UNIVERSITY OF SURREY

MPHYS DISSERTATION

Understanding the Relative Velocity Peak of Isolated Galaxy Pairs

---

Author:
Emily ST PIER

Supervisor:
Marcel PAWLOWSKI

A thesis submitted in fulfillment of the requirements for the degree of Master of Physics

This thesis and the work to which it refers are the results of my own efforts. Any ideas, data, images or text resulting from the work of others (whether published or unpublished) are fully identified as such within the work and attributed to their originator in the text, bibliography or in footnotes. This thesis has not been submitted in whole or in part for any other academic degree or professional qualification. I agree that the University has the right to submit my work to the plagiarism detection service TurnitinUK for originality checks. Whether or not drafts have been so-assessed, the University reserves the right to require an electronic version of the final document (as submitted) for assessment as above.

2024

UNIVERSITY OF SURREY

Abstract

Faculty of Engineering and Physical Sciences

School of Mathematics and Physics

Master of Physics

Understanding the Relative Velocity Peak of Isolated Galaxy Pairs

by Emily ST PIER

Following the discovery of a preferred relative velocity of isolated galaxy pairs, and its subsequent demonstration in two cosmological paradigms both as a prediction of MOdified Newtonian Dynamics (MOND) and a feature of cosmological constant cold dark matter (ΛCDM) simulations, this study investigated how the relative velocity peaks of pair catalogs at different magnitudes and redshifts vary in the Illustris TNG-300 ΛCDM simulation. Isolated galaxy pair catalogs with varying absolute magnitude limits were extracted from TNG-300, and their intervelocity peaks identified. These peaks have then been compared with those predicted by MOND given the distribution of absolute magnitude M_B of the pairs. It has been discovered that the relative velocity peaks in ΛCDM differ from those predicted by MOND in two key ways: they show less variance with magnitude, decreasing more slowly than those predicted by MOND, and they show far greater variation in their widths than those in MOND, the widths of which are approximately constant. The potential of these differences to be used as probes to differentiate between the two paradigms in the context of isolated galaxy pairs is then discussed.

Acknowledgements

Thank you to all the lovely folks at the Leibniz Institute for Astrophysics who I've had the chance to work with while producing this, and to my visiting tutor Dr M Collins.

Special thanks to Dr M Pawlowski for his invaluable advice and suggestions on this project and feedback on this manuscript.

Contents

For my lovely mum, and her tireless cheerleading. For my wonderful uncle, who made all this possible. And for my adoring partner, my rock and my reason.

Chapter 1

Introduction and Precis of Interim Dissertation

1.1 Introduction

Following the discovery of a peak in the relative velocities of galaxy pairs in 2020 by Nottale and Chamaraux[Nottale and Chamaraux, 2020], there has been some debate among physicists as to which model of gravity best explains its presence. Standard cosmology, ΛCDM or Cosmological Constant Cold Dark Matter, and MOND or Modified Newtonian Dynamics, are two models that have long been at odds, and in the case of isolated galaxy pairs there is no difference. Investigations of the intervelocity peak have, until now, been focused on the observational data - they have assessed its position, its significance, and searched for a theoretical explanation for its existence or some cosmological analog in simulation. They have not, however, explored in any detail the factors that may affect the existence and presence of the peak, nor its position.

This dissertation presents and interprets the results of an investigation into the nature of the intervelocity peak, beyond simply its existence. The study conducted was intended to shed some light onto the origin of the peak in ΛCDM, by identifying which variables and environments cause it to change in position or shape in cosmological simulations, with the hopes of making some predictions that may differentiate between the two models of MOND and ΛCDM.

1.2 Two Models of Gravity

It has long been apparent that there there is something missing from modern understanding of gravity.

The initial evidence for such a claim came from a study of the velocity dispersion of the Coma Cluster[Zwicky, 1933], wherein the velocity dispersion was found to be so high that by all understanding, the cluster should not have been bound and should have torn itself apart. Either the Coma Cluster was not bound, or there was some additional component to the gravitational forces within the cluster that held it together. Zwicky, to explain this, claimed there was some unseen matter within the cluster that increased the mass by such a degree that the cluster could, in fact, be bound, despite its high velocity dispersion. He named this substance dark matter, and from it came the cosmological model taught in most modern textbooks[Ryman, 2016]. The theory of dark matter was brought to the main scientific stage in the 1980s, when it was discovered that galactic rotation curves are flat[Einasto, Kaasik, and Saar, 1974; Ostriker, Peebles, and Yahil, 1974; Roberts, 1976; Rubin, Ford, and Thonnard, 1978] - that is, that the rotational velocities of stars around galaxies are constant to a large radius. To explain this with standard Newtonian gravity would imply a mass distribution far more extended than visible mass would allow - again, the need for dark matter arises.

However, there are some to whom the idea of a Universe dominated by unseen mass is not so attractive. In 1983, Mordehai Milgrom published a series of papers on a new theory of gravity, MOdified Newtonian Dynamics or MOND[Milgrom, 1983a; Milgrom, 1983b; Milgrom, 1983c]. These laws suggest that below a specific acceleration, $a_0 \approx 1 \times 10^{-8}$, gravitational behaviours change. In MOND, the acceleration a of a particle can be given by

$$\frac{a^2}{a_0} \approx \frac{GM}{r^2} \tag{1.1}$$

where $G = 6.67 \times 10^{-11}$ m^3 kg^{-2} s^{-2} is the gravitational constant, and M is the mass of the body being orbited. This modifications have had a range of successes in explaining physical phenomena much like ΛCDM has, most notably with reference to the baryonic Tully-Fisher relation, hereon BTFR[McGaugh, 2012] that describes a

tight and well-documented relationship between a galaxy's rotational velocity and the baryonic component of its mass. In ΛCDM, the dominant mass component of a galaxy is dark matter, but the BTFR suggests that its rotational dynamics are governed by the baryonic mass. While ΛCDM-based simulations have reproduced this, the fact remains that it is unintuitive in that framework, while something like MOND not only explains but predicts it.

1.3 Galaxy Pairs

1.3.1 An Isolated Galaxy Pair Catalog

The HyperLEDA database[Makarov et al., 2014] is a compilation of observational studies to create a more extensive catalog of galaxies than could ever be obtained using just one study. Using this, a catalog could be constructed of isolated galaxy pairs[Nottale and Chamaraux, 2018a] more extensive than any catalog produced before. This catalog, the Isolated Galaxy Pair Catalog or IGPC, was then analysed using Nottale and Chamaraux's newly developed deprojection algorithm[Nottale and Chamaraux, 2018b; Nottale and Chamaraux, 2020].

Galaxy pairs for this catalog were selected by the following criteria:

- Absolute magnitude M_B<-18.5.
- Radial velocity 3000<v_z km s^{-1}<16000
- Radial velocity difference ΔV<500 km s^{-1}
- Projected separation r_p<1Mpc
- Reciprocity - such that every galaxy is the closest to its neighbour, and its neighbour is closest to it.
- Isolation. An isolation parameter was defined as $\rho = \frac{r_{third}}{r_p}$, where r_{third} is the distance to the nearest neighbour. For inclusion in the catalog, a pair must have ρ>2.5.

To understand the relative motions of galaxy pairs, six things must be known: three dimensions of position, and three dimensions of velocity. However, an observer only knows their three projected values: two dimensions of position, those in

the plane of the sky; and one dimension of velocity, the radial or recessive velocity. Using Nottale and Chamaraux's deprojection algorithm, it became possible to create the distribution of 3D relative velocity - hereon intervelocity - and of 3D separation, using only the distributions of the known components. The only difficulty is that to do so requires a distribution. It cannot be done for a single pair, and there is no way to connect a deprojected intervelocity or separation to its specific galaxy - it requires a distribution as input, and produces only a distribution. Therefore, this sort of analysis can only be performed on large enough pair catalogs - hence the IGPC's importance.

To deproject intervelocity, the following algorithm is used:

$$P_v(v) = -v \left[\frac{dP_z(v_z)}{v_z} \right]_v \tag{1.2}$$

Where v is the 3D intervelocity and v_z is the radial or projected intervelocity. On applying this algorithm to the IGPC, Nottale and Chamaraux found an intervelocity peak at $\tilde{1}50$ km s^{-1}. Previously to this, no peak in the intervelocities of galaxy pairs had been decisively reported.

A previous work on the properties of galaxy pairs in the context of their environment[Moreno, Bluck, Ellison, et al., 2013] had found that galaxy pairs whose dynamics were dominated by their own binding energy did have an intervelocity peak between 100 and 200 km s$-$1. However, they also found that their isolated pairs had a far higher preferred intervelocity, more like 1000 km s^{-1}. In addition, isolated pairs were not the subject of their study, and their selection criteria differed enough from Nottale and Chamaraux that comparisons are not simple to make.

Following the publication of Nottale and Chamaraux's intervelocity result, further analysis of the catalog was carried out by Scarpa et al.[Scarpa, Falomo, and Aldo, 2022a]. Following an investigation of the statistical significance and false pair contamination of the peak, this study claimed the intervelocity peak's existence as a victory for MOND. They claimed that while Moreno's study had not found any conclusive results, a preferred intervelocity for galaxy pairs fell neatly out of the equations of MOND. In a follow-up paper[Scarpa, Falomo, and Aldo, 2022b], they applied the MONDian equation for the orbital velocities of two bodies in circular

motion [Milgrom, 1994] to the Isolated Galaxy Pair Catalog, and obtained an intervelocity peak of 141±5 km s^{-1}. This is, within the bounds of the uncertainty, the same as their result obtained by directly analysing the observational peak, which they found to be at 132±5 km s^{-1}.

The intervelocity of two bodies of different masses under circular motion in MOND is given by

$$\Delta V^4 = G a_0 m_{tot} B^2(m_1, m_2) \tag{1.3}$$

Where ΔV is the intervelocity of the pairs, G is the gravitational constant $G = 6.67 \times 10^{-}11$ m^3 kg^{-1} s, $a_0 = 1 \times 10^{-10}$ m s-2 is the MOND constant, m_tot is the sum of the masses of both pair members, and B is a function of the two masses such that

$$B = \frac{2}{3}\left(\frac{1 - x_1^{3/2} - x_2^{3/2}}{x_1 x_2}\right) \tag{1.4}$$

where x_i is the mass proportion of each member, such that $x_i = \frac{m_i}{m_{tot}}$.

This was followed in 2020 by a study on the cosmological simulation Illustris TNG-300[Nelson, Springel, Pillepich, et al., 2019; Naiman, Pillepich, Springel, et al., 2018a; Springel, Pakmor, Pillepich, et al., 2018; Nelson, Pillepich, Springel, et al., 2018; Naiman, Pillepich, Springel, et al., 2018b; Marinacci, Vogelsberger, Pakmor, et al., 2018] where isolated galaxy pairs were selected by the same criteria as for the IGPC and mock-observed to identify the intervelocity peak[Pawlowski et al., 2022].

An 'observer' was positioned on one of the faces of the simulation box, and this face determined which direction was radial velocity. Apparent magnitudes were calculated, and a cut-off applied at m_B=19 to mimic the observation. They also mimicked errors in the observation by randomly displacing each radial velocity by an error drawn randomly from the HyperLEDA database. Pairs were then selected according to Nottale and Chamaraux's selection criteria as listed above, and the resulting intervelocity distribution was deprojected to complete the mock observation.

The most important result from Pawlowski et al. was identification of an intervelocity peak at 131±1 km s^{-1} without deprojection, and 125±7 km s^{-1} with.

This proved that the intervelocity peak was present in ΛCDM simulations. In addition, that these two results - one obtained using mock observation and one without - were the same within the bounds of the uncertainty. This demonstrated that the mock observation was not necessary to obtain the intervelocity peak; however, that a far larger error would be obtained using mock observation data than true 3D data direct from the simulation.

Chapter 2

Methodology

2.1 Following on from Pawlowski et al.

2.1.1 The Simulation

Data for this project was taken from the Illustris TNG-300 simulation for a number of reasons. The TNG suite consists of a number of frequently used full-physics cosmological simulations, and the 300Mpc box size of the TNG-300 run allows for a useful balance between sample size and resolution. Smaller runs such as TNG-100 or TNG-50 may have allowed for a closer study of even fainter galaxies; however, the limitations on physical size would have also led to limited sample sizes, particularly in the brighter magnitudes. For the purposes of making observational predictions, they would then be rather impractical.

One of the earliest preliminary studies to be carried out was locating the absolute magnitude limit at which the resolution of the simulation would begin to result in 'missing' galaxies. Investigation into the magnitude distribution of galaxies from TNG-300 between -9.5<M_B<-25 shows that the distribution peaks at approximately MB=-11 for the z=0 snapshot - more on this in Results. This suggests that galaxies fainter than this are often poorly resolved, and may not be detected by TNG-300's halo finder; this will then lead to an incomplete catalog below this point. In contrast, the observational study imposed an absolute magnitude limit of -18.5, substantially higher than this point. The resolution on TNG-300 is therefore more than sufficient for the needs of this investigation, while maximising sample size.

2.1.2 Mock Observation

The simulated galaxies were retrieved from the friends of friends group catalogs, and are a complete account of all galaxies produced in the simulation. When considering predictions that may be made for observation, this is rather useless - many galaxies will not be able to be seen. Therefore, much as in Pawlowski et al.'s code, the reproduced code for this project mock-observes the galaxies present.

We positioned ourselves on one face of the simulation box, and measured distances to galaxies from that face. It is demonstrated by Pawlowski et al. that observations from the x, y, and z faces are, allowing for random fluctuations, indistinguishable from one another. From this, the recessional velocities were then calculated using the distances of each galaxy from the observer, along with the TNG-300 Hubble constant of H=67.76 km s^{-1} Mpc^{-1}.

The mock observational approach also imposed limits on position - on the distance to the galaxies, specifically. The region where recessional velocity falls into the required range, 3000<v_z km s^{-1}<16000 takes up approximately 60% of the simulation box, which left a lot of the data unused. However, despite this missing data, the position constraints were kept for two reasons. The first, on a more practical point, was to save on computation time. There were huge numbers of pairs, particularly in the fainter catalog - the pair finder often ran for over 24 hours at any given time to produce a single catalog. Keeping this distance constraint helped to keep this in a somewhat manageable region. The most important reason, however, was to ensure that the findings can be easily translated into potential predictions. By keeping as many constraints identical to the observational studies as possible, the data produced can still be treated as a mock observation with minimal adjustments - and this makes any predictions more practical for comparison to observation.

Pawlowski et al.'s mock observation conditions also impose an apparent magnitude limit. However, this condition has been removed from this study in order to thoroughly investigate the sample.

The final step of the mock-observation in the original study was deprojection of the intervelocity distribution, and this has been carried out in an identical way in

this study. The results of this deprojection, and how it may vary against the intervelocities read directly from the simulation, will be discussed in the Results section.

2.2 Selection Criteria

The Isolated Galaxy Pair Catalog required that a galaxy pair meet six criteria to be included. These are discussed in more depth in chapter 1. The criteria that relate to observability of a galaxy - the limit on redshift velocity and the limit on apparent magnitude, implicit in the observational study - are discussed in the Mock Observation section about. The selection criteria to determine whether or not a pair are interacting have been held constant in this study. To remind the reader, these are:

- Projected separation r_{sep}< 1 Mpc. This is retained in this study, rather than using a true 3D value, to again ensure compatibility of predictions with observational data.
- Difference in radial velocity ΔV< 500 km s^{-1}.
- Reciprocity
- Isolation parameter ρ>2.5 where ρ=r_{third}/r_p

The criteria on absolute magnitude has also been relaxed in this study, due to the simple fact that extending understanding of how the peak behaves with varying absolute magnitudes is one of the primary goals.

2.2.1 Application of the Selection Criteria - Single Galaxies

The pair selection algorithm was developed as follows:

First, the galaxies that met the criteria for selection were extracted from the friends of friends group catalog of TNG-300.

The mock observation strategy provided the line-of-sight distance to each galaxy. These were then used in tandem with the Hubble constant H_0=67.74 km s^{-1} Mpc^{-1} to produce recessional velocities, and galaxies with recessional velocities 2500<v km s^{-1}<16500 were selected. This range was chosen due to a combination of the redshift

range allowed by the observational catalog, 3000<v km s^{-1}<16000, and the pair selection criteria of $\Delta V < 500$ km s^{-1}. The intention of this was that all pairs meeting the criteria with at least one member in the desired distance range would be included - though this will be discussed again later. The absolute magnitude limit was also imposed at this stage.

Two catalogs were produced for each magnitude range. The first was the catalog for potential members. These catalogs spanned three magnitudes, to match the breadth of the observational study. The second was the extended catalog, that spanned an additional 2.5 magnitudes fainter than the standard. This was then used to ensure the pairs were isolated - details of this will be described with pair selection.

2.2.2 Application of Selection Criteria - Galaxy Pairs

Potential pairs were then identified by proximity. Potential primaries were selected by position, those that would have a recessional velocity 3000<v km s^{-1}<16000 and an absolute magnitude greater than the lower limit of the range (though this second requirement is implicit by catalog membership). Potential pairs were then selected by proximity, with only the closest galaxy to the primary being recorded as the potential partner. In this way, each galaxy can only be part of one pair - this is good for the study of isolated galaxy pairs, but difficult to extend to less isolated pairs.

Then, a reciprocity check was carried out. For each pair, member A must be the closest galaxy to member B, but member B must also be the closest to member A. This excludes multiplets, ensuring that the pairs are indeed pairs and not part of larger systems. This is the most significant difference between these selection criteria and those of Moreno et al [Moreno, Bluck, Ellison, et al., 2013], whose study focusing on the properties of galaxy pairs in cosmological context only selected for interacting galaxies, rather than specifically searching for pairs.

Finally, the isolation cuts were performed. The first cut searched for nearby galaxies that fell within the magnitude range of the pair or were brighter. If one was found within a distance equal to 10 times the pair separation, a value for isolation parameter ρ was recorded, alongside properties of the neighbouring galaxy - hereon the pair's third - such as its absolute magnitude M_B and its mass. For those where

no such galaxy could be located, ρ was simply recorded as 10, and no information on a third was logged. All galaxies with ρ>2.5 were retained.

The second isolation cut investigated galaxies fainter than the catalog. The lower limit on M_B for this second search was 2.5 magnitudes fainter than the faintest galaxy in the pair. The search then progressed much the same as the first, with all pairs with ρ>2.5 retained as catalog members.

2.3 Analysis Techniques

A series of investigations of these catalogs has been carried out, involving:

- How the intervelocity peak evolves with the magnitude range of the catalog.
- How the intervelocity peak changes shape and position with varying values of isolation parameter ρ.
- The effect of deprojection - whether each intervelocity peak can be recovered following deprojection of the recessional velocities, and quantifying by how much the peak shifts on from the true 3D intervelocity. This was applied in the same way as Pawlowski et al. - see Equation <X> in the Precis.
- What the MOND predictions are of the intervelocity for these catalogs, and how this compares with the intervelocity peaks extracted from TNG-300.
- How the intervelocity peak changes and evolves with redshift, and how this evolution compares with that predicted by MOND.

For each investigation, catalogs were produced, then assessed for their suitability for analysis. A number of the methods of analysis - most notably the deprojection - rely on statistics and distributions, so a significant number of galaxies was necessary to draw a reliable conclusion. For this study, this minimum pair count was set at N=1000.

For all intervelocity peaks produced, a double Gaussian was fit to the intervelocity distribution, and the covariance matrix produced alongside this curve fitting was used to estimate the uncertainties in the position of each peak. The only exception

to this were the MOND prediction peaks, that were best fitted with Gaussians - this will be discussed in more detail in the Results section.

2.3.1 Varying Isolation Parameter ρ

Following Pawlowski et al., it is ambiguous as to whether isolation has a specific impact on the peak position - this is due to vastly different trends in the observational catalog and the simulated catalog. For the observational catalog, the intervelocity after the standard ρ>2.5 cut was found to be significantly different to the intervelocity following cuts such that ρ>5 and ρ>10, but the latter two results were found to be the same to within their uncertainties. For the simulated data, there was no overlap in the uncertainties for the ρ>5 and ρ>10 sample, though the ρ>2.5 sample had a larger uncertainty and thus demonstrated overlap with both. One aim of this project was to clarify this relationship, and to conclude whether isolation affects the intervelocity peak, and if so, how.

2.3.2 Making Predictions for MOND

The predictions for the intervelocities using the equations of MOND were made using Equation 1.3, as quoted in the Precis.

The absolute magnitude distributions of the pairs were used to obtain a luminosity distribution. This was then used with a mass-to-luminosity ratio of 1.35, following the best fit of the observational catalog deduced by Scarpa et al. After applying this to obtain the distribution of masses, Equation 1.3 could be applied and the intervelocity distribution calculated.

2.3.3 Varying the Redshift

In order to vary the redshift of the sample, three different snapshots of TNG-300 are used. Snapshot 99, or redshift z=0, is discussed in the most depth in this dissertation. However, snapshots 50 and 33 (relating to redshifts z=1 and z=2 respectively) are also investigated. The full investigation - finding the intervelocities of different magnitude ranges, observing how these intervelocities vary with isolation, and

comparing these results to the MOND prediction - is then repeated for these higher redshifts.

2.3.4 Deprojection

To complete the analysis, each pair catalog had its radial intervelocities deprojected in order to investigate just how visible the patterns observed in this investigation might be in a real observation. This was carried out using Equation 1.2, and the need for large volumes of data in this step was the main reason for the minimum pair number for a given catalog being set at N=1000.

Chapter 3

Results

3.1 General Observations of the Galaxy Population

The distributions of masses and magnitudes of the galaxy catalog for redshift z=0 are shown in Figure 3.1. The mass distribution shows a fairly symmetrical peak about $10^{11}/M_{\odot}$, with masses typically ranging from $10^{8} < M/M_{\odot} < 10^{14}$. The important feature, however, is the decreasing number of galaxies at lower masses. This is not in agreement with expectation - the Schechter luminosity function[Schechter, 1976] predicts that the number of galaxies with a given mass should increase as mass decreases.

Similarly, the magnitude distribution shows a sharp peak followed by a decline for magnitudes. An interesting difference with the magnitude distribution, however, is its comparatively flat section between M_B=-20 and M_B=-15, something that is not

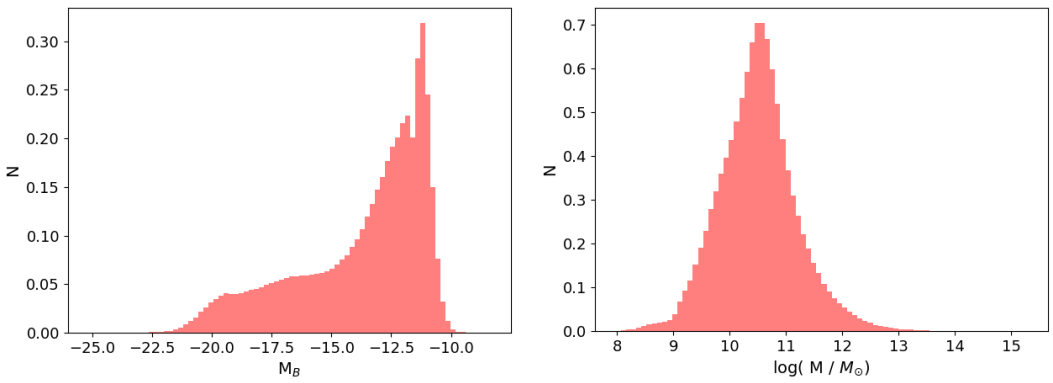

FIGURE 3.1: Properties of the z=0 galaxy sample. Left - the magnitude distribution. Right - the mass distribution. The presence of peaks in both distributions, and the sharp decline following M_B=-12.5 and log(M/$M_{\odot}$=11, reflect incompleteness of the sample due to the mass resolution of the TNG-300 simulation. <EMILY NOTE - EXPLAIN WHY.>

present in the mass distribution but will result in similarly sized pair catalogs in this region. The magnitude range itself spans approximately -22<M_B<-10, with the peak itself between -15<M_B<-10. However, the first signs of incompleteness come a little before the peak. Shortly before the decline begins, there is a fluctuation in the trend showing fewer galaxies of a given magnitude than would be expected. This occurs just a little above M_B=-12.5. Therefore, incompleteness of the catalog is treated in this investigation as affecting M_B>-12.5.

The root cause of the incompleteness is the mass resolution of the simulation TNG-300 has a limited mass of a particle that it might resolve, and that it might recognise as a galaxy. One particle of dark matter for this run has a mass of 5.9×10^7 $M_\odot$, while one baryonic particle has a mass of 1.1×10^7 $M_\odot$. This sets a lower limit on the mass of a galaxy, and means that formation of low-mass or low-brightness galaxies may not be entirely accurate. This leads to an incompleteness in the distribution of faint galaxies, seen in this distribution as the declining number of galaxies with masses M<10^{11} $M_\odot$ and magnitudes M_B>-12.5.

Considering this M_B>-12.5 limit was crucial to begin this investigation. As discussed in Methodology, to see how magnitude affected the intervelocity peak, ranges of 3 magnitudes were defined and the intervelocity peak calculated for each interval. The incompleteness defined the lowest magnitudes that could be investigated with reliable results. The six ranges used are defined in Table **??**, along with the number of pairs in each resulting catalog.

3.2 Properties of the Pair Catalogs

3.2.1 Pair Catalog Overview

Seven ranges were defined in total, in order to span the full range of magnitudes present in the base catalog. However, following pair selection, only thirteen pairs were found in the brightest range -24.5<M_B<-21.5. Therefore, it has been omitted from the analysis. The magnitudes and sizes of each pair catalog are presented in Table 3.33.

The brightest catalog remaining, -23<M_B<-20, is also substantially smaller than the other five catalogs, and so will have larger associated errors. However, it does contain enough galaxies to be analysed.

Absolute Magnitude Range	Number of Pairs
-23<M_B<-20	1055
-21.5<M_B<-18.5	7391
-20<M_B<-17	11490
-18.5<M_B<-15.5	9921
-17<M_B<-14	9533
-15.5<M_B<-12.5	14384

TABLE 3.1: The defined magnitude ranges of, and the number of pairs in, each pair catalog.

A point of note is the particularly large number of pairs found in the -20<M_B<-17 range. If only the number of galaxies in each base catalog is taken into account, a number of pairs somewhere between 7400 and 9900 - the total number of pairs in the ranges either side - would be expected. However, the reason for this is likely due to the nature of the isolation criterion. Only galaxies that have magnitudes M_{Bthird}<M_{Bcomp}+2.5 are considered as potential thirds. Therefore, the number of potential thirds increases significantly with every subsequently fainter catalog. This means pairs in this -20<M_B<-17 catalog are less likely to be discarded due to close proximity of an influential neighbour. In addition, beginning from the next catalog, with magnitude range -18.5<M_B<-15.5, the magnitude range for possible thirds will coincide with the peak in the magnitude distribution described above. This increases the potential number of thirds further.

The faintest catalog, and to some extent the second faintest, have one additional source of error: the incompleteness of the base catalogs as described earlier on. Both catalogs, but moreso the faintest -15.5<M_B<-12.5, have the problem where the magnitudes of their potential thirds fall into the incomplete region M_B>-12.5. This means that there is likely an overabundance of pairs in these regions, as pairs that may have a close neighbour in a world with better resolution would not have one in this simulation. However, while this third may not be resolved as a galaxy, it may still be resolved as a halo, something that will not be so frequently picked up by the SubFind algorithm. Therefore, more contamination will be observed as the magnitude range for potential thirds moves into the incomplete region.

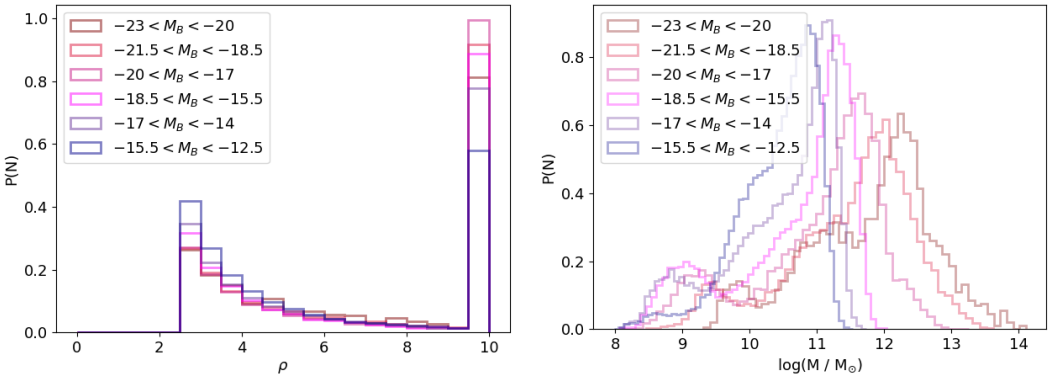

FIGURE 3.2: General properties of each pair catalog. Left: the distribution of isolation criterion ρ. Right: the distributions of masses for each catalog.

Comparing this to Pawlowski et al.'s study gives promising results. Their investigation found 7840 pairs with absolute magnitudes greater than -18.5 in this snapshot of TNG-300. This best aligns with the range -21.5>M_B>-18.5 in this investigation, where 7391 pairs were found; the 450 missing pairs are likely those that have one or both pair members with M_B<-21.5.

All six catalogs have very similar trends in the distribution of their isolation parameter ρ. The number of pairs decreases smoothly with the isolation parameter, with the exception to this being ρ=10 due to its use as the default value for highly isolated pairs. This is discussed more in Methodology. There are subtle differences between the catalogs, however, something that can be seen in Figure 3.2. The lowest bin for ρ shows a higher number of faint pairs than brighter pairs; in contrast, the highest bin shows more bright pairs. This is not a concrete trend - the brightest three catalogs all have approximately the same proportion of less isolated pairs, and the catalog with the highest proportion of highly isolated pairs is not the brightest. However, that the order of catalogs with the highest proportion of less isolated pairs is the same as the order of catalogs from faintest to brightest is an interesting feature, one that likely arises from the number density of galaxies. Fainter galaxies are far more numerous, as shown in Figure 3.1 - therefore, are more likely to be less isolated by chance. In contrast, the bright galaxies have a lower number density, and so will naturally be further apart.

The right hand panel shows the mass distributions of all galaxies in each of the six pair catalogs. There is a clear decrease in mass with increasing magnitude, as

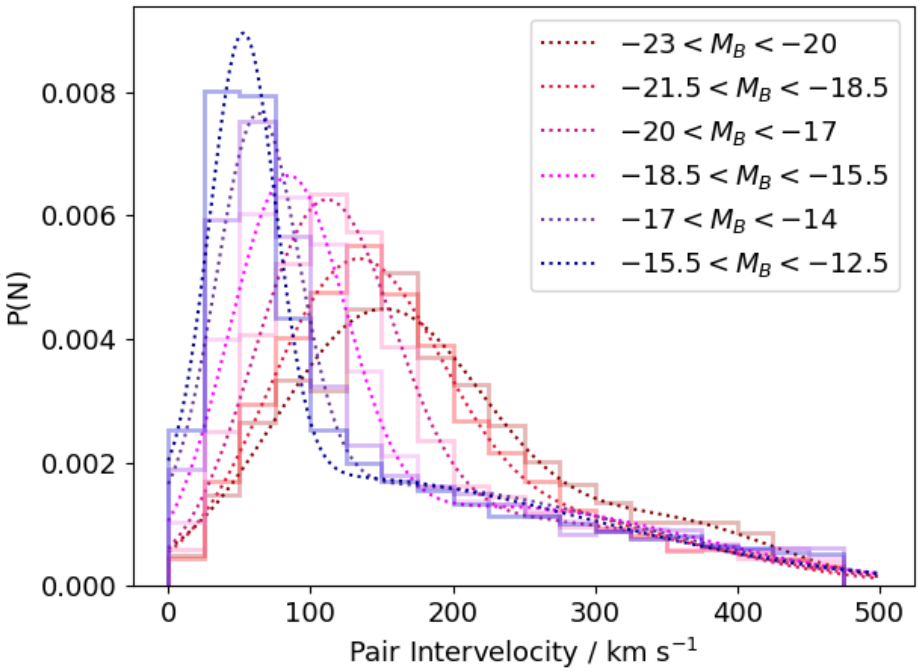

FIGURE 3.3: The intervelocity distributions of all six magnitude ranges. As the magnitude range of the catalog grows fainter, there is a clear trend as the intervelocity peak moves to the right and grows narrower.

would be expected. The times at which these galaxies accreted their mass can also be seen, by observing the different mass distributions for different redshifts.

3.2.2 Intervelocity of Galaxy Pairs

The intervelocities were calculated using the full data from the simulation, and the intervelocity distributions are shown in Figure **??**. The peaks show a gradual shift to the right as the galaxies grow fainter. The precise peak positions are as follows: 150 ± 9 km s^{-1}, 130 ± 3 km s^{-1}, 109 ± 2 km s^{-1}, 83 ± 2 km s^{-1}, 63 ± 2 km s^{-1}, and 53 ± 2 km s^{-1}.

The uncertainty in the peak position is seen to decrease with magnitude, and this is likely due more than anything to the sharpness of the peak. However, with the brightest catalog -23<M_B<-20, it may also be due to noise. With only 1055 pairs, it has fewer galaxies than the next least populated catalog by a factor of approximately 7. This naturally results in more random fluctuations.

The peak value for the magnitude range -21.5<M_B<-18.5 of 130 ± 3 km s^{-1} is of particular interest as it covers a similar range of magnitudes to previous works.

Scarpa et al.'s peak value at 132 ± 5 km s^{-1} was predicted for ρ>5 for a catalog that also included galaxies with slightly lower absolute magnitudes. However, there is an almost complete overlap of the uncertainties between the two - a promising result. Similarly, Pawlowski et al.'s analysis of the TNG-300 data produced a result of 132 ± 1 km s^{-1}. This, again, fully overlaps with the result produced in this study. While this is not a new result, given that this study was based on that of Pawlowski et al., this is not a particularly groundbreaking comparison. However, it does confirm that the results produced are accurate.

The other interesting feature of the intervelocity distributions is the decreasing width as absolute magnitude decreases. This is something that cannot be attributed to noise alone - if this were the cause, then the second most highly populated catalog, -20<M_B<-17, would be very narrow. This is not the case, something that suggests that that there is some dependency of the peak width on the magnitude. Further investigations show that the peak width is always 45%-50% of the peak height, something that points towards a universality of the structure- were the peak intervelocity value normalised to 1, the width would then be approximately constant across all values.

3.2.3 Investigating Isolation

Following the distributions of the isolation parameter shown in Figure 3.2, ranges to further investigate ρ were defined. These are $2.5 \leq \rho < 5$, $5 \leq \rho < 10$, and $\rho \leq 10$. These have similarities with Pawlowski et al.'s ranges for pairs that qualified for the catalog, then fairly isolated pairs, then highly isolated pairs, where ρ was defined as >2.5, >5, $\geq$10 respectively. However, the most significant difference in these ranges is that while Pawlowski et al. only removed the less isolated pairs, this investigation did not do so - this was with the aims of understanding the effect of lower isolation on galaxy pairs.

However, as shown in Figure 3.4, there does not seem to be large variation in the peak positions. Due to the size of the brightest catalog, -23<M_B<-20, the isolation-limited catalogs had fewer than 1000 pairs - this was the lower limit on pair total for analysis, so there was no analysis of the isolation for this sample. For all other samples, the shapes of the peaks are approximately similar - this is especially true

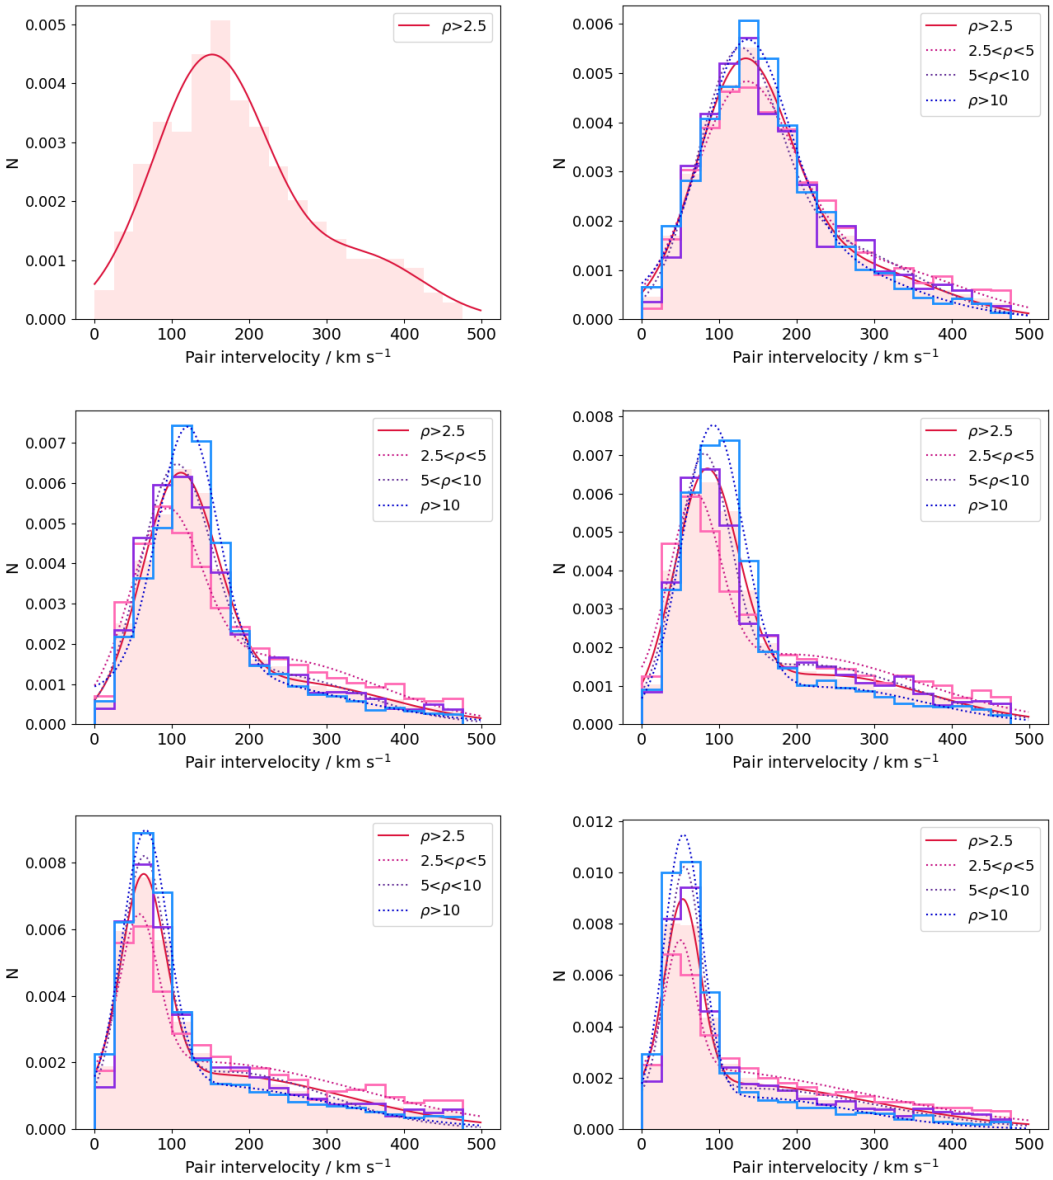

FIGURE 3.4: The intervelocity peaks of each magnitude range with varied cuts in isolation. It is seen that in the majority of cases, the limits imposed on isolation criteria do not appear to cause substantial variation in the peak position. Top left: -23<M_B<-20. Top right: -21.5<M_B<-18.5. Middle left: -20<M_B<-17. Middle right: -18.5<M_B<-15.5. Bottom left: -17<M_B<-14. Bottom right: -15.5<M_B<-12.5.

of two faintest samples. For the middle ranges, -20<M_B<-17 and -18.5<M_B<-15.5, there is a small change in peak position. As isolation increases, the peak shifts to the right. This aligns well with the findings of Pawlowski et al. - as their samples became stricter towards higher values of the isolation parameter, they also found that the peak shifted to the right.

A potential cause of this rightward shift is the bias in the method to exclude hierarchical pairs. Pairs with one member at the brighter end of the magnitude range and its partner at the fainter end will have a wider range of galaxies checked for their isolation criteria and a higher chance of exclusion, despite the fact that the high mass of the primary may actually mean the third found by the algorithm is not particularly influential with respect to the pairs dynamics.

It can also be seen that the noise around the peak increases with decreasing isolation, as evident in the lower probabilities of the peaks for lower ranges of isolation parameter. This suggests there are likely more false pairs in the lower isolation cuts, something that rather matches expectation.

3.3 Predictions in MOND

Using Equation 1.3, predicted values for the intervelocities in MOND could be calculated, and these distributions are shown in Figure 3.5. The most striking feature of the MOND distributions is the narrow shape. The MOND intervelocities have been derived solely from the absolute magnitudes, and so little noise outside of the peak would be expected due to the strict limitations imposed on the absolute magnitude - Due to this, the fits for the MOND curves are not double Gaussians as the other intervelocities, but single Gaussians.

This narrow shape is not something that is seen in the true intervelocity peaks. This, however, is more a failing of how the MOND formula was applied than the theory itself; Equation 1.3 refers to circular orbits. As discussed by Scarpa et al., the peaks would be expected to broaden for non-circular orbits - hence, peaks that are too narrow would be expected.

Regarding the peaks, however, the general trend in these predicted intervelocities is much the same as in the true 3D intervelocity. The peak position shifts to the

right with decreasing absolute magnitude. However, while for the true intervelocities the peak width decreases with magnitude, such a phenomenon is not observed in the MOND predictions. This, too, is likely an effect that can be ascribed to dependence on magnitude. The widths of the MONDian peaks, in comparison to the true intervelocity peaks,

The peak positions are as follows: 155 ± 1 km s^{-1}, 130 ± 1 km s^{-1}, 104 ± 1 km s^{-1}, 73 ± 1 km s^{-1}, 52 ± 1 km s^{-1}, and 34 ± 1 km s^{-1}. There is little overlap within the bounds of uncertainty for many of the catalogs, especially the fainter ones - at first, this may seem to make a promising prediction. The comparison between CDM and MOND predictions is shown in Figure 3.6.

The peak intervelocity values for MOND show a sharper decrease than those produced directly by TNG-300.

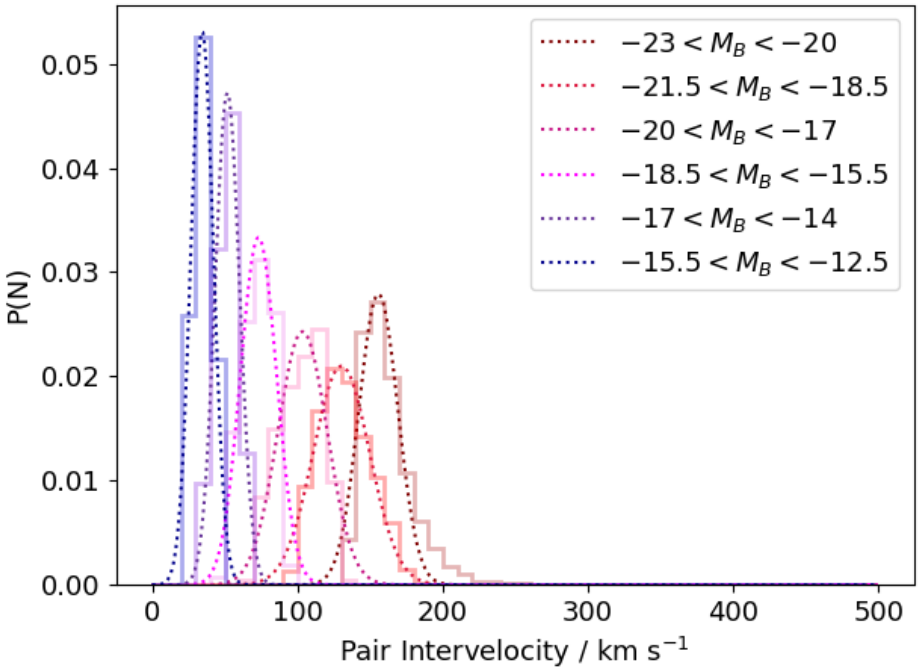

FIGURE 3.5: The intervelocities of each magnitude range as predicted by MOND.

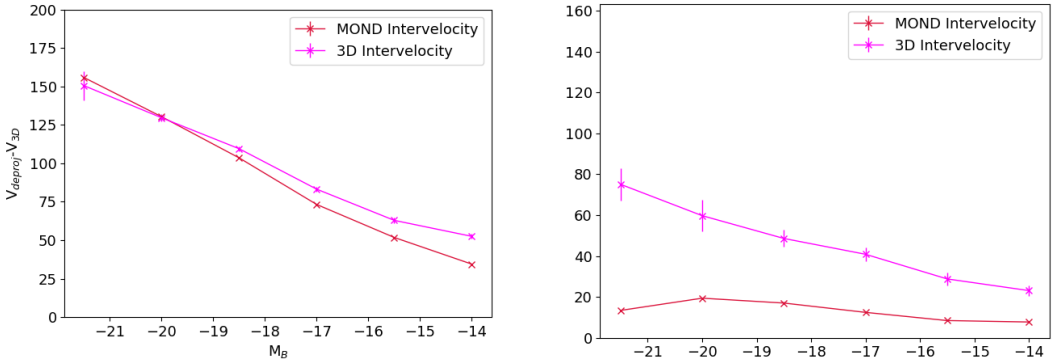

FIGURE 3.6: Parameters of the intervelocity peaks in MOND vs those of the true intervelocity peaks. Left: the position of the intervelocity peak. Right: the σ of each Gaussian.

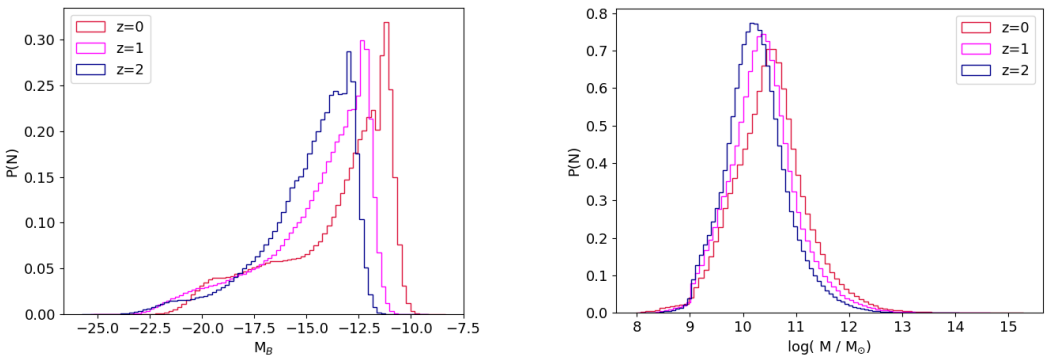

FIGURE 3.7: The mass and magnitude distributions for the z=1 and z=2 samples. Note the shift in both to the left - as expected, galaxies become brighter but less massive at higher redshifts.

3.4 Redshift Investigations

3.4.1 General Properties

Repeating the investigation at higher redshifts z=1 and z=2 gave two additional sets of catalogs, seven magnitude ranges for each snapshot. Comparing mass and magnitude ranges for the base catalogs at all redshifts, as in Figure 3.7, shows precisely the pattern theory would predict. At higher redshifts, both the mass distribution and the magnitude distribution are shifted to the left - this is due to hierarchical structure formation[Ryman, 2016], as galaxies had not yet accreted all their mass yet, and the high rates of star formation leading to the bright magnitudes. On average, halos are smaller at higher redshifts, but tend to have younger stellar populations - this makes them brighter.

The mass range moving to the left brought the peak - and thus, the incompleteness limit - to a lower mass. In the z=2 snapshot, this limit has been brought closer to $10^{10}M_{\odot}$, approximately ten times smaller than the limit in z=0. Similarly, the incompleteness limit for magnitude has also shifted to the left for both redshift values, with the limit for z=2 moving to approximately M_B=-14. This was critical as this meant that the lowest magnitude range has a large overlap with the incomplete region, leading both to missed pairs as one or both pair members are not present, and to false pairs as pairs that should have an influential nearby third do not have that third resolved, and so are included in the catalog. This leads to a greater level of contamination, something that should be considered during analysis of these catalogs,

It should also be noted that the shift to the left did result in a catalog for z=2 where -24<M_B<-21, a region that had not been sufficiently populated by either the z=1 or z=0 snapshots. This range was included in the intervelocity analysis.

Figure 3.8 goes a little further, showing that the transition to less massive galaxies does not correlate directly with the transition to brighter magnitudes. The mass distributions are shown to be moved to the left at higher redshifts even within a fixed magnitude range. While this is an established result, it is important to highlight as this demonstrates clearly that the different redshift catalogs for a set magnitude do have different masses.

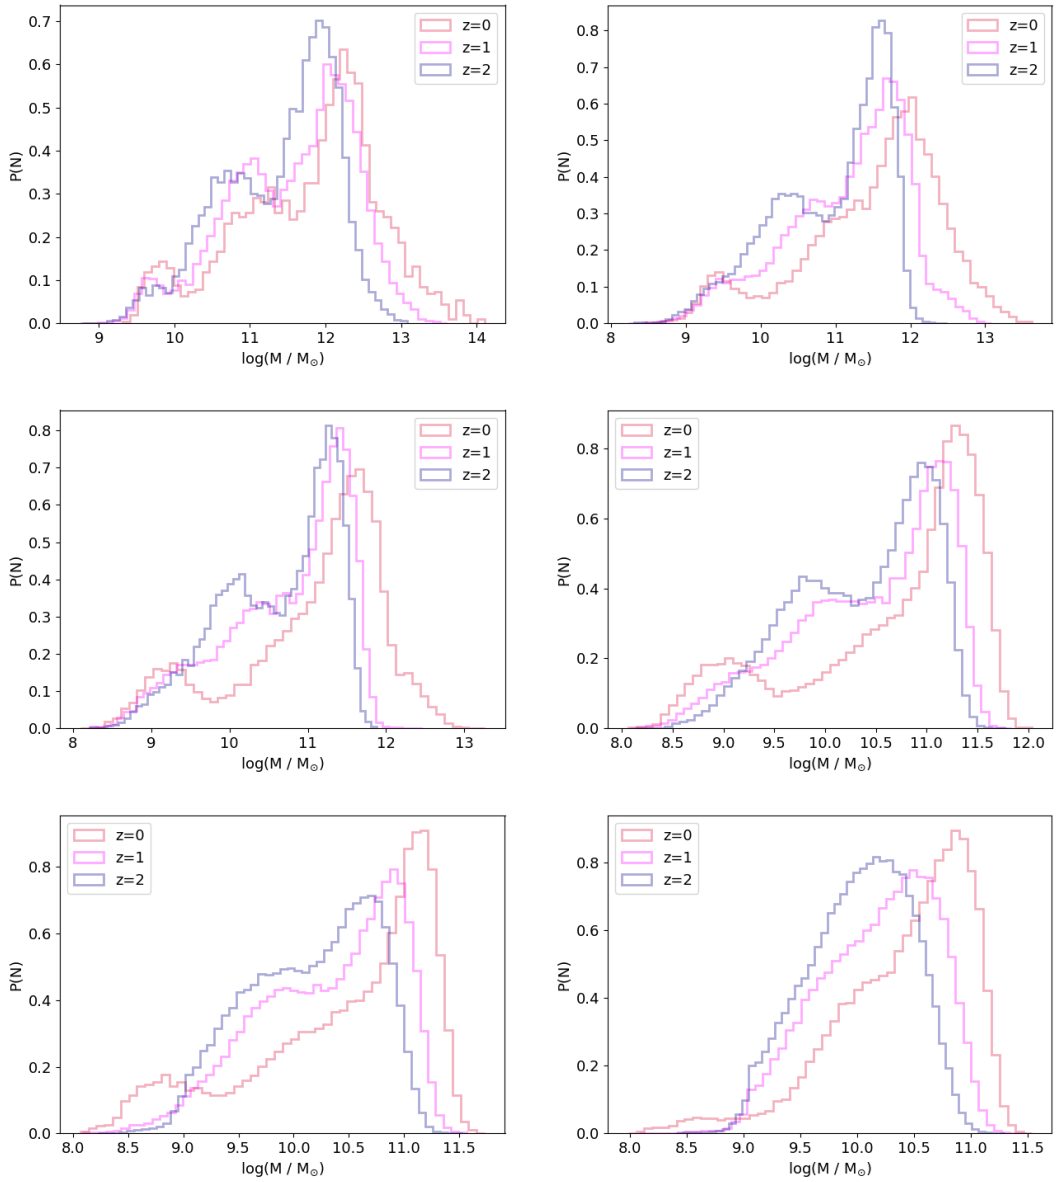

FIGURE 3.8: The mass distributions for the different redshifts, separated by magnitude catalog. Top left: -23<M_B<-20. Top right: -21.5<M_B<-18.5. Middle left: -20<M_B<-17. Middle right: -18.5<M_B<-15.5. Bottom left: -17<M_B<-14. Bottom right: -15.5<M_B<-12.5. Note the shift to the left in each case, confirming that galaxies were less massive regardless of magnitude.

3.4.2 Intervelocities

Investigating the intervelocity distributions for each additional redshift z=1 and z=2 gives rise to the top panels of Figure 3.9. In both cases, a series of peaks similar in appearance to those in Figure 3.3. Much the same as for z=0, the peak value for the intervelocity decreases as magnitude increases, as does the peak width. The positions of these peaks are laid out in the caption of Figure 3.9.

The comparisons between the three redshift snapshots are shown in the bottom panel of Figure 3.9. There is little to say on grounds of comparison for the brightest range, it only having reliable data for z=2.

For the second brightest range, the uncertainties are all noticeably larger than for all the ranges that follow. This is likely due to peak size, given that the brighter ranges are wider than the fainter ones, but could also be due to catalog size. The brighter catalogs have fewer pairs, a pattern established in z=0 that persists in z=1. Therefore, these peaks are likely to be noisier, and so have greater associated errors. An important note is that within the bounds of these uncertainties, for both the z=1 and z=2 snapshots it cannot be said that there is a difference between the intervelocity peaks for this magnitude range, -23<M_B<-20, and the next brightest, -21.5<M_B<-18.5.

The brightest magnitude range in z=2 would seem to contradict the second theory, having a smaller error than the second range. However, by the brightest range, the width of the intervelocity peak eclipsed the full velocity range 0 to 500 km s^{-1}, leaving little room for noise outside of the peak. Therefore, when the curve fitting function looked for a double Gaussian, it could not find the background peak, and instead found a fluctuation at a lower intervelocity to assign a peak to. This can be seen at approximately v=100 km s^{-1} right panel of Figure 3.9, and so forces a lower uncertainty in the identified peak value. This may also have an effect on the effectiveness of the peak fitting; closer observation of the peak suggests the true peak is slightly further left.

The range -21.5<M_B<-18.5, the one best matched to the observational catalog, shows overlap between the uncertainties of peak position in z=1 and both other snapshots. It does not, however, show overlap between the peak positions for z=0

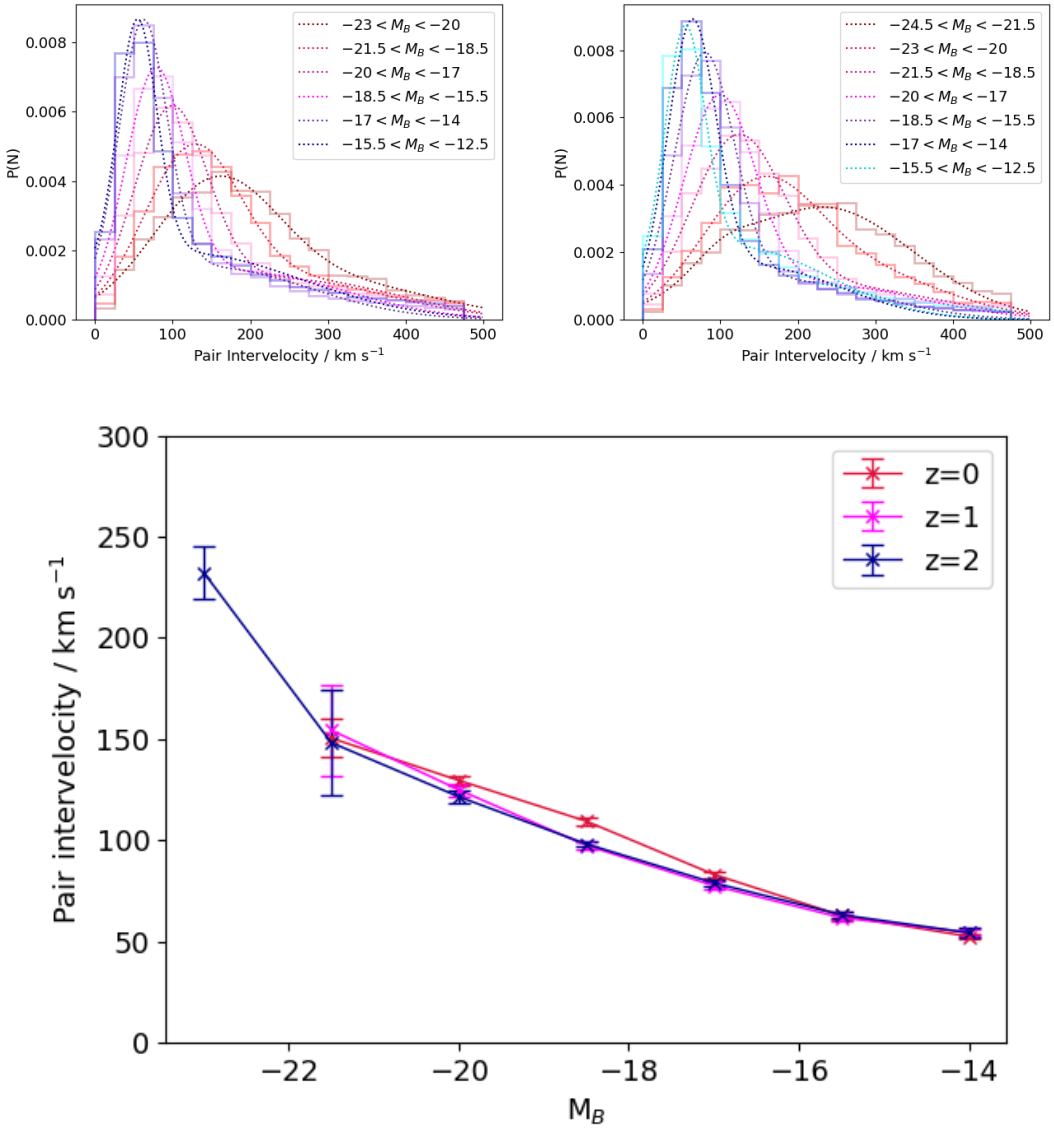

FIGURE 3.9: The intervelocity peaks for z=1 and z=2, plotted against absolute magnitude. Each catalog was assigned a single magnitude for the sake of this analysis, using the average of upper and lower limits of the catalog's magnitude range. Top left: The peaks for z=1. From the brightest to lowest magnitude ranges, the peaks are: 154±22 km s^{-1}, 124±3 km s^{-1}, 94±2 km s^{-1}, 78±2 km s^{-1}, 62±2 km s^{-1}, and 54±2 km s^{-1}. Top left: the corresponding peaks for z=2 are: 148±26 km s^{-1}, 122±3 km s^{-1}, 98±2 km s^{-1}, 79±2 km s^{-1}, 63±2 km s^{-1}, and 54±2 km s^{-1}. Finally, the brightest peak for z=2 is found at 232±12 km s^{-1}. Bottom: The peak intervelocity is seen to decrease with absolute magnitude across all redshifts. z=1 and z=2 show no difference within the bounds of their uncertainties. The z=0 snapshot does show a distinct difference from the others, most notably in the range averaging M_B=-19, but in the other ranges it also shows some overlap with one of both of the other catalogs.

and z=2. This is a significant difference; however, the difference is small enough that if the uncertainties have been underestimated even slightly, the difference will no longer be significant. It is not, therefore, so interesting a difference as that observed in -20<M_B<-17. While redshifts z=1 and z=2 overlap, as they do with all magnitudes, the snapshot for z=0 has an intervelocity peak that is significantly higher than the others for this magnitude band. However, it is still well within the bounds of the peak width for both z=1 and z=2 - as such, this can likely be attributed to sample variance.

The remaining catalogs all have peak intervelocities that are the same, within the bounds of their uncertainty. Thus, it can be concluded that redshift does not seem to have an influence on the position of the intervelocity peak.

3.5 Making observable predictions

The low differences in the intervelocities measured directly from the simulation introduce the lowest possible amount of error. No errors were taken into consideration in the production of these pairs - in contrast, for Pawlowski et al.'s mock observation, the full sample used errors extracted from the HyperLEDA catalog. Their mock observation also included a limit on apparent magnitude. However, the characteristic element of their method their deprojection algorithm. To investigate the effect of observation on these pairs, the deprojection algorithm has been applied to all catalogs. Figure 3.10 shows the deprojection intervelocity distributions for each snapshot.

The deprojection did generally recreate the pattern shown by the TNG velocities that were used exactly as they were. The intervelocity peak positions are all approximately the same as those found by the direct analysis, within the bounds of their uncertainty, and despite a far greater amount of noise, the peaks are still fairly visible. However, it can be seen that the pattern is not nearly so precise.

Studying the z=0 result, the smooth downwards curve in height has been entirely lost. Similarly, the relationship between peak position and the width has been partially obscured - the second brightest catalog is wider than the first, for example, and there is little discrimination between the widths of the faintest three catalogs.

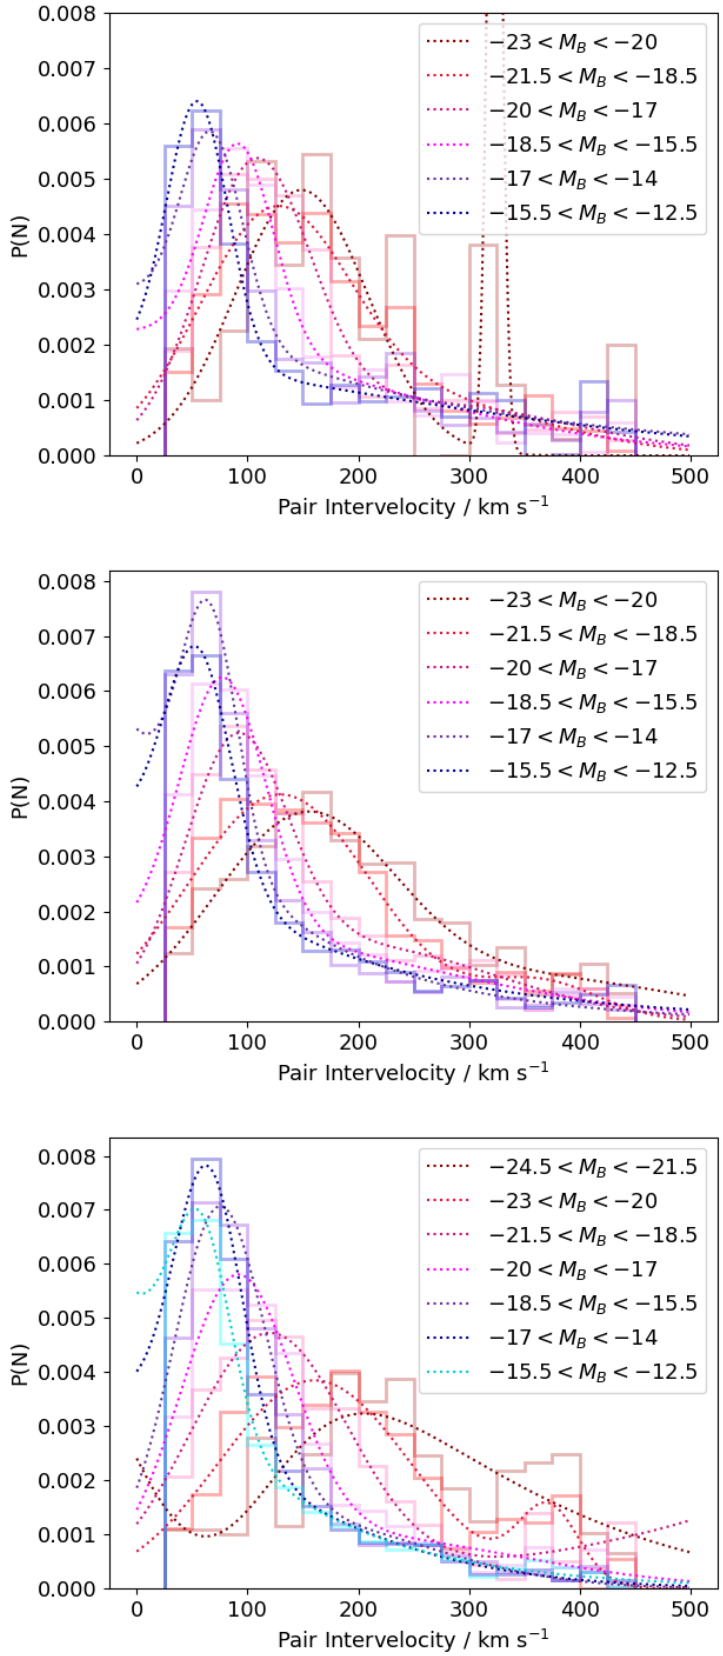

FIGURE 3.10: The histograms and fits for the deprojected intervelocities of redshifts z=0 (top), z=1 (centre), and z=2 (bottom). Comparison with the plots in Figure 3.3 quickly demonstrates the level of noise that is introduced when the intervelocities are deprojected; despite this, the peak positions have been recovered within the bounds of the uncertainties.

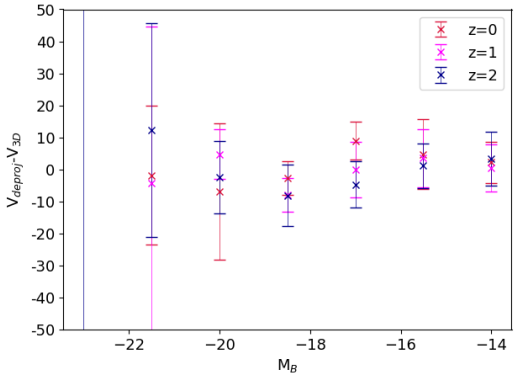

FIGURE 3.11: The differences between the deprojected velocity peaks and the true 3D peaks.

TABLE 3.2: Caption

Figure 3.11 then finally shows the avereage deviation of the deprojected peak from true TNG-300 intervelocity. While the fit of the top magnitude range, -24.5<M_B<-21.5, appears well-defined in Figure 3.10, it can be seen more clearly here that the curve fitting function was not able to find physical parameters for this curve. For the remainder of the data, however, there are fairly reasonable differences. Even for those values with high uncertainties, such as the -23<M_B<-20 region, the actual error introduced by the deprojection is little more than 10 km s^{-1}, and this is a trend that continues through the sample. The error introduced by deprojection is rarely more than 5 km s^{-1}, and this does align with the average uncertainty in peak position as calculated by the curve fitting algorithm.

ΔV km s^{-1}	z=0	z=1	z=2
-23<M_B<-20	149 $\pm$ 12	150 $\pm$ 26	160 $\pm$ 7
-21.5<M_B<-18.5	123 $\pm$ 18	130 $\pm$ 5	119 $\pm$ 8
-20<M_B<-17	106 $\pm$ 4	90 $\pm$ 3	90 $\pm$ 8
-18.5<M_B<-15.5	92 $\pm$ 4	78 $\pm$ 7	74 $\pm$ 6
-17<M_B<-14	68 $\pm$ 9	65 $\pm$ 7	64 $\pm$ 5
-15.5<M_B<-12.5	54 $\pm$ 5	55 $\pm$ 5	57 $\pm$ 6

TABLE 3.3: The deprojected intervelocities ΔV of each magnitude range at each redshift.

Chapter 4

Discussions, Conclusions and Outlook

4.1 Discussions

4.1.1 Redshift z=0 Intervelocity

The values for the intervelocity peaks at redshift zero were produced and these peaks plotted. It has been shown that the position of the intervelocity peak is dependent on the absolute magnitude of the pair, and it has also been shown that the width of the peak is dependent on the magnitude.

The widths have been shown to have an approximately constant ratio with peak position - that is, that the width of the peak is in every case between 45% to 50% of its peak velocity. This indicates some underlying universality to the dynamics. When normalised to a peak intervelocity v=1, the width of the peak would be approximately the same regardless of magnitude, despite the positions and widths themselves depending on magnitude. This adds to Pawlowksi et al.'s argument that the peak in ΛCDM is a physical effect, even if its exact origin is unknown. However, this effect does get lost during deprojection due to the additional noise - this will make it a difficult feature to verify through observation.

4.1.2 Variation with Isolation

In most catalogs, there is little variation of the intervelocity peak with isolation cutoff so long as that cutoff ρ>2.5. However, in certain scenarios, a gradual increase in the intervelocity is seen for different minimum values of ρ - this is most evident in the

ranges -20<M_B<-17 and -18.5<M_B<-15.5 for z=0. Interestingly, this is also where the intervelocity peak shows the most deviation from the trend set by the other redshift snapshots.

It is possible, however, that due to the bin size used - 25 km s^{-1}, a necessary limitation so that the deprojection algorithm might be used on the data later on - that the shift to the right does exist in those fainter magnitudes, but it so small a shift that it is no longer resolved in the analysis. For the faintest catalogs, the peak is only made of two bins, something that may not be enough to allow for the detection of a shift so subtle.

If this were simply statistical variance, there would more than likely be a sign of a catalog shifting to lower intervelocities with increasing isolation, but this is not the case. Therefore, a physical origin of this phenomenon is considered. Brighter galaxies, with there being fewer of them as discussed in Results, are far more likely to form pairs that, by the metrics considered in this study, are isolated. Fainter galaxies, however, are far less likely to form isolated pairs, and this is due to the large number of them.

An interesting further study of this would consider varying the definition of the isolation parameter, ρ. The study by Moreno et al. defined their isolation parameter in such a way that it accounted for the mass of the third, as well as its proximity, and this sort of approach may help shed some light on why the isolation parameter does not affect intervelocity peak position in some predictable way. A parameter defined as such would not only confirm whether or not a galaxy is isolated, as ρ does in this study, but would also allow insight into just how its neighbours affect it.

Other potential alterations would involve varying how the nearest third is searched for. An area of radius $10r_p$, where r_p is the pair separation, was searched for this investigation - this does make some attempt to take the pair's properties into account. However, extending this range for more massive galaxies while decreasing it for less massive galaxies may also serve to make the parameter ρ more effective at providing physical insights.

A final interesting note on the isolation is this: while it is almost certain that

the faintest catalog, -15.5<M_B<-12.5 is incomplete for z=2, with significant contamination by pairs that are not isolated but whose thirds are not resolved by the simulation, there is no significant difference between its intervelocity peak and those predicted by the other two catalogs. This can lead to one of two conclusions: either, the isolation criterion is not so influential over the intervelocity peak as has been implied; or, the selection criteria work better than have been so far discussed.

An important thing to note at this point is the other selection criterion focused on a pair's isolation: the reciprocity criterion. While this does not confirm that a pair is isolated, it does confirm that the pair is not part of any complex system. Referrnig back to the results of Moreno et al., the reciprocity criterion alone makes it likely that the gravitational interactions of the pair are dominated by its members more than any nearby galaxy. An investigation of the binding energy of these pairs would provide a greater insight into this suggestion - potentially even developing a parameter for isolation that takes this energy into account, much like was done in Moreno et al.

4.1.3 Redshift Catalogs - Varying Mass

The varying masses of the individual magnitude catalogs, despite their constant velocities across different redshifts, might strike as concerning for some. If intervelocity does not vary with halo mass, that would appear to suggest against ΛCDM and towards MOND. However, what this conclusion does not take into account is that while the galaxies will accrete mass over time, it is likely that the mass accreted was already close by to the pair before it became bound to one halo or the other. This would mean that these particles would already be influencing the gravitational dynamics of the pair, before they registered in TNG-300 as being a part of one of the halos. An incredibly useful investigation to follow this theory up is discussed in Outlooks.

4.2 Conclusions

The results of this investigation are somewhat inconclusive. It has been shown that through pure simulated data, there is a significant difference between the predictions of MOND and those of ΛCDM in both the slope of the intervelocity peak, with MOND decreasing more quickly, and the variation in the widths. Even considering the widening of the MONDian distributions given non-circular orbits, the relationship between peak position and width in ΛCDM is not present in MOND. However, through deprojection alone these differences are likely to be obscured in observation due to the nature of statistical deprojection, and the noise it introduces. A large enough catalog of galaxy pairs may avoid this - however, this will take significant advances in observational data compared to what is currently available.

4.3 Outlook

Potential future investigations could take the study of galaxy pairs in a great many directions. Consideration of galaxy masses to improve and expand upon the isolation parameter might provide key information as to just how isolation affects the dynamics of these galaxy pairs, something that due to the gradual increase seen in two catalogs from z=0 is still yet unclear.

Beyond that, investigations of the mass accretion of galaxy pairs through redshift, studying how the environment changes and how this may affect their gravitational interactions and their preferred intervelocities may show that the intervelocity does or does not depend on dark matter halo mass. By identifying this, a critical distinction between MOND and ΛCDM might yet be made.

And further improvements upon the predictions of MOND made in this study could lessen or extend the disagreements found.

Finally, additional observation data at fainter magnitudes will help to verify the predictions made.

Appendix A

RYID

A.1 Background

A.1.1 Introduction

Perhaps one of the greatest open questions in physics is to the nature of dark matter. It is the centre of the standard model for modern cosmology, an unknown substance that makes up 80% of matter in the universe [Ryman, 2016], and its effects are seen in gravitational interactions from the scales of galaxies and larger. However, at the time of writing, it has not yet been directly observed; strictly speaking, it is not yet known what it actually is. Much of what is known of its nature is discussed in Cirelli's review [Cirelli, Strumia, and Zupan, 2024]. There are those with whom this idea is found to be rather disagreeable, who are motivated to find alternative descriptions as a result - the most prevalent of these is Modified Newtonian Dynamics, or MOND [Milgrom, 1983a; Milgrom, 1983b; Milgrom, 1983c].

There are a great many systems in which MOND has argued against the existence dark matter and validity of its associated model, cosmological cold dark matter (hereon ΛCDM). A handful of these tensions will be discussed in Section A.1.2. This paper will briefly discuss the two paradigms and summarise their contentious history; it will not, however, discuss the technical details of either. To do either justice in that regard would require a dedicated paper; for such things, an interested reader is encouraged to review textbooks [Ryman, 2016] or reviews such as [Cirelli, Strumia, and Zupan, 2024; Bertone and Hooper, 2018] for ΛCDM - it being the dominant cosmology - and the pilot papers of MOND [Milgrom, 1983a; Milgrom, 1983b; Milgrom, 1983c], as well as recent reviews [Famaey and McGaugh, 2012; Banik and Zhao, 2022] that summarise its development, achievements, and compare them with ΛCDM.

The focus of this paper, however, is to review the recent works that have brought to light a peak in the intervelocities of galaxy pairs, and the ways in which the existence of that peak has been interpreted and investigated in the contexts of ΛCDM and MOND. We also propose future work that hopes to make testable predictions as to how this peak will behave and evolve, with the aim of investigating whether ΛCDM can or cannot explain the existence of the intervelocity peak.

A.1.2 The need for new gravity

Both ΛCDM and MOND provide answers to the same problem: there is something fundamentally missing in the current understanding of gravity. At present, gravity is best described by Einstein's general relativity [Einstein, 1916], one of the most thoroughly tested theories in physics that has withstood a century of experiment [Will, 2014]. It is so well-respected that obeying its laws is a requirement for any new theory - in fact, this was a problem for MOND until Bekenstein's TeVeS [Bekenstein, 1984]. However, the evidence suggesting something is missing is incredibly damning.

It first became apparent in the 1930s, when observations of the Coma Cluster led Zwicky to conclude that there must be significantly more mass present in the cluster than can be observed [Zwicky, 1933]. He measured the velocity dispersion in the cluster and found it far too high. The observed values were approximately 1000 km s^{-1}, in stark contrast to the predicted value of 80 km s^{-1}. This difference is significant enough, in fact, that it jeopardises the stability of the cluster; if the velocity dispersion were truly so high, the cluster would not be gravitationally bound if only the luminous mass were present. The only way it could be bound would be if there were some additional mass that could not be seen with a telescope - and Zwicky referred to this as 'dark matter'.

40 years later, it was discovered that galactic rotation curves are, at large radii, flat. This is in contrast with what is known about ordinary Newtonian systems; rotation curves describe how the rotational velocity of a galaxy varies with its radius, and for systems such as the Solar System, are expected to decrease with radius. However, a series of groundbreaking results in the 1970s - with key papers on the flatness of rotation curves [Roberts, 1976; Rubin, Ford, and Thonnard, 1978] and the implications of these curves [Ostriker, Peebles, and Yahil, 1974; Einasto, Kaasik, and Saar, 1974] - all lead to one conclusion: that the masses of galaxies are far higher than what luminous matter suggests. The laws of general relativity held in such high regard that rather than question those laws, physicists came to the agreement that there must simply be matter we cannot see - this 'missing mass hypothesis', however, was not so readily accepted by all.

In 1983, Mordehai Milgrom published a series of three papers [Milgrom, 1983a; Milgrom, 1983b; Milgrom, 1983c] outlining his controversial new approach to the missing mass problem. Namely, that mass was not missing at all. He instead postulated that something was missing from the laws of gravity; that much like how general relativity only becomes noticeable in high-gravity systems, there exists some modification

to Newtonian gravity that dominates interactions in very low gravity. The most important constant in MOND is the boundary for this behaviour, a_0, with a value of 1.2×10^{-10}m s^{-2} as set by [Begeman, Broeils, and Sanders, 1991]. By changing the dependence of gravitational forces on distance to $1/r$ in this low acceleration regime - mathematically, when $a << a_0$ - he was able to explain the flat rotation curves of galaxies without needing any additional matter at all.

More recent victories of MOND include phenomena such as the baryonic Tully-Fisher relation [McGaugh, 2012], a well-established correlation between the total baryonic mass of a galaxy and its rotational velocity. In an ΛCDM universe, such a relationship is rather counterintuitive - if the dark matter quantities are several times that of the baryons, why would the dynamics be governed by the non-dominant component? While there have been complex numerical simulations within the context of ΛCDM that have been capable of replicating this relationship (as discussed here [McGaugh, 2012]) it is still considered by many to be explained far more naturally with MOND than with ΛCDM. Similarly, observations of colliding galaxy clusters have calculated some rather high velocities on collision - often too high for ΛCDM, as discussed in [Angus and McGaugh, 2008] with respect to the Bullet Cluster.

The Bullet Cluster itself is an interesting debate, however. A pointedly open question for MOND proponents and perhaps the most compelling argument for ΛCDM is the gravitational lensing around the Bullet Cluster, discussed neatly [Clowe, Gonzalez, and Markevitch, 2004]. The two main visible components of the cluster - the gaseous component, and the galaxies themselves - have been separated out by the collision. The gas, observable due to its X-ray emissions, has been slowed by ram pressure, while the individual galaxies pass through one another as as collisionless fluid. The gas is the most significant contributor to the mass of the cluster; as such, gravitational lensing around the cluster should be around the gas. However, this is not the case. Instead, the lensing centres on the galaxies. The bending of light in this manner is far trickier to explain with modified laws of gravity than it is simply using mass that cannot be seen, and it is because of this that the Bullet Cluster is often regarded as one of the most critical pieces of evidence for ΛCDM and against MOND.

It should be noted that this issue has been addressed in defence of MOND [Angus and McGaugh, 2008], where it is argued that as it is already established even within MOND that galaxy clusters must contain some missing mass, it only follows that mass in the Bullet Cluster does not behave as might be expected. This, however, suggests that while MOND can reduce the need for dark matter to explain the behaviour of galaxy clusters, it cannot erase it entirely - a clear victory for ΛCDM.

The tensions between MOND and ΛCDM have existed for quite some time, and even today the debate still is not settled. The baryonic Tully-Fisher relation is still a point of interest for MOND [Mistele et al., 2024], but ΛCDM is still the textbook cosmology. At present, the most hotly debated topic in the field is that of wide binary stars, where acceleration $a < 2a_0$. This should be low enough to demonstrate MOND behaviour, while being on a small enough scale to exclude dark matter (something not usually seen on the scale of individual stars).

According to two prominent 2023 papers by Chae [Chae, 2023] and Hernandez [Hernandez, 2023], statistical studies of wide binaries do show decisive evidence of non-Newtonian behaviour, displaying a change in behaviour as they transition into the low acceleration regime. This would pose a serious challenge to ΛCDM, as there is no predecent for dark matter acting on such a scale. An interesting note here is that both papers show a fundamental challenge faced by MOND - while both describe how this behaviour fits particularly well with AQUAL, a MOND-type theory set out by Bekenstein and Milgrom in 1984 [Bekenstein and Milgrom, 1984], they also refer to a range of similar MOND derivatives, a great number of which are potential relativistic theories of MOND. This lack of consensus is of course an issue with MOND and theories derived thereof. However, it in no way detracts from the claim that wide binaries display distinctly non-Newtonian behaviour at low accelerations.

This, fortunately for ΛCDM, is contested by a 2024 paper by Banik et al [Banik et al., 2024], which supports ΛCDM with a 16σ confidence. This is an incredibly decisive result - some may be concerned that it is too decisive for a single study. However, it is always worth noting when someone such as Banik, who has advocated for MOND in the past, claims to have a discovery with such confidence in the favour of ΛCDM.

It is with this debate in mind that we move on to the discussion of galaxy pairs. MOND is,

as has been demonstrated, a fearsome challenger - should behaviour be found that ΛCDM does not predict, there will soon appear a theory of MOND that does. The dynamics of galaxy pairs are a new way of testing ΛCDM - if predictions can be made that observations cannot match, it opens up the possibility of new physics, or perhaps opens the door for a theory such as MOND to explain what ΛCDM cannot.

A.2 Galaxy Pairs

A.2.1 An aside on redshift

Within this review, redshift will be discussed in units of velocity. Here and in the papers discussed herein, redshift is defined as

$$v = H_0 d + v_{pec} \tag{A.1}$$

where v is the redshift given as a recessional velocity. v_pec is the peculiar velocity of a galaxy, or its velocity that is unrelated to the expansion of the Universe, caused by its environment. H_0 is the Hubble constant, the value of which has been taken as .Note that for each paper that will be discussed, the respective authors have chosen slightly different values: 73 km s^{-1} Mpc^{-1} in Moreno's paper (see Section A.2.3); 70 km s^{-1} Mpc^{-1} in Nottale and Chamaraux's paper (see Sections A.2.2, A.2.3); and 67.74 km s^{-1} Mpc^{-1} in Pawlowski et al. [Pawlowski et al., 2022]. (see Section A.2.4).

The form of this equation was taken from Pawlowski et al, and any independent calculations undertaken as part of this review will assume the same value $H_0 = 67.74$ km s^{-1} Mpc^{-1} as defined there.

A.2.2 The Isolated Galaxy Pair Catalog

In order to understand gravity in its simplest form, the simplest system is the most useful subject. Nottale and Chamaraux in 2018 [Nottale and Chamaraux, 2018a] used the HyperLEDA database [Makarov et al., 2014] to produce a catalog of over 13 000 galaxy pairs that could be considered isolated in order to investigate their masses, mass-luminosity ratios, and the implications thereof. While a number of galaxy pair catalogs had been constructed in the past [Karachentsev, 1972; Karachentsev and Makarov, 2008], there had been none published to date with any more than 1000 galaxies. The Isolated Galaxy Pair Catalog (IGPC) was therefore a critical step forward with respect to sheer amount of data.

By using the HyperLEDA database to construct the catalog, they could incorporate data from multiple surveys together, covering a far larger fraction of the sky than if they had focused on just one. It does, however, mean that the errors can vary quite significantly throughout the catalog, depending on which survey the information is taken from - this is addressed within their paper, and is worth noting.

To select the galaxy pairs they introduce six criteria, and these will be reviewed in twos. The most important two first are the criteria that ensure a full enough sample of the galaxies can be observed. They are the boundaries for redshift, specifically from 3000 km s^{-1} $< v <$ 16000 km s^{-1}, and the absolute magnitude limit, $M < -18.5$.

The lower bound on redshift is to ensure that the redshift itself can be an accurate measurement of distance, and will not have significant errors from a galaxy's peculiar velocity. The upper limit, 16 000 km s^{-1}, should be taken in tandem with the magnitude limit, and Fig. A.1 depicts why. Until $M = -20.5$, the distributions of the galaxies all match the Schechter luminosity function [Schechter, 1976], which suggests that all the galaxies with those magnitudes have been identified and accurate redshift data recorded. Up until $M = -18.5$, it can be seen that while they no longer follow the Schechter luminosity function (demonstrating the lack of complete redshift data for these galaxies) the distributions appear to be consistent within reason regardless of redshift. This means that while galaxies are missing, a uniform sample can still be created without a preference for one distance range over another. Beyond this limit, the distributions begin to behave rather differently - the furthest distances drop off rapidly, suggesting that only a handful are observed beyond this point, while the closest remain fairly consistent until $M = -17$. To use these data would introduce a preference for closer galaxies - while $M = -18.5$ does mean that some galaxies are missing from the sample, the benefits from a wider range of data more than offset this. However, line 5 is already fast dropping away by $M = -18.5$. Any further galaxies would certainly have a far patchier sample, and this could interfere with the analysis. The discrepency between the observed galaxies and the luminosity function is noted to have an effect on the catalog, and this will be discussed later on.

The later criteria aren't so complex, but lead to interesting discussion. The next decision is how close two galaxies must be to be considered

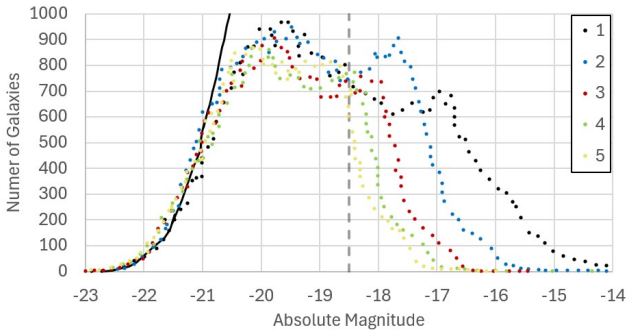

FIGURE A.1

After Nottale and Chamaraux [Nottale and Chamaraux, 2018a]. Figure 1 shows the number of galaxies of each magnitude found in different velocity ranges, against the solid black line - the Schechter luminosity function [Schechter, 1976], showing the number of galaxies expected at each magnitude. Lines have been labelled in line with the original figure, and their associated redshifts in km s^{-1} are as follows. 1) 3000 to 7000. 2) 7000 to 10 000. 3) 10 000 to 12 000. 4) 12 000 to 14 000. 5) 14 000 to 16 000.

a pair. Nottale and Chamaraux made the decision to cast the net wide, rather than too small; the projected separation can be up to 1Mpc, while the radial velocity difference must only be less than 500 km s^{-1}. Other papers that have been cited in the study of galaxy pair intervelocites, such as Moreno et al.'s 2013 study of the Millennium simulation [Moreno, Bluck, Ellison, et al., 2013; Springel, White, Jenkins, et al., 2005], have chosen much smaller values - their separation, for example, was only $250h^{-1}$kpc. This paper in particular is worth note as later papers will compare its results with those obtained from the IGPC as discussed in Section A.2.3, and so it is useful to draw comparisons between the two where appropriate.

Nottale and Chamaraux's criteria meant that even galaxy pairs with abnormally large separations or abnormally high differences in velocity would be caught; however, it also meant there was a significant chance for contamination by spurious pairs. This is assessed first by Nottale and Chamaraux themselves, who used the flattening of their radial velocity distribution at high velocities to estimate contamination to be approximately 10%. The issue is later addressed by two further papers that will be discussed in Sections A.2.3 and A.2.4 respectively.

The final criteria determine whether the galaxy pair is truly isolated or not. First, the exclusion of multiplets. Again drawing contrast with Moreno et al who does not do this, but refers within to many studies that do, this would mean that pairs are exclusive; galaxy A must be the closest to galaxy B, but galaxy B must also be the closest to galaxy A. It ensures that any gravitational interactions between the two are not dominated by a nearby third galaxy. To assess the crowdedness of the local region, then, Nottale and Chamaraux introduced an isolation criterion, something defined as follows:

$$\rho = \frac{r_3}{r_p} \tag{A.2}$$

where ρ is the isolation criterion, r_3 is the distance to the nearest galaxy in the M-limited HyperLEDA catalog to the pair, and r_p is the separation of the pair. All galaxy pairs with an isolation criterion $\rho > 2.5$ are kept. Nearby galaxies with $M > -18.5$ are treated slightly differently; again any pair with $\rho_{faint} > 2.5$ is kept, but now if $L_{weakest}/L_{faint}$ >10, the pair will also be kept, $L_{weakest}$ being the luminosity of the fainter pair member. In effect, this allows for the presence of dwarf satellites and similar that would have minimal influence on the gravitational interactions, and allows for the consideration of more pairs.

An interesting comparison can again be struck here with Moreno et al, who calculated an isolation criterion of their own, their proximity parameter, Φ. It is given by $\Phi = m_3/r_3$, where m_3 is the dynamical mass of the third galaxy, and r_3 is the pair-third separation as before. This gives an idea of the third galaxy's gravitational influence on the pair, but does not compare that influence to the pair's on one another, while Nottale and Chamaraux's isolation criterion directly compares the influence of the third galaxy.

A great strength of this catalog is the flexibility of the isolation criterion. Most commonly created are three samples: the 'weakly isolated' pairs, with $\rho > 2.5$; the 'fairly isolated' pairs, with $\rho > 5$; and the 'highly isolated' pairs with $\rho > 10$. These terms will be used further on in this review to refer to different analyses of the IGPC.

Reviewing this catalog, it is clear that it marks a significant advance in the availability of galaxy pair catalogs. The number of pairs - a tenfold increase on what was previously available - is a monumental achievement, facilitated by use of the HyperLEDA database. However, to return the discussion of Fig. A.1, the existence of missing galaxies at the lower magnitudes has a measurable effect on the number of pairs detected. When the distribution of pairs with redshift is

compared with the distribuition of galaxies with redshift, an excess of pairs at high redshift is observed, and this is ascribed to the missing galaxies described in Fig. A.1. Galaxies missing at higher redshifts would result in two effects, as discussed in the paper. First, that pairs could be missed entirely due to a lack of redshift data for one or both pair members. This would result in a deficit of pairs when compared to the distribution of galaxies, but this is not what is observed. Then, the second effect must be stronger, that pairs will be introduced that do not truly have $\rho > 2.5$; while physically there is some third galaxy close enough to invalidate the pair, the lack of redshift data for it excludes it from the analysis. This shows the decreasing efficiency of the isolation criterion at higher redshifts.

A.2.3 Analysis of the Isolated Galaxy Pair Catalog

Nottale and Chamaraux, 2020

At the same time as the catalog was being published, Nottale and Chamaraux published a second paper, outlining their new method of deprojecting galaxy pairs to obtain their full 3D intervelocities and interdistances using the observed data and statistics [Nottale and Chamaraux, 2018b]. When observing galaxies, only three of the six required values can be obtained: the two components of the interdistance that are on the plane of the sky, and the radial component of the intervelocity. However, the methods put forward in this paper allowed the PDFs of full 3D intervelocites and interdistances to be derived using only PDFs of the observed data, and the assumption that the orientation of a galaxy pair is random. The most useful result for this review is the deprojection formula obtained for the intervelocity, given as

$$P_v(v) = -v \left[\frac{dP_{v_z}(v_z)}{dv_z} \right]_v , \qquad \text{(A.3)}$$

where $P_v(v)$ is the PDF of the 3D intervelocity v, and $P_{v_z}(v_z)$ is the PDF of the radial intervelocity v_z.

To assess the effectiveness of their methods, full galaxy pair intervelocities were drawn from a Gaussian distribution, randomly projected, and then deprojected using the new algorithm.

Two years later, Nottale and Chamaraux applied this technique to the IGPC [Nottale and Chamaraux, 2020]. A number of properties of galaxy pairs were investigated in this paper, the investigation being aimed more towards gathering data than answering any specific question. Data on the intervelocities, interdistances, and mass-to-light ratios within the IGPC were among the most notable results; however, their results for the deprojection of the intervelocities has since been by far the most influential.

Deprojection was performed for three separate samples: $\rho > 2.5$, $\rho > 5$, and $\rho > 10$. Furthermore, the $\rho > 2.5$ sample was split again into low errors, with error $\delta v < 20$ km s^{-1}, and higher errors, with $\delta v < 70$ km s^{-1}. The intervelocity PDF was obtained for each of these samples, and showed the same pattern each time: a peak at $\sim$150 km s^{-1}. This was a remarkable result, and one that was fairly new; while there had been hints at such a thing, such as Moreno et al's 2013 paper [Moreno, Bluck, Ellison, et al., 2013] finding a peak at around 200 km s^{-1} for galaxy pairs of a central galaxy and its satellite, there had certainly not been something found for so general a sample. Yet, the peak was present in all subsamples tested. In the highly isolated sample, there were also indications of a second peak at around 350 km s^{-1}, but this did not appear in other subsamples, and so nothing certain could be said of it.

The deprojection was then further carried out on a previous catalog compiled by Nottale and Chamaraux from the Uppsala Galaxy Catalog (UGC), the UGC Galaxy Pair Catalog (UGPC) [Chamaraux and Nottale, 2016], and both the existence and position of the peak were verified. The UGPC has similar pair selection criteria to the IGPC, with the same limits on separation, radial velocity difference, and isolation criterion limit $\rho > 2.5$ with ρ defined as in Equation A.2. Multiplets are excluded as in the IGPC. However, the magnitude limits for the source catalog do differ. The limit for the UGPC is an apparent magnitude limit of $m < 14.5$. Converting the IGPC's absolute magnitude limit of $M < -18.5$ to apparent magnitude a modified version of equation A.1, assuming $v_{pec} = 0$. Using $v = 16500$ km s^{-1}, the recessional velocity limit on the background catalog (this to allow for the pair selection limit of 16 000 km s^{-1} and the radial velocity difference limit of 500 km s^{-}1), the upper limit on distance is therefore to 3 significant figures 244 Mpc. This value can then be used to convert the absolute magnitude limit into an apparent magnitude limit - the resulting value comes out to be $m = 18.4$, suggesting that the IGPC will be able to see fainter galaxies.

The most important takeaway from this paper, then, should be the existence of a peak in

the intervelocities of galaxy pairs. The aim of the study was to gain the broad dynamical properties of the catalog, rather than calculate any specific value; while the peak's position was estimated by eye, there was no fitting. This is not done until a later work by Scarpa et al. [Scarpa, Falomo, and Aldo, 2022b], whose pivotal papers not only obtained a value for the peak position, but found its existence could be predicted using the equations of MOND.

Scarpa et al., 2022

The first of the studies to follow up the discovery of an intervelocity peak was conducted by Scarpa et al. in 2022 [Scarpa, Falomo, and Aldo, 2022a], with its first and clearest advantage being that by the time the study was conducted, the HyperLEDA database had grown by 20%; when it was searched for pairs by Nottale and Chamaraux's selection criteria (detailed earlier), the number found was 25% higher than it had been in Nottale and Chamaraux's 2020 analysis.

Using this catalog, they then repeated Nottale and Chamaraux's deprojection, obtaining similar results with the peak at approximately 150 km s^{-1} - see Fig. A.2. The first deprojection they carried out used the fairly isolated pair subsample, $\rho > 5$, and just like Nottale and Chamaraux they did not observe any 350 km s^{-1} peak for this sample. They did, however, obtain a substantial peak at approximately 500 km s^{-1}. This, however, is explained as being a quirk of the analysis; it is likely formed from the galaxy pairs where though their radial velocity difference meets the selection criterion of $\Delta V < 500 kms^{-1}$, their true 3D intervelocity difference is higher, and so cannot be correctly deprojected. Once they had confirmed Nottale and Chamaraux's analysis with an extended catalog, using a range of bin widths and observing a stable peak for each (see again Fig. A.2), they then went one step further to assess the significance of the peak. To do so, they began with the weakly isolated pairs ($\rho > 2.5$) and selected half of them at random. These were then used to reconstruct the peak, a background value calculated - using bins adjacent to the peak - and then the strength of the peak calculated. The peak was found to be present ∼50% of the time at the 5σ level, and ∼90% at the 3σ level - an astounding result, and a confidence that only improves once the galaxy pairs with radial velocity errors > 70 km s^{-1} were excluded. Scarpa et

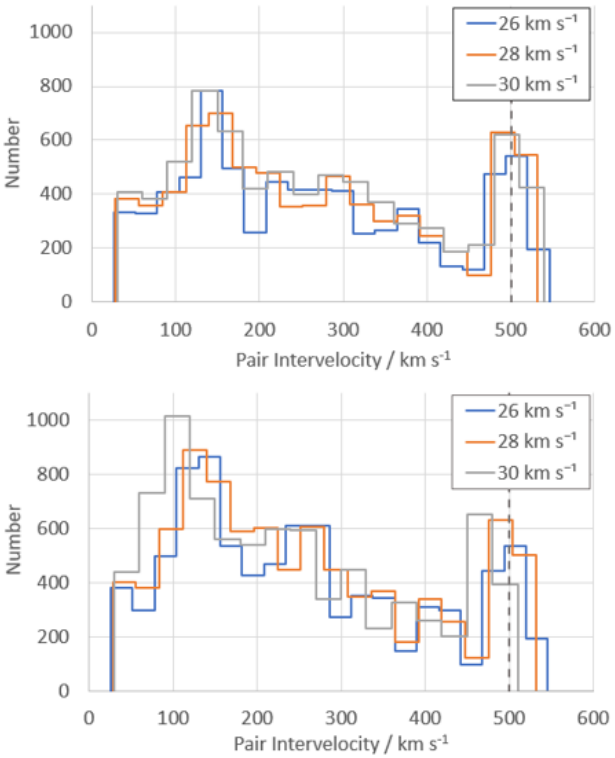

FIGURE A.2: After Scarpa et al [Scarpa, Falomo, and Aldo, 2022a]. *Top:* Scarpa et al.'s deprojection of Nottale and Chamaraux's original catalog, with three different bin widths: 26 km s^{-1}, 28 km s^{-1}, and 30 km s^{-1}. *Bottom:* Deprojection of the larger, updated pair catalog, with three different bins widths. A more prominent peak is present at approximately 150 km s^{-1}, confirming Nottale and Chamaraux's original findings. In both, a grey dashed line represents the limit for the radial velocity difference, and the peak that comes about from this.

al.'s thorough analysis demonstrates beyond reasonable doubt that there is an intervelocity peak present in the IGPC.

However, what this paper is yet to be convinced of is the validity of the IGPC altogether. A section of their paper is dedicated to the identification of false pairs, estimating their contribution and effect. Nottale and Chamaraux originally used the radial velocity distribution, and its asymptote at $\sim$400 km s^{-1}, to estimate the false pair contamination to be approximately 10%. Scarpa et al. strongly disagree with this analysis, and discuss at length different methods that would show it to be much higher.

They calculate false pair contamination in a number of ways, firstly by assuming that the galaxy catalog was distributed uniformly across the sky. As a result of this, for each galaxy examined, there would be on average 2 other galaxies in a radius of 1Mpc. Then, using the redshift range of 13000 km s^{-1} and the radial velocity difference criterion of $\pm$500 km s^{-1}, they calculated a probability of 1/13 that any second galaxy is in the correct redshift range for the examined galaxy to meet the pair criterion. After adjusting for double-counting, they concluded that even a perfectly uniform distribution of galaxies would find 6800 pairs by chance. Even though some of these would be excluded due to the multiplet exclusion criterion, this would not have too strong an effect, and the conclusion was still a lower limit on the contamination level of $\sim$40%. This is a lower limit due to the catalog only in truth covering a fraction of the sky. A second estimate was made by randomly permuting the catalog's velocities, and this randomly produced $\sim$10 000 pairs, so an overall contamination level of $\sim$60%. No matter the method, the conclusions are similar: the that level of contamination is significantly higher than that estimated by Nottale and Chamaraux, and is likely to be close to if not above 50%.

Then, however, the paper becomes a little more controversial. In order to assess galaxy pairs through the lens of ΛCDM, Scarpa et al. here refers back to Moreno et al's 2013 analysis of galaxy pairs in the Millennium Simulation [Moreno, Bluck, Ellison, et al., 2013; Springel, White, Jenkins, et al., 2005]. This has been discussed previously within this paper, but has not yet been given thorough treatment. Scarpa et al. claim that Moreno et al.'s work 'clearly indicates' that there is no precedent for an intervelocity peak in ΛCDM; however, as suggested in Pawlowski et al.'s 2022 rebuttal [Pawlowski et al., 2022], it is not quite so black and white.

The aim of Moreno's paper was to assess the dynamics of galaxy pairs in a cosmological setting - that is, to investigate the effects of a galaxy pair's surroundings on its dynamics, and to assess the assumption that galaxy pairs can be treated in isolation. This is repeatedly referred to as one of the most important things to take away throughout the paper; as such, they often have vastly different selection criteria. Firstly, the upper limit on distance separation for a galaxy pair is significantly lower for Moreno et al. than it is Nottale and Chamaraux, and the studies that followed them. Moreno et al. set their limit at $250h^{-1}$kpc, approximately 30% of the 1Mpc set by Nottale and Chamaraux.

Secondly, there is no exclusion of multiplets in Moreno et al.'s paper, which is perfectly justifiable in light of their paper's aim. However, it does not necessarily map well to the isolated pair catalog produced by Nottale and Chamaraux.

Aside from the selection criteria, it is noted that Moreno et al. are studying a simulation, with minimal effort taken to relate this simulation to observation. Thus, any comparisons with observation must take in the limitations of observation - finite resolution, errors, only able to measure half of the values needed for accurate intervelocities and interdistances, etc. In addition to this, the mass ratios of pairs are substantially smaller in Nottale and Chamaraux's catalog, with Scarpa et al. calculating a median mass ratio of only $\sim$1.2, in contrast with the mass ratios in Moreno's study spanning several orders of magnitude.

The analysis of Moreno et al. shows an interesting approach to galaxy pairs, with the information from the simulation allowing for far more detailed analysis than an observational catalog would allow. The most obvious example of this is the original classification of pair 'flavours', describing five distinct types of galaxy pairs with relation to their dark matter haloes. The most important of these classes, with by far the highest fractional contributions to the overall sample were central-satellite (where the pair consists of a central galaxy and its satellite in the same dark matter halo) making up 52.7% of the sample, and satellite-satellite (where the pair consists of two satellite galaxies of the same central galaxy in the same dark matter halo) making up 46.1%. The remaining 1.2% comes from the other classes, but only the two most significant classes will be discussed here.

When plotting the intervelocity distributions for their galaxy pairs, the satellite-satellite sample showed a peak at 600 km s^{-1}, which is not

likely to be connected to the observed peak. However, the central-satellite sample showed a peak at 200 km s^{-1}, and this is not so drastically different from those peaks observed in Nottale and Chamaraux's catalog. Given the differences in selection criteria, an exact analog would not be expected - therefore, 200 km s^{-1} is not a bad match.

Similarly, further in the analysis, Moreno et al. reclassify their pairs in terms of their potential energies relative to a nearby third galaxy - it is worth noting that these results are not to be taken as law when discussing isolated galaxy pairs as, by definition, there should not be any significant influence from a third galaxy. The systems are classified as one of: pair only, where a bound pair is not bound to a third galaxy; pair dominated, where a bound pair is bound to a third galaxy, but more tightly bound to one another; third dominated, where a bound pair is bound to a third galaxy, more tightly to the third than one another; and third only, where the pair is not bound to one another, but only to the massive third galaxy. It is worth noting that Moreno et al. later define another class, IsoPO, which describes pair only systems that do not have a third galaxy within 10 h^{-1} Mpc. The closest analog to these in Nottale and Chamaraux's catalog would be those pairs for which isolation criterion $\rho > 10$.

When the velocity distributions for these classes are plotted, a surprising result emerges. The pair only, pair dominated and third dominated classes all have an intervelocity peak that falls between 100 km s^{-1} to 200 km s^{-1}, much in line with what is seen in Nottale and Chamaraux's galaxy pairs. However, not only does the third only class have much much higher intervelocities (closer to 1000 km s^{-1}), the IsoPO category also has a much higher peak, at approximately 700 km s^{-1}.

While the result for third only pairs can be said to be unrelated, as two galaxies bound to the same third are not at all the sorts of galaxy pairs one would expect to see in Nottale and Chamaraux's catalog, the result for IsoPO is quite shocking considering what has been observed in the IGPC. This, perhaps, is where Scarpa et al.'s claim has its greatest strength - for the system most similar to the IGPC, its results in Moreno et al's paper do not match at all what has been calculated by Scarpa et al. However, to say there is no evidence for a velocity peak is perhaps a little harsh, as no less than three of the energy classes have peaked in the correct region.

When investigating isolated galaxy pairs, the dynamics of galaxy pairs in crowded cosmological settings are not necessarily all that relevant, due to the significant differences in environment. However, Moreno et al.'s findings on the isolated pair only (IsoPO) class are certainly surprising, and in part motivate the future investigation of the Millennium simulation discussed in Section A.3 - it will determine with certainty whether the disagreement between the IGPC and Moreno et al.'s pairs from the Millennium simulation is simply a matter of selection criteria, or something more fundamentally wrong with the assumptions of the simulation itself.

When considering what might have caused the discrepancy with Moreno et al.'s results, Scarpa et al. look to the cosmology of the simulation to explain. As the ΛCDM simulation cannot reproduce these pairs, they turn to ask whether MOND can. In this first paper, this was carried out with an approximation - a formula for the velocity difference of a galaxy pair with equal masses under circular motion was proposed by Milgrom [Milgrom, 1983c] as

$$\Delta V^4 = 2GM_{tot}a_0 \tag{A.4}$$

where G is the gravitational constant, $M_{tot} = 2M$ the total mass, a_0 is the MOND acceleration constant taken as 1.2×10^{-10} m s^{-1}, and ΔV is the estimated velocity difference. To call back to the earlier discussion of MOND, Scarpa et al. notes the lack of dependence of ΔV on separation - this is the key feature that reproduces the Tully-Fisher relations. This analysis produced a distribution with a peak that was a little higher than that of Nottale and Chamaraux at ∼150 km s^{-1}, with the average being 172 km s^{-1} and a FWHM of 104 km s^{-1}. However, considering both that it is simply an approximation and that there had still not yet been any more attempts to solidly quantify the peak position from the IGPC, this is a strong comparison; certainly enough to warrant further investigation.

And that further investigation was carried out in a second paper by Scarpa et al [Scarpa, Falomo, and Aldo, 2022b]. The most important information left to retrieve was the exact position of the peak in Nottale and Chamaraux's catalog; however, Scarpa et al. note that the peak position varies slightly depending on the precise details of the method of deprojection. The deprojection itself is done in bins, the widths of which can be varied. However, a caveat with the algorithm - specifically, that it requires a monotonically decreasing function - means that bin sizes must be large enough that statistical fluctuations are accounted for. It is not, then, so simple as to just look at the smallest reasonable bin, as might be

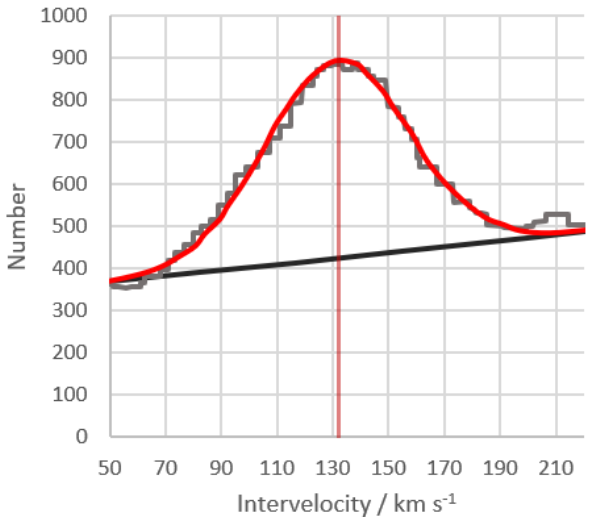

FIGURE A.3: After Scarpa et al. [Scarpa, Falomo, and Aldo, 2022b]. The grey histogram shows the resampled distribution, while the red line shows the fit. This is to date the most careful calculation of the position of the intervelocity peak. The black line shows Scarpa et al.'s assumption of the background.

done in a finite difference method.

Scarpa et al. instead used deprojections with ten different bin widths, ranging from 23 km s^{-1} to 32 km s^{-1}, and averaged these deprojections to create a far smoother distribution - and one that could have a Gaussian fitted to it with minimal deviations around the peak. Using this method, Scarpa et al. found the position of the peak to be 132 ± 5 km s^{-1}, and the fitted Gaussian is shown in Fig A.3. This is a remarkable result, it being the first value for the peak of the observational data to be calculated rather than estimated; it does not, however, fit well with their previous theoretical calculation placing the peak at 172 km s^{-1}.

The remainder of the paper is then dedicated to laying out a more thorough derivation of the expected intervelocity distribution in MOND, the details of which are beyond the scope of this review. What is important is that using this more precise formula, with notable differences including removing the need for the two galaxies to have equal masses and the changing of the prefactor 2 from equation A.4 to 0.610, the predicted peak shifts back for a new peak position of 141 ± 5 km s^{-1}. There is, therefore, a small overlap in the uncertainties in the peak positions for the observed catalog and the MOND predictions; within the bounds of the uncertainties, MOND does predict the observed intervelocity peak.

A.2.4 A ΛCDM Simulation

While there is little denying the success of MOND in describing the peak, there is something to be said as to Scarpa et al.'s claim that ΛCDM frameworks do not also predict one. This is the focus of a 2022 study by Pawlowski et al., whose criticisms of Scarpa's conclusion - that ΛCDM has no evidence of such an intervelocity peak - motivated a study of the intervelocities of galaxy pairs in the TNG300 simulation of the IllustrisTNG project [Nelson, Springel, Pillepich, et al., 2019; Naiman, Pillepich, Springel, et al., 2018a; Springel, Pakmor, Pillepich, et al., 2018; Nelson, Pillepich, Springel, et al., 2018; Naiman, Pillepich, Springel, et al., 2018b; Marinacci, Vogelsberger, Pakmor, et al., 2018]. Due to the physical distancse to the furthest galaxies in the IGPC (see calculation in Section A.2.3) being 244Mpc, the 300Mpc box length of TNG300 is a sensible value, and the choice of a full-physics simulation in contrast with Moreno et al.'s use of the Millennium Simulation (it being the only other major simulation-based result in the field) ensures that baryonic physics can be correctly accounted for.

The most striking thing about Pawlowski et al.'s methods, however, was how they treated the simulation. Simulation data is significantly different from observational as it is complete - every property of every galaxy is known and stored. This difference makes a comparison between the two slightly challenging, as one would expect some error due to the faintest/closest galaxy pairs only being measured by the simulation, but that cannot be detected through observation. Additionally, and particularly relevant for this project, the lack of need for deprojection eliminates any biases or inaccuracies introduced by said deprojection; while Nottale and Chamaraux's random sampling found that a 3D velocity distribution could be fairly well recovered, investigations such as this require comparable data. In an attempt to alleviate some of this difference, Pawlowski et al. used a mock observational approach. An 'observer' was placed on the $x, y, z = 0$ planes in turn, and the projected values for interdistance and radial velocities were then used. Redshift was determined as described in equation A.1, and limits were imposed such that $2500 < v < 16500$ km s^{-1}. This, similarly to the magnitude calculation in section A.2.3, uses Nottale and Chamaraux's pair selection limit of $3000 < v < 16000$

km s^{-1} with a further 500 km s^{-1} on either end to account for the 500 km s^{-1} radial velocity difference limit. This, therefore, will contain all required galaxies for a complete sample of galaxy pairs at the same redshifts as Nottale and Chamaraux, and constitutes the "base" sample.

The "full" sample went further. For this, they first introduced an apparent magnitude cutoff limit at $m_B < 19$. For the distance limit imposed by the redshift, this corresponds to an absolute magnitude limit of $M_B < -18$. They also introduced errors to their velocities, in order to ensure compatibility with the observational catalog - the errors were drawn with replacement from the HyperLEDA catalogue, and used as the standard deviation for a Gaussian. An error with which to displace the velocity was then drawn from each Gaussian.

Two further samples, where the velocities and positions were randomised respectively, were also created, and this was in order to confirm physical origin. This is a similar approach to that Scarpa et al. used to determine false pairs; however, unlike Scarpa et al. these samples are used throughout the analysis.

Pairs were then selected based on the same criteria as Nottale and Chamaraux, outlined in Section A.2.2, and this approach was verified by compared the distributions of projected separations, radial velocity differences and isolation criteria of each sample with those from Nottale and Chamaraux's IGPC. The full and base samples matched the distributions of the observed samples well at most values - it is noted, however, that the simulated galaxy pairs had a far greater abundance at small separations for all three variables than the observed pairs. This is one significant difference that the mock-observation cannot account for - in simulations, the existence of every galaxy is already known, whereas observers may struggle to resolve a close pair. This verification was followed by the full deprojection, the results of which can be seen in Fig. A.4. The blue and red lines represent the base and full samples respectively, and both can be seen to display a peak at similar points to the grey histogram representing the observational data. Fits are also shown in dotted lines, with the peak positions found to be: 88 ± 13 km s^{-1} for the simulated galaxy pairs with $\rho > 2.5$; 125 ± 7 km s^{-1} for $\rho > 5$; and 125 ± 4 km s^{-1} for $\rho > 10$. These are then compared with obtained values for the observed peaks, where: 113 ± 18 km s^{-1} for $\rho > 2.5$; 113 ± 5 km s^{-1} for $\rho > 5$; and 132 ± 13 km s^{-1} for $\rho > 10$. Each of these shows some overlap within their uncertainty with their respective TNG300 predictions - quite the promising result. This alone is enough to show that the peak is also present in ΛCDM, but Pawlowski et al. went further.

First, they observed that the peaks in both the base and full samples are significantly more prominent than those in the observed data. To investigate this, they created a new sample, identical to the full sample but this time with the errors doubled - this is shown in Fig A.4 as the teal lines. Given that the teal lines show much better agreement with the observed data than the full and base samples, it would be reasonable to suggest that the errors in the observational study could have been underestimated.

The most important addition to Fig A.4, however, was the red histogram. It shows the 3D intervelocities as they were taken directly from the full sample, without being deprojected. Not only did this confirm the deprojection analysis in the similarities between the red line and the red histogram, but it also confirms that an intervelocity peak is indeed present in ΛCDM, found in the 3D intervelocities for the full sample ($\rho > 2.5$) at 132 ± 1 km s^{-1}, even if its physical origin is not yet understood. While this value is noticeably higher than the observed peak for $\rho > 2.5$ in Pawlowski et al.'s analysis, due to the high uncertainty in the observed peak the uncertainties still overlap.

The remainder of the paper is dedicated to investigating the nature of the peaks, with two key results shown in

First, the randomised samples were used to confirm that the appearance of the peak is indeed a physical effect. A broad peak at approximately 500 km s^{-1} is present for the intervelocities of both randomised samples (see Fig. A.5), with no evidence of any peak close to 150 km s^{-1} as would be expected if it were simply a random effect. It shows that the peak emerges due to some relationship between individual galaxy pairs. Notably, a peak at around 500 km s^{-1} was also present in Scarpa et al.'s analysis (see Fig. A.2); this would further suggest that the velocity limit of 500 km s^{-1} results in an artificial peak.

A particularly useful piece of analysis pertains to the debate over the influence of false pairs - while Nottale and Chamaraux estimated it to approximately 10%, Scarpa et al. estimated it to be at least 40%, and likely as high as 60%. The simulation approach has, within reason, matched all observations - and any discrepancies have been

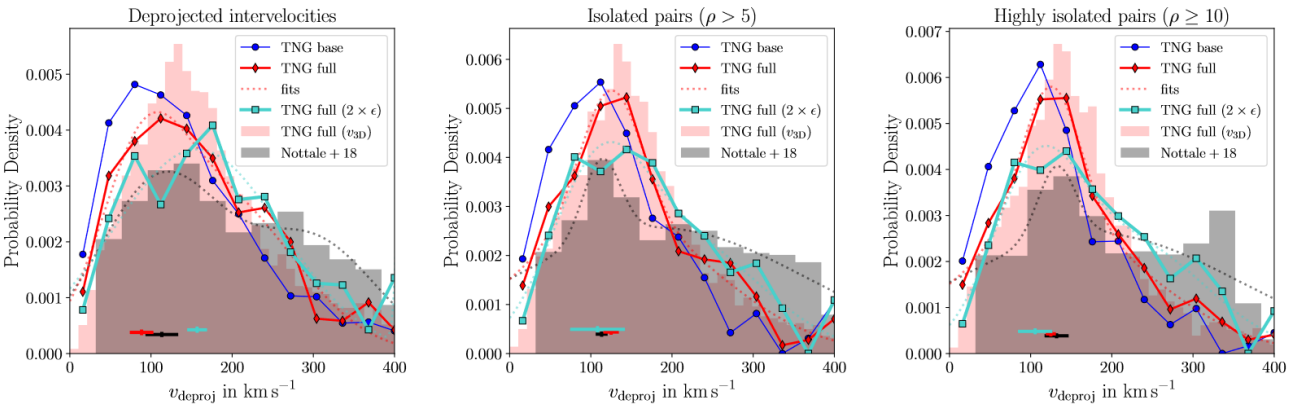

FIGURE A.4: Taken from Pawlowski et al. [Pawlowski et al., 2022] with permission. PDFs for the intervelocity distributions are shown for three different ranges of isolation criterion, similarly to the considerations in Nottale and Chamaraux's analysis (see Section A.2.3). An intervelocity peak is observed at all three scales, for all samples shown. The blue lines show the base samples, and the red the full. The grey histogram shows the observed data - and it can be seen that the base and full samples show a similar peak to this regardless of isolation criterion. Interesting samples are shown in teal and in the red histogram. The teal line shows a second version of the full sample, where the errors the velocities are displaced by are doubled. The red histogram shows the full 3D intervelocities as taken from the simulation, without deprojection - and it is critical to note that peaks are also present in each of these.

easily explained. Pawlowski et al. investigated the effect of spurious pairs by using the simulation to find out whether or not each pair shared a halo ID, using this as an indicator of physical association. The interdistance of these pairs, shown in Fig. A.5, demonstrates just how apt this assessment is - there is a clear distinction between the interdistances of galaxies in the same halo, that are predominantly very low, and galaxies in different halos, that are almost entirely greater than 500 kpc and often greater than the projected interdistance limit, 1 Mpc. It is highly likely that this are spurious pairs, with no real physical association. It is interesting, however, that when the intervelocities for shared and different halo IDs were plotted, peaks were present for both. The pairs with different halo IDs had an intervelocity peak at 88 ± 2 km s^{-1}, something Pawlowski et al. ascribes to their greater average separation. The pairs with shared halo IDs had a peak at 138 ± 1 km s^{1}, far closer to the predicted intervelocity peaks.

Finally, the comparison of how intervelocity peak varies with absolute magnitude, also seen in Fig. A.5. This is last to be discussed as it is particularly relevant for Section A.3, shortly ahead. Pawlowski et al. calculated peak positions for different bins of absolute magnitude of the brighter galaxy, and found that increasingly luminious galaxies led to increasingly higher intervelocities. The dimmest band, with $-17 > M_1 > -18.5$ had an intervelocity peak of 86 ± 1 km s^{-1}; the middle band $-18.5 > M_1 > -20$ had a peak at 117 ± 1 km s^{-1}; and the brightest band $-20 > M_1$ peaked at 153 ± 1 km s^{-1}. This shows a clear pattern, though one that is fairly reasonable; brighter galaxies will on average have higher masses, and so will again on average cause higher accelerations and velocities. Extending this pattern down to dimmer and dimmer galaxy pairs is a primary object of the project proposed in Section A.3. However, there is one significant criticism of Pawlowski et al.'s study, and that is to do with the simulation's assumptions. It is apparent that the peak is of physical origin, with the strongest piece of evidence for that being that it moves with luminosity, and so by extension, moves by mass. Every simulation must be tuned to reproduce something, and TNG300 is tuned to, amongst other things, reproduce the stellar mass halo mass (SMHM) relation through abundance matching. However, this technique is consistent with the baryonic Tully-Fisher relation, something that has been discussed at length in A.1.2 as a clear victory of MOND. The criticism is that by tuning for the SMHM relation, the simulation has mimicked MOND before it runs - that the MOND effect has been baked into the code, as it were, and so the observed peak is not truly a victory of ΛCDM but simply a coincidence based on tuning. To investigate this

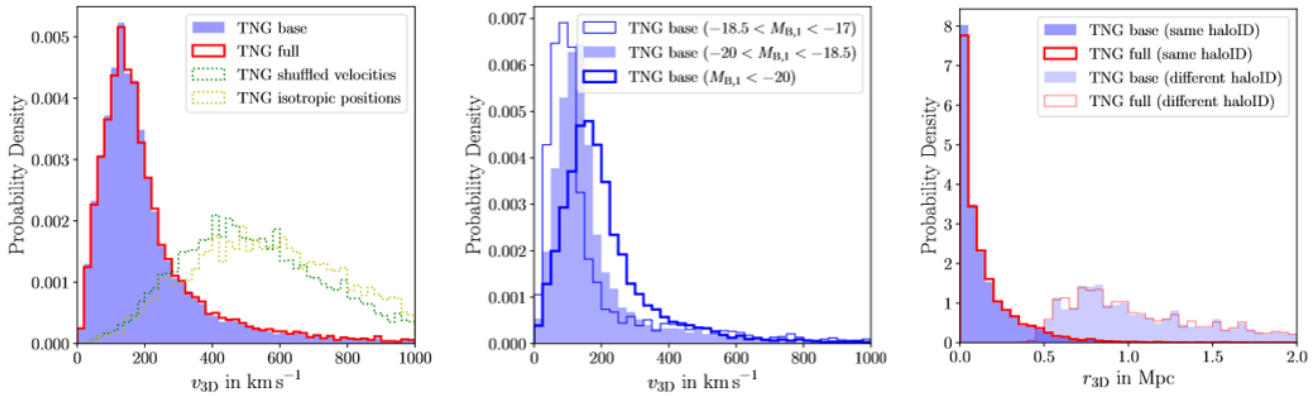

FIGURE A.5: Taken from Pawlowski et al. with permission. *Left:* The intervelocity distribution of the base and full samples compared with the randomised samples. Note the broad peak in both randomised samples, with a peak position of approximately 500 km s^{-1}. *Centre:* The variation of peak position with absolute magnitude of the brighter galaxy. There is an apparent shift in the peaks to higher preferred intervelocities as absolute magnitude increases, and this is likely connected to their (on average) increasing mass. This trend, however, will be the subject of future study, as discussed in Section A.3. *Right:* The full interdistances of galaxy pairs with the same and different halo IDs. Only those with the same halo ID have low interdistances; conversely, those without have far higher interdistances. Those that do not share a halo ID are likely spurious pairs.

criticism, the proposed project discussed in Section A.3 also considers investigations of galaxy pairs with different simulations with different assumptions. These, then, could prove or disprove the claims that the TNG300 simulation has already accounted for MOND.

A.3 Future Work

Thus far, it has become evident that a peak is present in the intervelocity of galaxy pairs, in observation and in both theories of ΛCDM and MOND. The position of this peak is agreed upon in all three cases, within the bounds of the uncertainty. To discriminate between the models, testable predictions must be made for ΛCDM that significantly differ from those made in MOND. Beyond that, MOND still sees stronger success than ΛCDM in explaining the physical origin of the peak; while ΛCDM clearly predicts it, and Pawlowski et al. have made a strong argument relating its position to the absolute magnitudes of the component galaxies, its origin is not nearly so neatly explained as in MOND.

Therefore, future work must focus on these two areas. It may also be useful to test for the presence of this intervelocity peak in different ΛCDM simulations, due to the criticisms of Pawlowski et al. discussed in Section A.2.4. One simulation of interest would be the Millennium simulation, to investigate whether the application of Nottale and Chamaraux's selection criteria makes the peak evident there as well, or whether the assessment of Scarpa et al. was correct in stating that the Millennium simulation does not show an intervelocity peak in this way. In that case, the simulations would need to be directly contrasted in order to explain any discrepancies. However, given that it is highly likely that the differences in results between the Millennium simulation and TNG-300 are due in the most part to the vast differences in selection criteria, what is far more important first is to broaden understanding of the intervelocity peak as it has been obtained - in TNG-300.

For testable predictions for the intervelocity peak in ΛCDM, it would be most prudent to focus on fainter galaxy pairs. Pawlowski et al's investigation of the position of the intervelocity peak for different magnitude ranges suggested that the peak may move as the absolute magnitude of the galaxies increases. By increasing the magnitude cutoff point for analysis of the simulation, this can be more thoroughly investigated, and a plot of intervelocity peak against absolute magnitude produced. This trend, whatever it may be, could then be compared to future observations of fainter and fainter galaxies; it may also be interesting to see what predictions would be made in a MOND system, and to compare the three plots much as has been done with the peaks themselves.

If MOND and ΛCDM make mutually exclusive predictions, this could then become a test of the paradigms, beyond simply strengthening

ΛCDM's ability to describe the peak. It is also possible that relationship might emerge that sheds some light on the physical origin of this peak in a ΛCDM context, and this should be kept in mind throughout the investigation - however, given that variations with fainter galaxies are the most easily testable against future observations, the predictions will be the most important result.

To gain insight into the physical origin of the peak, it would be far more appropriate to observe the dynamics of galaxy pairs throughout their evolution. One of the greatest benefits of studying simulations is the ability to access 'snapshots', freeze-frames of the simulation throughout time. Using these, it becomes possible to investigate the intervelocity peak through the evolution of the simulation, searching for how it originated and how it has evolved through time. Beyond this, use of the snapshots would also make it possible to track specific galaxy pairs back through their evolution using their halo IDs. By observing the development and environments of galaxy pairs both in and outside the peak, it may be possible to gain insight into the types of environments that do/do not result in intervelocities around the 135 ± 4 km s^{-1} peak. This may then provide clues as to how the peak has come about in the full context of the simulation.

The investigations will be carried out in order of priority, with the magnitude investigation being undertaken first. As this will provide results that will be possible to test against observations in the near future, this is by far the most promising route to meaningful predictions. Should some pattern come of this, the nature of this pattern will be investigated - its stability, its statistical significance, and perhaps even its origin. The nature of these further tests will be designed with advice from colleagues, and will likely depend on the nature of the pattern that arises.

Should there be no significant pattern, it will instead be followed by the redshift investigation - its results will yield the same crossroads as magnitude. The completion of these studies will mark a pivotal point in the project, at which point the direction of the project will need to be discussed with supervisors and a new plan made. This plan will likely be either specialised investigations following patterns in the magnitude or redshift results, or identifying different simulations to repeat the studies within with the aim of verifying our results. However, given the significant dependency on the results obtained in the first sections, little can be said with certainty about this period.

The studies will be undertaken using modified versions of Pawlowski et al.'s code, and data from the TNG-300 simulation. Should any notable results be produced, such as significant predictions that can be tested against observation, it is of the utmost importance that these are replicated using analysis of a different simulation; this is to ensure that the criticisms of TNG-300, such as those faced by Pawlowski et al. (see earlier discussion), are properly addressed. If multiple ΛCDM simulations predict one pattern, it strengthens the prediction, and by extension any comparisons with future observations.

It should be noted, however, that no observational work has been specifically planned in tandem with this paper. Thus, predictions made will have to wait for the results of future observational studies to be confirmed or contradicted; even should this study produce a prediction for ΛCDM that differs from MOND, there can be no conclusions drawn as there will be no observations to ground them in. This will require future work. Beyond that, while it would be a significant achievement to find differing predictions between ΛCDM and MOND, it may be that no such predictions can be found in this context - in this case, it may be interesting to explore the possibility of further tests probing the depths of the agreement, and potential causes for said agreement. Again in this case, investigations with different simulations may be enlightening.

A.4 Discussions and Conclusions

The intervelocity peak evident in galaxy pairs is a new phenomenon with significant potential for discrimination between gravitational models.

We have reviewed the general debate between MOND and ΛCDM, and followed the discovery and development of understanding of an apparent peak in the intervelocities of galaxy pairs, found first through observation by Nottale and Chamaraux, then firmly located by Scarpa et. al. at 132 ± 5 km s^{-1}, and the debate between the two paradigms as to which can better describe it.

Of the two, ΛCDM has made a wider range of predictions and calculations as to how the peak behaves in different circumstances. Pawlowski et al.'s paper locates the peak for different isolation criteria, different magnitudes, and even uses the simulation to identify physical association and thereby the contamination of the samples by false pairs. Scarpa et al. only locate one peak, for isolation criterion $\rho > 5$, and though they too

give thorough consideration to the level of contamination, it is not so decisive as Pawlowski et al.'s due to the extensive amount of data available through a simulation.

This does not, however, change that ΛCDM's predictions only come from a simulation, with little understanding as to the precise physical origins of the peak beyond having some relation to the mass. Nor does it change the criticisms of Pawlowski et al.'s study relating to the assumptions of the simulation. MOND's predictions of the peak coming directly from the equations themselves is a significant success, and certainly not one that should be overlooked.

As things stand, there is no clear victor in the galaxy pair debate. A project has been proposed that aims to strengthen ΛCDM's claim by investigating the origins of the peak, using its variations with absolute magnitude and redshift as probes into its nature and history, and investigating a wider range of simulations in order to assess the validity of the criticisms against Pawlowski et al. Eventually, however, the part of theory is only half of the information required; any predictions made by this upcoming project will require far more detailed observational data than is current available to draw any conclusions with confidence.

A.5 Acknowledgements

I would like to thank my supervisor for their insightful critiques of this review, and the numerous helpful discussions that have helped me understand the work in this field.

Bibliography

Angus, Garry W and Stacy S McGaugh (2008). In: *MNRAS* 383, pp. 417–423.

Banik, Indranil and Hongsheng Zhao (2022). In: *Symmetry* 14(7), p. 1331.

Banik, Indranil et al. (2024). In: *MNRAS* 527, pp. 4573–4615.

Begeman, KG, AH Broeils, and RH Sanders (1991). In: *MNRAS* 249, pp. 523–537.

Bekenstein, J. and M. Milgrom (1984). In: *ApJ* 286, pp. 7–14. DOI: 10.1086/162570.

Bekenstein, Jacob D. (1984). In: *ApJ* 286, pp. 7–14.

Bertone, Gianfranco and Dan Hooper (2018). In: *Rev. Mod. Phys.* 90(4), p. 045002.

Chae, Kyu-Hyun (2023). In: *ApJ* 952.2, p. 128.

Chamaraux, Pierre and Laurent Nottale (2016). In: *Astrophys. Bull.* 71, pp. 270–278.

Cirelli, Marco, Alessandro Strumia, and Jure Zupan (2024). arXiv: 2406.01705 [hep-ph]. URL: https://arxiv.org/abs/2406.01705.

Clowe, Douglas, Anthony Gonzalez, and Maxim Markevitch (2004). In: *ApJ* 604, pp. 596–603.

Einasto, Jaan, Ants Kaasik, and Enn Saar (1974). In: *Nature* 250, pp. 309–310.

Einstein, Albert (1916). In: *Annalen der Physik* 354(7), pp. 769–822.

Famaey, Benoît and Stacy S. McGaugh (2012). In: *Living Rev. Relativ.* 15(1), p. 10.

Hernandez, X (2023). In: *MNRAS* 525.1, pp. 1401–1415.

Karachentsev, ID (1972). In: *Soobshcheniya Spetsial'noj Astrofizicheskoj Observatorii* 7, pp. 1–92.

Karachentsev, ID and DI Makarov (2008). In: *Astrophys. Bull.* 63, pp. 299–345.

Makarov, Dmitry et al. (2014). In: *A&A* 570, A13.

Marinacci, Federico, Mark Vogelsberger, Rüdiger Pakmor, et al. (2018). In: *MNRAS* 480(4), pp. 5113–5139.

McGaugh, Stacy S (2012). In: *AJ* 143, p. 40.

Milgrom, Mordehai (1983a). In: *ApJ* 270, pp. 365–370.

— (1983b). In: *ApJ* 270, pp. 371–383.

— (1983c). In: *ApJ* 270, pp. 384–389.

— (July 1994). "Modified Dynamics Predictions Agree with Observations of the H i Kinematics in Faint Dwarf Galaxies Contrary to the Conclusions of Lo, Sargent, and Young". In: *ApJ* 429, p. 540. DOI: 10.1086/174341. arXiv: astro-ph/9311031 [astro-ph].

Mistele, Tobias et al. (2024). In: *ApJ (Letters)* 969, p. L3.

Moreno, Jorge, Asa F. L. Bluck, Sara L. Ellison, et al. (2013). In: *MNRAS* 436(2), pp. 1765–1786.

Naiman, Jill P., Annalisa Pillepich, Volker Springel, et al. (2018a). In: *MNRAS* 475(1), pp. 648–675.

— (2018b). In: *MNRAS* 477(1), pp. 1206–1224.

Nelson, Dylan, Annalisa Pillepich, Volker Springel, et al. (2018). In: *MNRAS* 475(1), pp. 624–647.

Nelson, Dylan, Volker Springel, Annalisa Pillepich, et al. (2019). In: *Comput. Astrophys. Cosmol.* 6, p. 2.

Nottale, Laurent and Pierre Chamaraux (2018b). In: *A&A* 614, A45.

— (2018a). In: *Astrophys. Bull.* 73(3), pp. 310–317.

— (2020). In: *A&A* 641, A115.

Ostriker, Jeremiah P, Philip JE Peebles, and Amos Yahil (1974). In: *ApJ (Letters)* 193, pp. L1–L4.

Pawlowski, Marcel S. et al. (2022). In: *A&A* 664, p. L6.

Roberts, MS (1976). In: *Comments on Astrophysics, Vol. 6, p. 105* 6, p. 105.

Rubin, V. C., Jr. Ford W. K., and N. Thonnard (1978). In: *ApJ (Letters)* 225, pp. L107–L111.

Ryman, Barbara (2016). *Introduction to Cosmology*. Cambridge University Press.

Scarpa, Riccardo, Renato Falomo, and Treves Aldo (2022a). In: *MNRAS* 510(2), pp. 2167–2172.

— (2022b). In: *MNRAS* 512(1), pp. 544–547.

Schechter, Paul (1976). In: *ApJ* 203, pp. 297–306.

Springel, Volker, Rüdiger Pakmor, Annalisa Pillepich, et al. (2018). In: *MNRAS* 475 (1), pp. 676–698.

Springel, Volker, Simon D. M. White, Adrian Jenkins, et al. (2005). In: *Nature* 435, pp. 629–636.

Will, Clifford M. (2014). In: *Living Rev. Relativ.* 17(1), p. 4.

Zwicky, Fritz (1933). In: *Helv. Phys. Acta* 6, pp. 110–127.