D. R., & Torrey, P. 2020, *MNRAS*, 493, 3716
 Hopkins, P. F., Cox, T. J., Hernquist, L., et al. 2013, *MNRAS*, 430, 1901
 Hopkins, P. F., Croton, D., Bundy, K., et al. 2010, *ApJ*, 724, 915
 Hopkins, P. F., Hernquist, L., Cox, T. J., & Keres, D. 2008, *ApJS*, 175, 356
 Jackson, R. A., Kaviraj, S., Martin, G., et al. 2022, *MNRAS*, 511, 607
 Jackson, R. A., Martin, G., Kaviraj, S., et al. 2019, *MNRAS*, 489, 4679
 Ivezic, Z., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, 873, 111
 Lin, L., Koo, D. C., Willmer, C. N. A., et al. 2004, *ApJL*, 617, L9
 Lin, L., Patton, D. R., Koo, D. C., et al. 2008, *ApJ*, 681, 232
 Lotz, J. M., Davis, M., Faber, S. M., et al. 2008, *ApJ*, 672, 177
 Lotz, J. M., Jonsson, P., Cox, T. J., et al. 2011, *ApJ*, 742, 103
 Lubert, N., Pearson, S., Putman, M. E., et al. 2022, *AJ*, 163, 49
 Man, A. W. S., Zirm, A. W., & Toft, S. 2016, *ApJ*, 830, 89
 Mao, Y.-Y., Geha, M., Wechsler, R. H., et al. 2021, *ApJ*, 907, 85
 Marinacci, F., Vogelsberger, M., Pakmor, R., et al. 2018, *MNRAS*, 480, 5113
 Martin, G., Bazkiaei, A. E., Spavone, M., et al. 2022, *MNRAS*, 513, 1459
 Martin, G., Jackson, R. A., Kaviraj, S., et al. 2021, *MNRAS*, 500, 4937
 Martin, G., Kaviraj, S., Devriendt, J. E. G., Dubois, Y., & Pichon, C. 2018, *MNRAS*, 480, 2266
 Martin, G., Kaviraj, S., Devriendt, J. E. G., et al. 2017, *MNRAS*, 472, L50
 Mihos, J. C., & Hernquist, L. 1996, *ApJ*, 464, 641
 Moster, B. P., Naab, T., & White, S. D. M. 2013, *MNRAS*, 428, 3121
 Munshi, F., Brooks, A. M., Applebaum, E., et al. 2021, *ApJ*, 923, 35
 Naiman, J. P., Pillepich, A., Springel, V., et al. 2018, *MNRAS*, 477, 1206
 Nelson, D., Pillepich, A., Springel, V., et al. 2018, *MNRAS*, 475, 624
 Nelson, D., Springel, V., Pillepich, A., et al. 2019, *ComAC*, 6, 2
 Patel, E., Besla, G., & Sohn, S. T. 2017a, *MNRAS*, 464, 3825
 Patel, E., Besla, G., & Mandel, K. 2017b, *MNRAS*, 468, 3428
 Patton, D. R., Pritchett, C. J., Carlberg, R. G., et al. 2002, *ApJ*, 565, 208
 Patton, D. R., Qamar, F. D., Ellison, S. L., et al. 2016, *MNRAS*, 461, 2589
 Patton, D. R., Torrey, P., Ellison, S. L., Mendel, J. T., & Scudder, J. M. 2013, *MNRAS*, 433, L59
 Patton, D. R., Wilson, K. D., Metrow, C. J., et al. 2020, *MNRAS*, 494, 4969
 Paudel, S., Smith, R., Yoon, S. J., Calderon-Castillo, P., & Duc, P.-A. 2018, *ApJS*, 237, 36
 Pearson, S., Besla, G., Putman, M. E., et al. 2016, *MNRAS*, 459, 1827
 Pearson, S., Privon, G. C., Besla, G., et al. 2018, *MNRAS*, 480, 3069
 Pillepich, A., Nelson, D., Hernquist, L., et al. 2018, *MNRAS*, 475, 648
 Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016, *A&A*, 594, A13
 Ramos Almeida, C., Tadhunter, C. N., Inskip, K. J., et al. 2011, *MNRAS*, 410, 1550
 Robertson, B., Dickinson, M., Ferguson, H. C., et al. 2019a, *BAAS*, 51, 30
 Robertson, B. E., Banerji, M., Brough, S., et al. 2019b, *NatRP*, 1, 450
 Rodriguez-Gomez, V., Genel, S., Vogelsberger, M., et al. 2015, *MNRAS*, 449, 49
 Rodriguez-Gomez, V., Sales, L. V., Genel, S., et al. 2017, *MNRAS*, 467, 3083
 Sales, L. V., Wang, W., White, S. D. M., & Navarro, J. F. 2013, *MNRAS*, 428, 573
 Satyapal, S., Ellison, S. L., McAlpine, W., et al. 2014, *MNRAS*, 441, 1297
 Snyder, G. F., Lotz, J. M., Rodriguez-Gomez, V., et al. 2017, *MNRAS*, 468, 207
 Snyder, G. F., Pena, T., Yung, L. Y. A., et al. 2023, *MNRAS*, 518, 6318
 Snyder, G. F., Rodriguez-Gomez, V., Lotz, J. M., et al. 2019, *MNRAS*, 486, 3702
 Spergel, D., Gehrels, N., Baltay, C., et al. 2015, arXiv:1503.03757
 Springel, V., Pakmor, R., Pillepich, A., et al. 2018, *MNRAS*, 475, 676
 Springel, V., White, S. D. M., Tormen, G., & Kauffmann, G. 2001, *MNRAS*, 328, 726
 Stewart, K. R., Bullock, J. S., Barton, E. J., & Wechsler, R. H. 2009, *ApJ*, 702, 1005
 Stierwalt, S., Besla, G., Patton, D., et al. 2015, *ApJ*, 805, 2
 Theureau, G., Hanski, M. O., Coudreau, N., Hallet, N., & Martin, J. M. 2007, *A&A*, 465, 71
 Treister, E., Natarajan, P., Sanders, D. B., et al. 2010, *Sci*, 328, 600
 Wang, L., Pearson, W. J., & Rodriguez-Gomez, V. 2020, *A&A*, 644, A87
 Williams, C. C., Curtis-Lake, E., Hainline, K. N., et al. 2018, *ApJS*, 236, 33
 Wuyts, S., Forster Schreiber, N. M., Nelson, E. J., et al. 2013, *ApJ*, 779, 135