The Lopsided Distribution of Satellite Galaxies

Noam I. Libeskind, Quan Guo, Elmo Tempel, Rodrigo Ibata

The Astrophysical Journal, 830:121 (11pp), 2016 October 20 doi:10.3847/0004-637X/830/2/121

THE LOPSIDED DISTRIBUTION OF SATELLITE GALAXIES

Noam I. Libeskind[1], Quan Guo[1], Elmo Tempel[2], and Rodrigo Ibata[3]
[1] Leibniz-Institut für Astrophysik, Potsdam, An der Sternwarte 16, D-14482 Potsdam, Germany
[2] Tartu Observatory, Observatooriumi 1, 61602 Tõravere, Estonia
[3] Observatoire astronomique de Strasbourg, Université de Strasbourg, CNRS, UMR 7550, 11 rue de l'Université, F-67000 Strasbourg, France
Received 2015 November 10; revised 2016 May 11; accepted 2016 June 2; published 2016 October 17

ABSTRACT

The distribution of smaller satellite galaxies around large central galaxies has attracted attention because peculiar spatial and kinematic configurations have been detected in some systems. A particularly striking example of such behavior is seen in the satellite system of the Andromeda galaxy, where around 80% are on the near side of that galaxy, facing the Milky Way. Motivated by this departure from anisotropy, we examined the spatial distribution of satellites around pairs of galaxies in the Sloan Digital Sky Survey. By stacking tens of thousands of satellites around galaxy pairs, we found that satellites tend to bulge toward the other central galaxy, preferably occupying the space between the pair, rather than being spherically or axis-symmetrically distributed around each host. The bulging is a function of the opening angle examined and is fairly strong—there are up to ∼10% more satellites in the space between the pair than expected from uniform. Consequently, it is a statistically very strong signal, being inconsistent with a uniform distribution at the 5σ level. The possibility that the observed signal is the result of the overlap of two halos with extended satellite distributions is ruled out by testing this hypothesis by performing the same tests on isolated galaxies (and their satellites) artificially placed at similar separations. These findings highlight the unrelaxed and interacting nature of galaxies in pairs.

Key words: cosmology: observations – galaxies: dwarf – galaxies: general – galaxies: halos – Local Group

1. INTRODUCTION

The current paradigm of structure and galaxy formation, known as the ΛCDM model, asserts that the early universe started as a near-homogenous density field, seeded by inflation with small density fluctuations. These small primordial density fluctuations constitute a Gaussian random field whose statistics were famously first computed by Bardeen et al. (1986). The Gaussian nature of the overdensities implies that they are, to first order, modeled as spherically symmetric perturbations. However, as structure forms due to gravitational instability, the initial anisotropy of the fluctuations becomes exacerbated such that by the present day, the dark matter halos that form out of this process are fairly aspherical. Halos of mass $10^{12}\,M_{\odot}$ have axis ratios of around 0.6 (the axis ratio depends on many factors, including mass, redshift, and cosmology; see Allgood et al. 2006).

The shape of dark matter halos can be probed by examining the location of satellite galaxies that inhabit them. It is a fairly well established empirical finding that satellite galaxies tend to not be isotropically distributed around their hosts. The "flattening" of satellite distributions is seen both by large surveys (i.e., in the 2dF or the Sloan Digital Sky Survey [SDSS]; Sales & Lambas 2004, 2009; Brainerd 2005; Azzaro et al. 2007; Agustsson & Brainerd 2010, see also Zaritsky et al. 1997) and in the local universe: satellites of the Milky Way (e.g., Lynden-Bell 1976; Majewski 1994; Kroupa et al. 2005, to name a few), our Local Group partner M31 (Conn et al. 2013; Ibata et al. 2013; Pawlowski & Kroupa 2013), and Centaurus A (Tully et al. 2015) all exhibit a pronounced flattening.

Although the prevailing ΛCDM paradigm of structure formation does not obviously reproduce such extreme behavior (Metz et al. 2008; Pawlowski et al. 2014), a number of studies have been able to accommodate such setups by relying on either the triaxial nature of host halos or large-scale filaments (e.g., Libeskind et al. 2005, 2009; Zentner et al. 2005; Wang et al. 2013; Buck et al. 2015; Cautun et al. 2015; Gillet et al. 2015; Phillips et al. 2015; Sawala et al. 2016). However, debate still exists on whether the simulations are being appropriately compared with the observations. Other issues such as the stability of satellite planes and their formation mechanisms or how satellite galaxy planes relate to other well-known problems of low-mass objects (such as the "missing satellite" [Klypin et al. 1999] or the "too big to fail" [Boylan-Kolchin et al. 2012] problem) imply that it is clearly important to scrutinize the orientation of satellite galaxies around hosts.

The anisotropic distribution of satellites is often couched in terms of two-dimensional and implicitly axisymmetric systems. Yet in the case of M31, at least 21 out of 27 galaxies in the region covered by the PAndAS survey (Conn et al. 2013) appear to be on the near side of M31, preferentially occupying the space in between the two main galaxies. Therefore, not only are ∼50% of M31's satellites confined to a thin, apparently spinning disk of satellites, but ∼80% constitute a lopsided, highly azimuthally axis-asymmetric distribution. Quantifying such a lopsidedness around the Milky Way is complicated not only by the zone of obscuration due to the Galactic disk but also by the lack of complete sky coverage for ultrafaint dwarfs discovered by the SDSS. Despite the limitations and bias, if we consider just the "classical" 11 brightest Milky Way satellites, around 4 of them are on the side toward Andromeda.

It is not clear whether such a lopsidedness is commonplace among galaxies beyond the Local Group. This issue has motivated the work that follows. Do satellites of pairs beyond the Local Group show a statistically significant tendency to have a lopsided distribution with respect to the geometry of their hosts? A hint comes from work such as Faltenbacher et al. (2008), which examined the tendency of satellites, albeit in *N*-body simulations, to be closer to or further from principal directions of the matter distribution. Yet such studies implicitly

gloss over the issue of lopsidedness, because they usually examine angles in the range (0°, 90°), instead of (0°, 180°). Regardless of this important subtlety, studies that find alignments between neighboring groups' shapes often consider objects orders of magnitude more massive than the Local Group, over larger scales in distance as well. Here we confine ourselves to galaxy pairs that roughly mimic the Local Group and quantify the tendency for satellites to fill the space between them.

The location of satellites around central galaxies—even in binary pairs like the Local Group—is surely related to the distribution of matter on larger scales, namely, filamentary accretion along the so-called "cosmic web" (Bond et al. 1996). On satellite scales, much work has focused on mainly two different types of anisotropic alignments: those with respect to the cosmic web (i.e., Paz et al. 2011; Lee et al. 2014; Lee & Choi 2015; Libeskind et al. 2015a; Tempel et al. 2015), and those with respect to the central galaxy (i.e., Yang et al. 2006; Wang et al. 2008, 2014). Since in both of these cases the environment is often assumed to be the main driver of alignments, in this work we examine situations where the tidal field is dominated by a close partner. In order to isolate the effects of an overlapping satellite distribution, we compare our results with a carefully curated overlap sample.

2. DATA AND METHODS

In this section we explain how pairs of galaxies and their satellites are selected from the SDSS DR10 (York et al. 2000; Ahn et al. 2014) for our analysis. We select two samples: an "observational sample" used to examine the hypothesis of this paper, and an "overlap sample" used to examine the effect of overlapping satellite distributions.

2.1. Observational Sample

The first step to examine azimuthal asymmetry in satellite galaxy distributions is to identify pairs of galaxies in the sample. This is done by using the group finding method published by Tempel et al. (2014c). Essentially, a Friends-of-Friends (FoF) algorithm with variable linking length in the radial and transverse direction to account for redshift space distortions is applied to the galaxy distribution. Only galaxies in the spectroscopic sample are considered for pair selection. We adopt cosmological parameters consistent with Planck Collaboration et al. (2014) cosmology, the Hubble constant $H_0 = 67.8\ \mathrm{km\ s^{-1}\ Mpc^{-1}}$, the matter density $\Omega_\mathrm{m} = 0.308$, and the dark energy density $\Omega_\Lambda = 0.692$. Groups of only two galaxies are retained for further analysis, as these constitute a sample of galaxy pairs. With this method we obtain a set of 48,917 two-galaxy pairs, which we subject to further cuts, outlined below. The preparation and details of the spectroscopic galaxy sample are described in Tempel et al. (2014c).

We select pairs where each member has a magnitude in the r band between $(-23.5 < M_r^{0.0} < -21.5)$ as our fiducial magnitude cut (this cut is examined in more detail in Section 3.2.1). This magnitude range is motivated by estimates of the MW and M31's r-band magnitude from Liu et al. (2011), which, although not identical (-21.2), is not too far off either. The magnitude range is purposefully kept wide in the interest of not prejudging our sample. Such that only binary systems where the two galaxies have similar magnitudes are considered, a further cut is performed and only pairs that are within a magnitude difference of each other are retained.

In order to avoid close pairs of galaxies, the two primaries are required to be separated by $0.5 < d_\mathrm{sep}/\mathrm{Mpc} < 1.0$ (in projection). The mean redshift of the two primaries is taken as the distance to the pair, and it is at this redshift that distances are computed. Changing d_sep has little effect since it is difficult to get groups of just two galaxies when d_sep is increased (few pairs have larger values of d_sep). The value of d_sep is considered in more detail in Section 3.2.4.

Once these cuts have been applied to the host, a (projected) search radius within which satellites are retained is examined. The fiducial search radius is set to $r_\mathrm{search} = 250$ kpc, roughly equal to the virial radius of a Milky-Way-mass galaxy halo. The search radius is examined in more detail in Section 3.2.3.

The potential satellite galaxies are searched from the photometric catalog of the SDSS. Since galaxy pairs are selected from the spectroscopic redshift catalog and by construction they do not have other spectroscopic galaxies - nearby, the potential satellite galaxies do not have measured (spectroscopic) redshifts. Because of this, the selection is done only in projection; hence, not all of these will be satellites in the strict sense (namely, within a virial radius). Indeed, many have poorly constrained photometric redshifts that are inconsistent with their alleged hosts. However, any interlopers are expected to be homogeneously distributed and therefore cannot dominate the signal. Therefore, any signal obtained ignoring the photometric redshifts of the potential satellites should be seen as a lower limit.

In sum, there are four criteria that determine which hosts and satellites constitute the fiducial sample:

1. host magnitude: $-21.5 < M_r < -23.5$;;
2. host separation $0.5 < d_\mathrm{sep}/\mathrm{Mpc} < 1.0$;;
3. magnitude difference of pair <1 mag;
4. satellites within a (projected) search radius: $r_\mathrm{search} = 250$ kpc;
5. no consideration of satellite photometric redshift.

In Section 3.2, a section devoted to systematic effects, we examine how the signal changes if these fiducial values are altered.

Our final fiducial sample consists of 12,210 pairs of galaxies, which among them have 46,043 potential satellites. In Figure 1(a) we show the redshift difference of the pairs included in the observational sample (namely, $\Delta z = |z_\mathrm{P1} - z_\mathrm{P2}|$, where z_P1 and z_P2 are the redshifts of the brighter [P1] and fainter [P2] primary). These pairs tend to be close in redshift (by construction) with redshift differences of $\Delta z < 0.003$. Figure 1(b) shows the redshift distribution of the mean redshift of the pair, while Figure 1(c) shows how many satellites are found per host. The median satellite number per host is three, while the mode of the distribution is two satellites.

2.2. Overlap Sample

An immediate area of concern regarding the effect considered in this paper is contamination due to extended overlapping satellite distributions. Suppose that two galaxies have accompanying satellite systems that extend well beyond the 250 kpc considered here. Consider placing these two galactic systems at the separations considered here, namely, between $500\ \mathrm{kpc} < d_\mathrm{sep} < 1000$ kpc. In such a situation one would expect that the projected overlap of the satellite systems

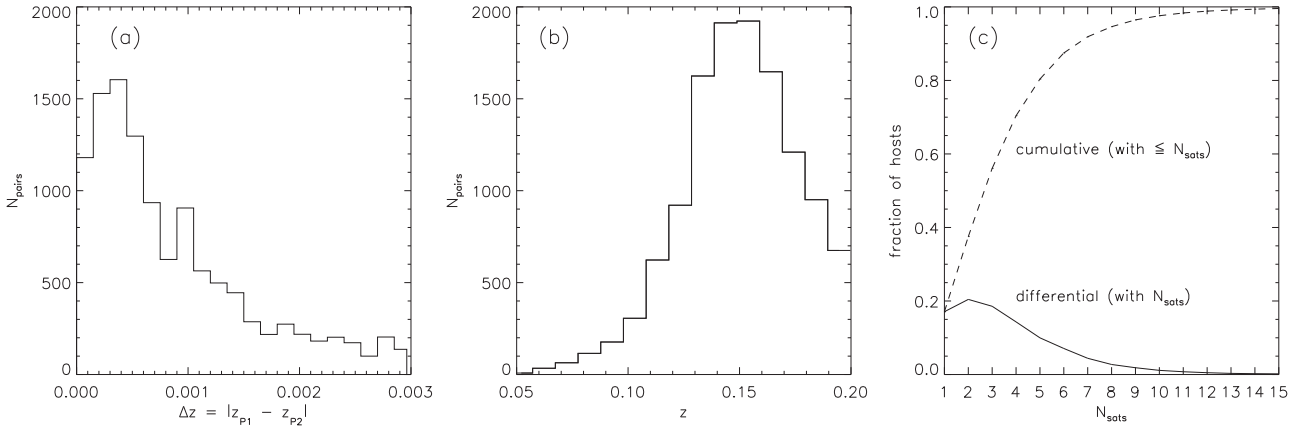

Figure 1. Distribution of redshift differences $z_{P1} - z_{P2} = \Delta z$ (left panel) and of redshifts z (middle panel) for the pairs of galaxies considered here (in the observational sample). By construction these are identical to the overlap sample. Right panel: number of satellite galaxies around hosts. The fraction of hosts that have N_{sats} is shown cumulatively (dashed line) and differentially (solid line).

of these galaxies alone would result in a similar signal to the bulging signature we wish to isolate.[4] In order to control for this overlap effect, we construct an "overlap sample." In practice, we wish to replace each pair member from our observational sample with an isolated "doppelgänger" galaxy (and satellites).

Constructing a sample of overlapping pairs begins with a list of isolated primaries drawn from the spectroscopic sample of the SDSS. "Isolated" galaxies are defined as those FoF groups (as described above; see, e.g., Tempel et al. 2014c) composed of just one single galaxy. All satellite galaxies, within a projected distance of 2.5 Mpc of each isolated primary, are identified from the photometric survey.

For each pair in the observational sample, we select five pieces of data: the redshifts (z_{P1} and z_{P2}), the magnitudes ($M_{r,P1}$ and $M_{r,P2}$), and the separation (d_{sep}) of the two primaries. For each primary (Pi, where $i \in \{1, 2\}$) in each pair in the observational sample, we wish to find an isolated system whose redshift (z_{Di}, D for "doppelgänger") and magnitude ($M_{r,Di}$) are as close as possible to z_{Pi} and $M_{r,Pi}$. We specifically force $|z_{Pi} - z_{Di}| = \Delta z < 0.001$ and $|M_{r,Pi} - M_{r,Di}| = \Delta M_r < 0.05$, where $i \in \{1, 2\}$ represent the two members of the pair.

Once two isolated primaries have been identified as the pair's doppelgängers, these galaxies (and all their satellites within 2.5 Mpc) are given a projected separation equal to d_{sep}. Note that the sense of orientation of their two satellite systems is completely arbitrary. Therefore, in order to sample just the effect of extended radial distributions, the angular position of each satellite is randomized 100 times while keeping its radius with respect to its host fixed. In this way, the effect of the radial distribution of each pair member's satellite system and the strength of any overlap effect can be quantified, isolated, and subtracted from the observed signal. Note that we could have randomized the angular position of the satellites an infinite number of times, since the overlap sample just measures the deviation from a uniform signal. In principle, if we measure the statistical significance for the observed sample, we should compare it with the overlap sample, not with the uniform sample. In practice, it is equivalent to drawing the error bars around the overlap sample's curve and not around the uniform signal (i.e., unity). From a practical point of view, this does not change our results or conclusions.

[4] We note that such an overlapping signal should have a maximum exactly in between the two primaries. If the luminosity of the primaries (or radial concentration of the satellites) is different, then the maximum overlapping signal should be slightly shifted toward the fainter primary.

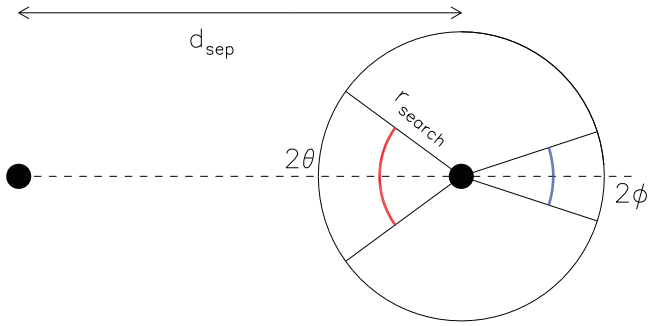

Figure 2. Schematic diagram indicating the geometry of galaxy pairs in the plane of the sky. A pair of primary galaxies (black points) are identified according to the algorithm defined in Tempel et al. (2014c) and described in Section 2. Satellite galaxies within the projected search radius r_{search} are retained as satellites. The line connecting the two primaries, d_{sep}, is used to compute the angles (θ, ϕ), which face toward or away from (respectively) the other primary.

The same cuts are then applied to the overlap sample, namely, only those satellites that fall within 250 kpc of each primary are considered. This is now a combination of each primary's "own" satellites and some contaminants from the other primary system. Regardless of which host the galaxies originally belonged to, the angles (θ, ϕ) are then computed for all satellites within 250 kpc of each host. Note that by construction the isolated primaries used to build the overlap sample have the same magnitude, redshift, and separation distribution as the observational sample.

The "true" signal is thus the observed signal minus the overlap signal. Although in many of our plots we show all three (observed, overlap, and true signals), we emphasize the biases caused by overlapping satellite systems.

2.3. Naming Convention

Two data sets are used in the paper: the observational (Section 2.1) and overlap (Section 2.2) samples. We adopt the following convention (see Figure 2). An opening angle facing the other primary is denoted θ, while an opening angle facing directly *away* from the other primary is denoted ϕ. The angles are measured on the plane of the sky. These angles are centered

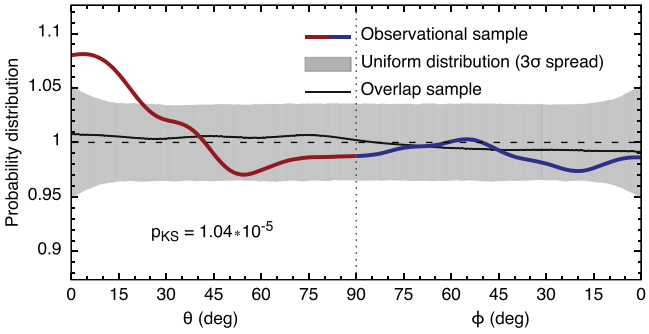

Figure 3. Probability distribution of the angles formed between the position vector of each satellite and the line connecting the two primaries. The red curve represents θ (the angle facing the other primary; see Figure 2), while the blue curve represents the distribution of ϕ (the angle facing away). The gray band shows the 3σ region expected from a uniform distribution of the same size. The excess above this is seen at low θ, and the fact that ϕ always remains well within 3σ is the signature of a lopsided distribution. The K-S probability that the distribution is drawn from a random one is given as $p_{KS} \sim 10^{-5}$ and rules out the null hypothesis at the $\sim 5\sigma$ level. The distribution of angles from an artificially constructed overlapping signal (see Section 2.2) is shown in green and constitutes a minor effect.

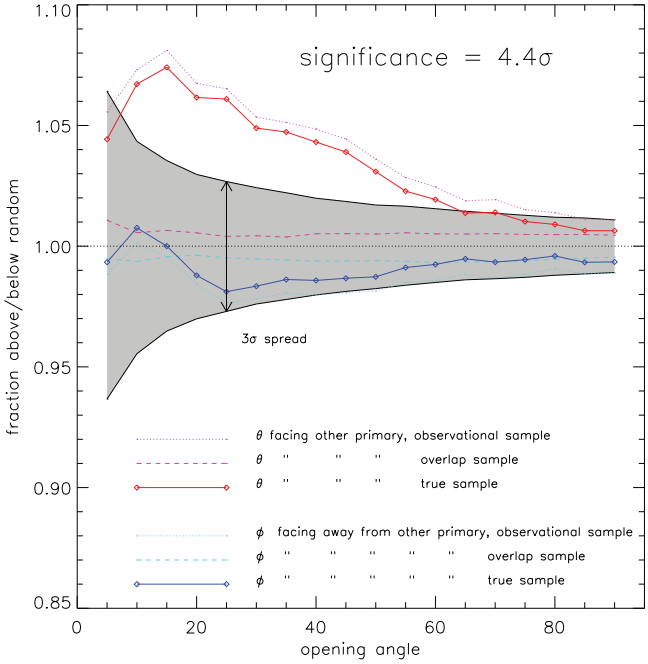

Figure 4. Overabundance of satellites within a cone facing (red curve) or facing away (blue curve) from the other primary shown by simply dividing the measured excess by that expected from a uniform distribution. Clearly, angles facing the other primary always have more satellites than angles facing away—even at 90°, the hemisphere facing the other pair has more galaxies than the hemisphere facing away. The robustness of this claim is made by counting the number of galaxies within the same angle given uniform randomized positions. We randomize the position angle (with respect to the primary) of each satellite 100,000 times and examine the spread in expected over- or underabundances. The 3σ spread is shown by the gray band. Within around 20°–40° the bulging effect is a 5σ deviation from random. The dashed lines show the effect expected if satellite distributions simply overlap. For each pair in the observational sample two isolated galaxies (that have similar redshifts and magnitudes) are placed at the same separation. The angles, with respect to the line connecting the two galaxies, of all the "satellites" within 2.5 Mpc of each host are then randomized, while keeping the projected distance of the satellite to the host fixed. Such an overlap effect constitutes a $\sim 1\%$ effect, is always well within 3σ, and thus cannot explain the bulging seen in the solid curves.

on and hence symmetric about the line connecting the two primaries. Thus, for example, $\theta = 90°$ corresponds to the hemisphere facing the other primary, while $\phi = 90°$ is the hemisphere facing away from the other primary. In what follows we will simply count the number of galaxies within (θ, ϕ) for the observational and overlap samples and compare these to that expected from a uniform distribution.

3. RESULTS

We quantify the anisotropic distribution of satellites by counting how many satellites are within a given opening angle (θ, ϕ). The statistical significance of any anisotropy measured at each opening angle can be gauged via a simple Monte Carlo test. We perform 100,000 trials where each satellite galaxy's radial distance is kept fixed but its angular position around its host is randomized. For each of these 100,000 trials we count the number of (randomized) satellites within a given angle. This provides the expected deviation from a uniform distribution and allows for a statistical quantification of any measured signal. In what follows, these Monte Carlo tests are what is used to statistically quantify the measured signals.

The distribution of angles made between the position vector of each satellite and the line connecting the two primaries is shown in Figure 3. In red we show θ, the angle measured facing the other primary, while in blue we show ϕ, the angle measured facing away from the other primary (see Figure 2). The gray band represents the 3σ spread one expects from a uniform distribution of (θ, ϕ). The distributions are plotted on a single x-axis to highlight that because of the nature of our geometry (θ, ϕ) will never overlap. Since θ shows a strong excess at low values (say, below $\sim 40°$), while ϕ is flat and well within that expected from a random, uniform distribution, we conclude a strong lopsidedness: satellite galaxies tend to preferentially inhabit the regions in between galaxy pairs. There are up to 8% more satellites than expected from a uniform distribution at low values of θ. A similar but significantly weaker effect is seen in the overlap sample (plotted in green), where the probability distributions always lie well within the 3σ interval expected from a uniform distribution, at around 1%. Therefore, we conclude that the lopsided effect seen in the observational sample is not the result of an overlapping satellite distribution. A Kolmogorov–Smirnov (K-S) test can be performed on the full set of angles to quantify the likelihood that these are consistent with being derived from a uniform distribution and is indicated within Figure 3 as $p_{KS} = 1.04 \times 10^{-5}$. Therefore, for the observational sample, this hypothesis is rejected at the 5σ level. The overlap sample is fully consistent with uniform, and the same hypothesis is not rejected by the K-S test.

The bulging effect can be quantified by asking a complimentary question: "how many more satellites are seen, *within* a given opening angle, than expected from a uniform distribution?"—essentially a cumulative version of Figure 3. This is shown by the magenta (for θ) and cyan (for ϕ) curves in Figure 4 for the observed (dotted line) and the overlap (dashed line) samples. The so-called "true" signal is simply the observed distribution divided by the overlap sample and is shown by the solid blue and red curves. The true signal in Figure 4 shows that within the strongest, most overabundant region, $\sim 30°$, bulging approaches an 8% effect. As opening angles increase, the bulging effect decreases, becoming just consistent with uniform (at the $\sim 2\sigma$ level) at 180°, namely, within the hemisphere.

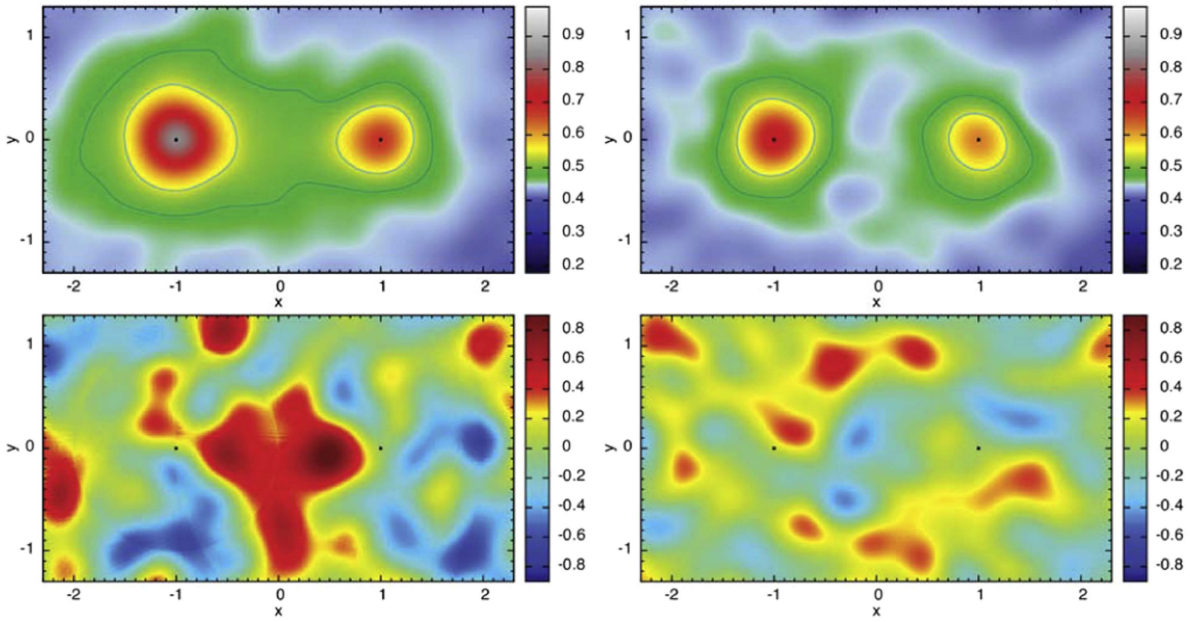

Figure 5. The bulging nature of satellite galaxies can be visualized in a two-dimensional "sky plot," which mimics the view of these systems from the SDSS. Each primary is assigned a position (the bright one is placed at $x = -1$, and the fainter one at $x = 1$), and the number density of satellites is simply contoured with red being high density and blue being low density (top panels). The color bar shows the number of satellites per arbitrary unit area. In the bottom panels the spherically averaged radial distribution computed while masking out the other primary and its satellites is subtracted from each primary to highlight the deviation from spherical symmetry and anisotropic morphology of the satellite population. The color bars in the bottom panels are arbitrary representations of where more satellites are found with respect to the spherically averaged distributions. Left: observed sample; right: overlap sample with the same number of objects as in the observed sample.

The strength of the lopsided signal can be ascertained by computing, in units of σ, the average deviation from uniform of a given signal. For the fiducial sample shown in Figure 4, the significance of the observed sample is more than 4σ.

The relative abundance of satellites with respect to the angle facing *away* from the other primary (namely, ϕ) is roughly constant below unity but within $\sim 3\sigma$ for the observed sample. This means that there is always a deficit of satellite galaxies on the far side of the pair, but this deficit is well within 3σ of what one would expect from uniform distributions of the same size. Any deficit is therefore not statistically significant and is consistent with a uniform distribution. Note that within all angles facing toward and away from the other primary, the overlap sample shows complete consistency with uniform distributions and is thus not sufficient to explain the measured signal. The overlap signal represents a departure from uniformity at the 1.09σ level, while the observational signal for some angles ($\sim 20^\circ$) is over a 5σ effect.

When the effect of overlapping satellite distributions is subtracted from the observed signal, the final true signal is on average 4.4σ away from uniform. Although the true signal "only" reaches a maximum of an 8% increase over uniformity, this 8% is statistically significant and robust. Furthermore, it is also a lower limit, because most of the "satellite galaxies" are interlopers.

The situation can be visualized by stacking all satellites of the galaxy pairs and examining the projected 2D number density of satellites, shown in the top left panel of Figure 5. Here, the brighter primary is given a position at $(x, y) = (-1, 0)$, and the fainter primary is placed at (1,0). The (projected) number density of satellites is then (arbitrarily) contoured and displays a barbell morphology with a clear bulging of satellite galaxies toward each other. The number density of satellites in Figure 5 is the combination of foreground/background galaxies and true satellites. The foreground/background density can be estimated from the outer parts of the figure. While counting satellites around host primaries and subtracting the foreground/background density, we are left with a true number of satellites. Altogether around 15% of all galaxies are true satellites of the central galaxies; most of the galaxies are foreground/background objects.

The averaged radial distribution of satellites (masking out the other primary and satellites belonging to it) can be calculated for each host, directly from the top left panel of Figure 5. If this radial profile is subtracted from the satellite density map, we are left with a morphological impression of the anisotropic spatial distribution of the satellite population, shown in the bottom left panel of Figure 5. Here the overdensity of satellites between the two primaries is clearly visible. In order to compare with the overlap sample, the same plots are shown on the upper and lower right of Figure 5. The morphologies are markedly different: the overlap sample shows no significant contaminatory signal from overlapping galaxies. Moreover, the morphology of the lopsided signal is different than expected from an overlapping satellite distribution.

3.1. Filamentary Environment

One possible interpretation of our findings is that the galaxy pairs, which can be separated by $\Delta z \sim 10^{-3} \approx 300$ km s$^{-1} \approx 4$ Mpc, are aligned along the line of sight with a filament. In this interpretation, the lopsidedness results from the filament's potential itself, rather than from the binary interaction of the galaxy pair. In order to test this hypothesis, we identified which of our galaxy pairs were located in filaments identified with the Bisous method (see Tempel et al. 2014b, 2016). The Bisous method is a marked point process used to identify curvilinear patterns (such as filaments) in the galaxy distribution. It has been employed to

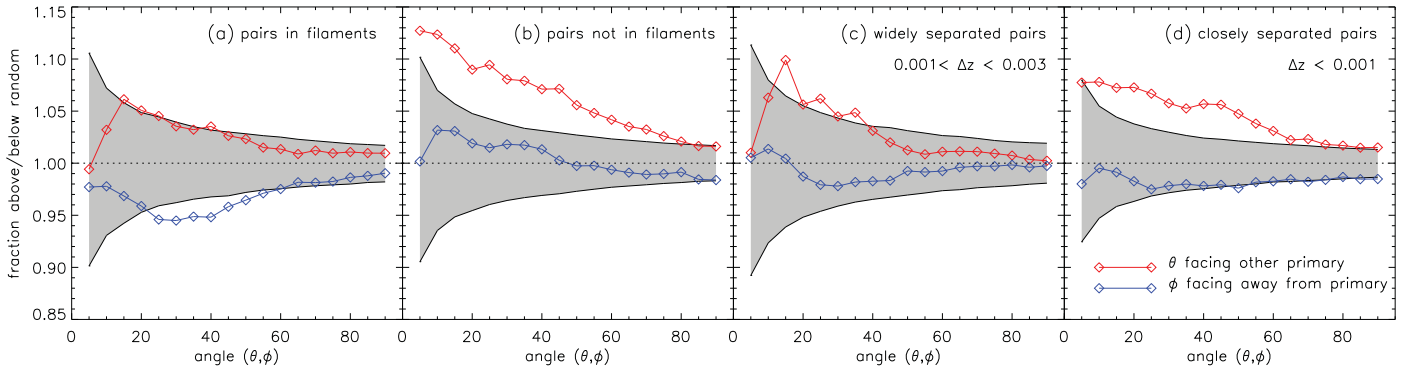

Figure 6. Examining the filament hypothesis. After identifying pairs that are in Bisous filaments (see Tempel et al. 2014b), we compute the lopsided signal for (a) pairs in such filaments and (b) pairs not in Bisous filaments. Pairs in filaments do not show a statistically significant signal. We also examine how the pair's separation (measured as the redshift difference of the two members Δz) affects the signal, with (c) widely ($0.001 < \Delta z < 0.003$) and (d) closely ($\Delta z < 0.001$) separated pairs, under the assumption that if the signal was dominated by filaments, the signal would be stronger for widely separated pairs. The signal is roughly the same in both cases, although in the former (c), the sample size is smaller, so the statistical significance of any signal is weaker.

study, among other things, satellites and galaxy alignments in the SDSS before (Tempel & Libeskind 2013; Tempel et al. 2015; Tempel & Tamm 2015; Pahwa et al. 2016) and tested against other methods for identifying filaments (Tempel et al. 2014a; Libeskind et al. 2015b). In general, it is accepted as a reliable way of identifying such structures in the community.

When applied to the sample here, galaxy pairs are ascribed a simple "within" or "not in" filament status in a similar way to that defined in Guo et al. (2015). The lopsidedness signal for these two subsamples is shown in Figures 6 (a) and (b). Pairs in filaments (Figure 6(a)) show a weaker signal than those not found in filaments (Figure 6(b)), although we note that the sample of the former is smaller and thus any statistical statement is weaker. The most straightforward interpretation of this result is that the lopsided signal is stronger for more isolated pairs, and that when pairs are found within the more dynamically active environment of filaments, the lopsided signal can be weakened. We note that we also examined the strength of the lopsided signal for in-filament pairs depending on whether the filament was oriented along or perpendicular to the line of sight (not shown). In this case when the filament is along the line of sight, the signal is weaker, and when it is perpendicular, the signal is stronger, further supporting this hypothesis. Again, however, such statements are statistically weak, given the reduced sample size.

Another way of testing the filament hypothesis is to divide the sample by Δz into "closely separated" pairs ($\Delta z < 0.001$) and "widely separated" pairs ($0.001 < \Delta z < 0.003$). In Figures 6(c) and (d) we show that the lopsided signal is driven by pairs that are close in line-of-sight separation. Further separated pairs show a weaker, although not altogether absent, lopsidedness in their satellite distribution. That is, our results do not support the filament hypothesis on filament scales of more than $\sim$4 Mpc (which corresponds roughly to $\Delta z = 0.001$).

Our method for finding filaments in the galaxy distribution—the Bisous model—could also be misidentifying (or simply missing) filaments along the line of sight. In other words, there could be filaments causing the effect reported here that we simply do not identify. Although our tests cannot rule out this hypothesis conclusively, the evidence we have presented in Figure 6 is inconsistent with it.

3.2. Systematic Effects

In this section we examine, systematically, how our choice of fiducial parameters may be driving the lopsided effect reported here.

3.2.1. Host Magnitude

The reader will recall that we select only pairs that are within 1 mag of each other, but that fall in a bin 2 mag wide such that our samples are not dominated by the low end of the luminosity function. In Figure 7 we examine how the permitted brightness affects the observed signal, allowing it to vary in half-magnitude increments from $[-19.5, -21.5]$ to $[-22, -24]$ in Figures 7(a)–(f). Note that all other parameters ($r_{\rm search}$, $d_{\rm sep}$, etc.) are kept fixed at this stage and only the host magnitude is varied.

There is no lopsided satellite distribution seen for lower-magnitude pairs (Figures 7(a)–(c)), at least for the other parameters studied here. The signal starts to develop and slightly creep above the 3σ significance for larger pairs (like the LG) in the bin $[-21, -23]$ (Figure 7(d)) and is clearly strongest for the fiducial value $[-21.5, -23.5]$ (Figure 7(e)). Since few galaxies in the SDSS are brighter than $-23.5 < M_r^{0.0}$, changing the bright end limit has little effect on our results. As expected, the highest magnitude range, Figure 7(f), contains few galaxy pairs, and thus the statistical significance is weak in this sample.

3.2.2. Foreground and Background Galaxies as Interlopers

The lopsided signal measured in this work is a lower limit since it is undoubtedly contaminated by a myriad of foreground and background galaxies. Indeed, we have ignored the photometric redshift of the alleged "satellites." However, if these are considered, galaxies whose photometric redshift is inconsistent with the spectroscopic redshift of the pair can be safely eliminated as potential satellites.

In Figure 8 we show how the lopsided signal depends on whether or not the photometric redshift of a satellite is consistent with the spectroscopic redshift of the pair. Specifically, in Figures 8(a)–(d) we only consider a galaxy as a possible satellite if the spectroscopic redshift of the alleged host is within 0.5σ, 1σ, 2.5σ, or 5σ (respectively) of the photometric redshift, where σ is the error on the photometric

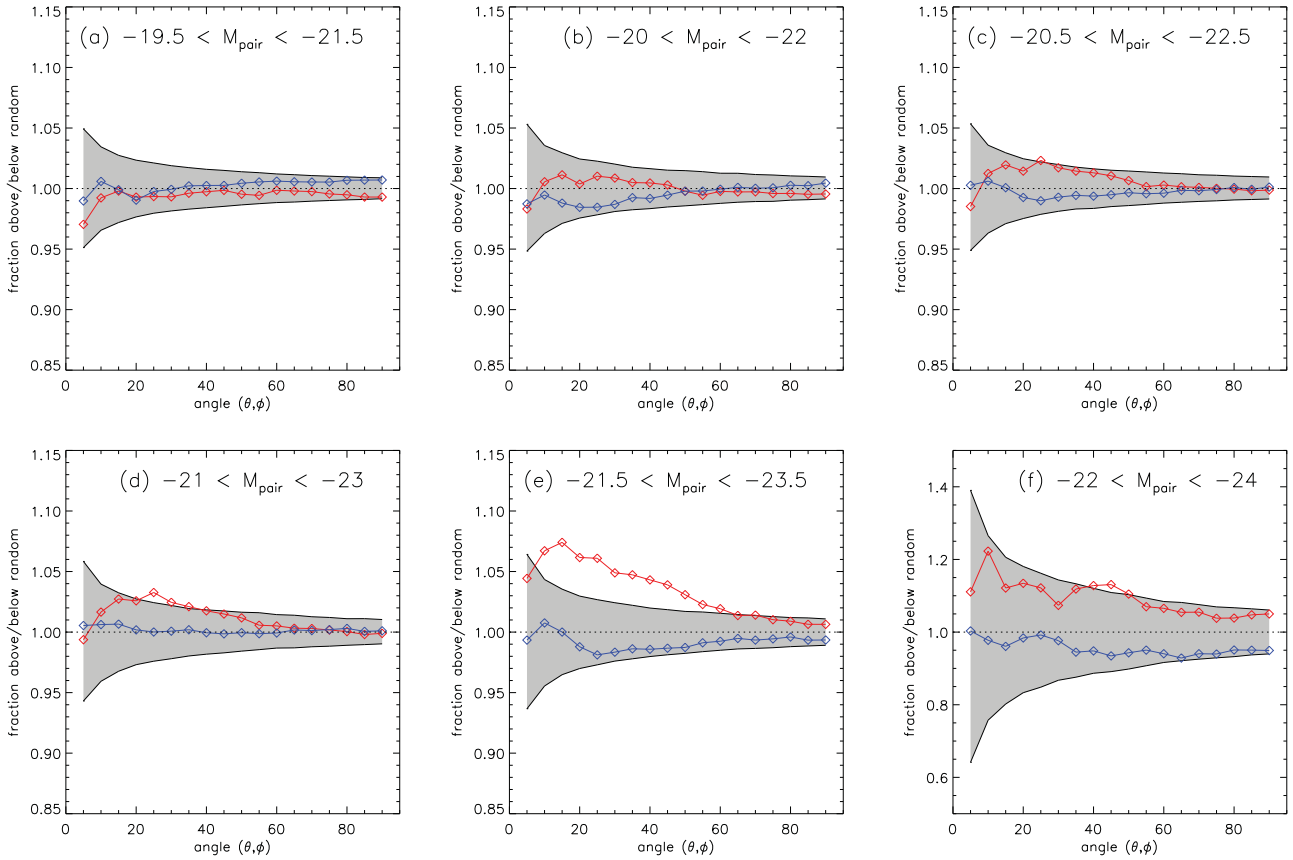

Figure 7. The signal is examined as a function of host magnitude. Pairs where both galaxy magnitudes fall within the range $-19.5 < M_r < -21.5$; $-20 < M_r < -22$; $-20.5 < M_r < -22.5$; $-21 < M_r < -23$; $-21.5 < M_r < -23.5$ and $-22 < M_r < -24$ are shown in (a)–(f). The number of hosts that make the cut drops significantly in the highest magnitude bin (f).

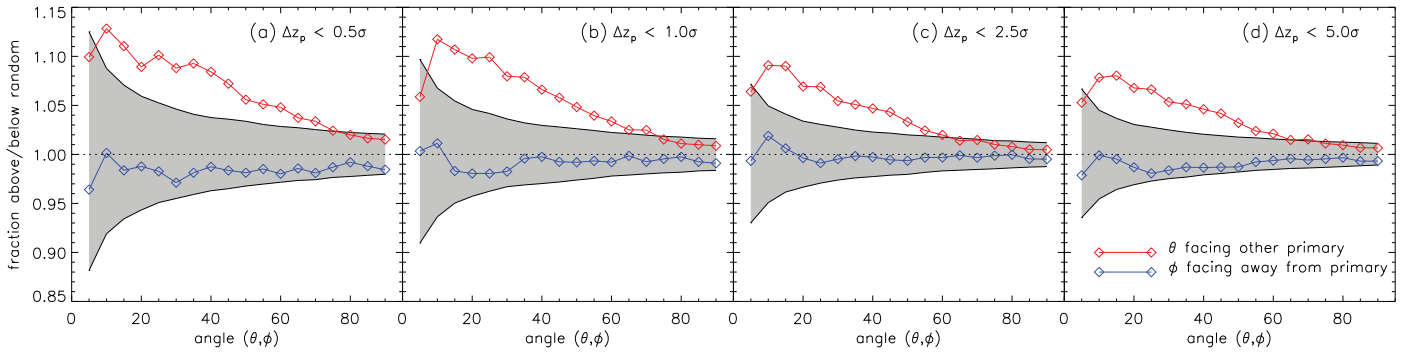

Figure 8. How photometric redshift cuts on the satellite sample affect the signal reported here. We attempt to assign satellite status to a galaxy based on whether the difference between its photometric redshift and the pair's spectroscopic redshift, Δz_p is within (a) 0.5σ, (b) 1.0σ, (c) 2.5σ, or (d) 5σ from the pair, where σ is the error on the photometric redshift determination. The reader will note that the selection of galaxies with $\Delta z_p \lesssim 2.5$ is generous enough that nearly all galaxies are included.

redshift. When photometric redshifts are used to discriminate between interlopers and real satellites, the signal strength increases. Indeed, when only galaxies whose redshift estimation is within 0.5σ of the host, the signal is strongest. As expected, as more and more interlopers are allowed to contaminate the signal, the signal weakens. For all galaxies whose photometric redshift is within 2.5σ of the host, we essentially converge to the "no photo-z cut," since most interloping galaxies have redshifts consistent at this level. The reader will also note that when interlopers are excluded, the sample size is reduced, and thus so is the statistical strength of any statement. In fact, the significance of the signal in Figures 8(a)–(d) is 4.4σ, 4.6σ, 4.2σ, and 4.4σ, respectively. Since the most conservative approach is to present a lower limit, we choose to preferentially present the case where the photometric redshift is ignored, noting that this is a lower limit.

The lopsidedness of just the background galaxies can also be qualified by examining galaxies whose photometric redshift implies that they are *inconsistent* with the spectroscopic redshift of the pair—the opposite of Figure 8. In Figures 9(a)–(d), we examine the signal for galaxies whose photometric redshift is 0.5σ, 1σ, 2.5σ, and 5σ (where σ is the error of the photometric redshift) *greater* than the spectroscopic redshift of the pair, Δz_p. Since the $\Delta z_p > 0.5\sigma$ cut includes some satellites, in this regime the signal is still statistically significant ($\sim 2.3\sigma$). However, for larger cuts the uniformity of

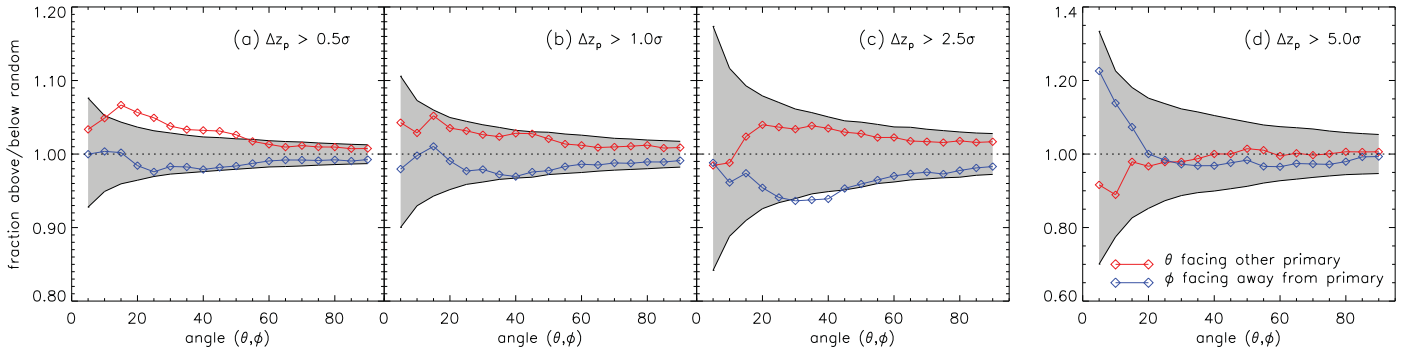

Figure 9. Similar to Figure 8, but here we show the uniformity of background objects, using photometric redshifts. As opposed to Figure 8, where galaxies were *included* as satellites based on the difference between their photometric redshift and the pair's spectroscopic redshift (Δz_p), here we *exclude* galaxies whose Δz_p is too large, specifically greater than (a) 0.5σ, (b) 1σ, (c) 2.5σ, and (d) 5σ. The reader will note that only when we (falsely) include galaxies with $\Delta z_p < 0.5\sigma$ does the "background" show any signal. In all other cases the signal is perfectly consistent with the 3σ spread one expects from uniform distributions. Note the difference in y-axis for this plot, compared with other plots in this paper.

the background becomes apparent as all the curves fall well within the 3σ spread expected for a uniform distribution.

3.2.3. Search Radius

Perhaps the most interesting systematic we examine is the effect of the search radius $r_{\rm search}$ on the lopsided signal. In this case we remind the reader of the following methodology. A galaxy is first and foremost ascribed to the host to which it is closest. Namely, in the case in which one host has a satellite distribution that extends well into its partner's environs, such satellites will be assigned to the partner (i.e., the overlap effect). Also, proximity in and of itself is not enough to be classified as a satellite; one must also be within $r_{\rm search}$. Because galaxies are deemed to belong to those pair members to which they are closest, in the case of closely separated pairs or large $r_{\rm search}$ where $r_{\rm search} > d_{\rm sep}/2$, the search radius effectively becomes half the separation.

In Figures 10(a)–(i) we examine how the signal changes as $r_{\rm search}$ increases. We include the observed, overlap, and true signal here so that the reader can see how all three of these vary with increasing $r_{\rm search}$. We highlight the salient points of this figure: both the observed and the overlap signal are effectively nonexistent at $r_{\rm search} < 100$ kpc and both increase significantly as $r_{\rm search}$ approaches 500 kpc, the largest distance allowed given the range of separations. This behavior is expected: in the closest parts around a host, the overlap effect is minimal as is any asphericity. The lopsided distribution of satellites gets stronger as $r_{\rm search}$ is increased, becoming visible (albeit statistically insignificant) already at 150 kpc. Statistical significance, defined as being greater than 3σ, emerges at 200 kpc and is maximum at 250 kpc (the fiducial value). The significance drops as $r_{\rm search}$ increases, but is still visible (yet statistically weak) at 500 kpc. Figures 10(a)–(i) imply that the signal is driven by anisotropy both in the 200–300 kpc range and within an opening angle of around 45°.

3.2.4. Pair Separation

The final criterion that could introduce a systematic effect is the pair separation $d_{\rm sep}$, chosen in the fiducial sample to be between 500 and 1000 kpc. Recall that in Section 3.2.3, when $r_{\rm search}$ is allowed to be large, we take as our effective search radius $\min(r_{\rm search}, d_{\rm sep}/2)$, a practice we follow here as $d_{\rm sep}$ is allowed to shrink. In Figure 11 we show the effect of changing the (projected) separation from closely separated pairs to widely separated pairs. Signals are seen for pairs separated by [0.25, 0.5] Mpc (11(a)), [0.5, 1.0] Mpc (fiducial value, 11(b)) and [1.0, 1.5] (11(c)), at the 2.2σ, 4.4σ and 1.9σ levels, respectively. Note that only a few pairs are widely separated and make it into the final bin; here the sample size is small (a mere 440 satellites versus, e.g., 46,043 satellites in the fiducial sample), and no signal is found. Although our results do not rule out lopsidedness being found in the slightly smaller or slighter larger separation bins, given this sample the statistical significance is less than the conventional 3σ threshold. Perhaps given larger pair samples from future surveys, more statistically significant results may be obtained in these separation ranges.

4. SUMMARY AND DISCUSSION

Satellites of galaxies located in pairs identified in the SDSS show a statistically significant tendency to bulge toward one another. The identification of pairs is chosen to roughly mimic Local-Group-like systems, namely, to have two central galaxies of similar magnitudes and separation. All galaxies within 250 kpc (projected) of each primary—a region similar to the virial radius—are then chosen, and the angle formed with the line connecting the two primaries is examined. In fact, two angles are examined: the angle facing the other primary and the angle facing away from it. The number of galaxies within a given angle is then compared to that expected from a uniform distribution.

Around 8% more galaxies are seen within a ∼15° angle facing the other primary than expected from a uniform distribution. Roughly the same number are seen in regions 180° away, namely, in the direction opposite to the other primary. This discrepancy is the signature of a lopsided or bulging distribution. Monte Carlo and K-S tests quantify the strength of the bulging signal as upward of 4.4σ—simply put, the chance that such a bulging is produced in random trials is less than 1 in 10 million. The lopsided distribution of satellites of pairs is thus a robust statistical signal.

Since galaxies can have extended satellite distributions, it is possible that this effect is being driven by a simple overlap of the two hosts' satellite populations. In order to distinguish a gravitational effect from an overlap, we conduct the exact same study on a purpose-built "overlap sample." This is constructed by placing two isolated galaxies (chosen to have the same redshift and magnitude distribution as the main sample) and their satellites at the same separation of each observed pair.

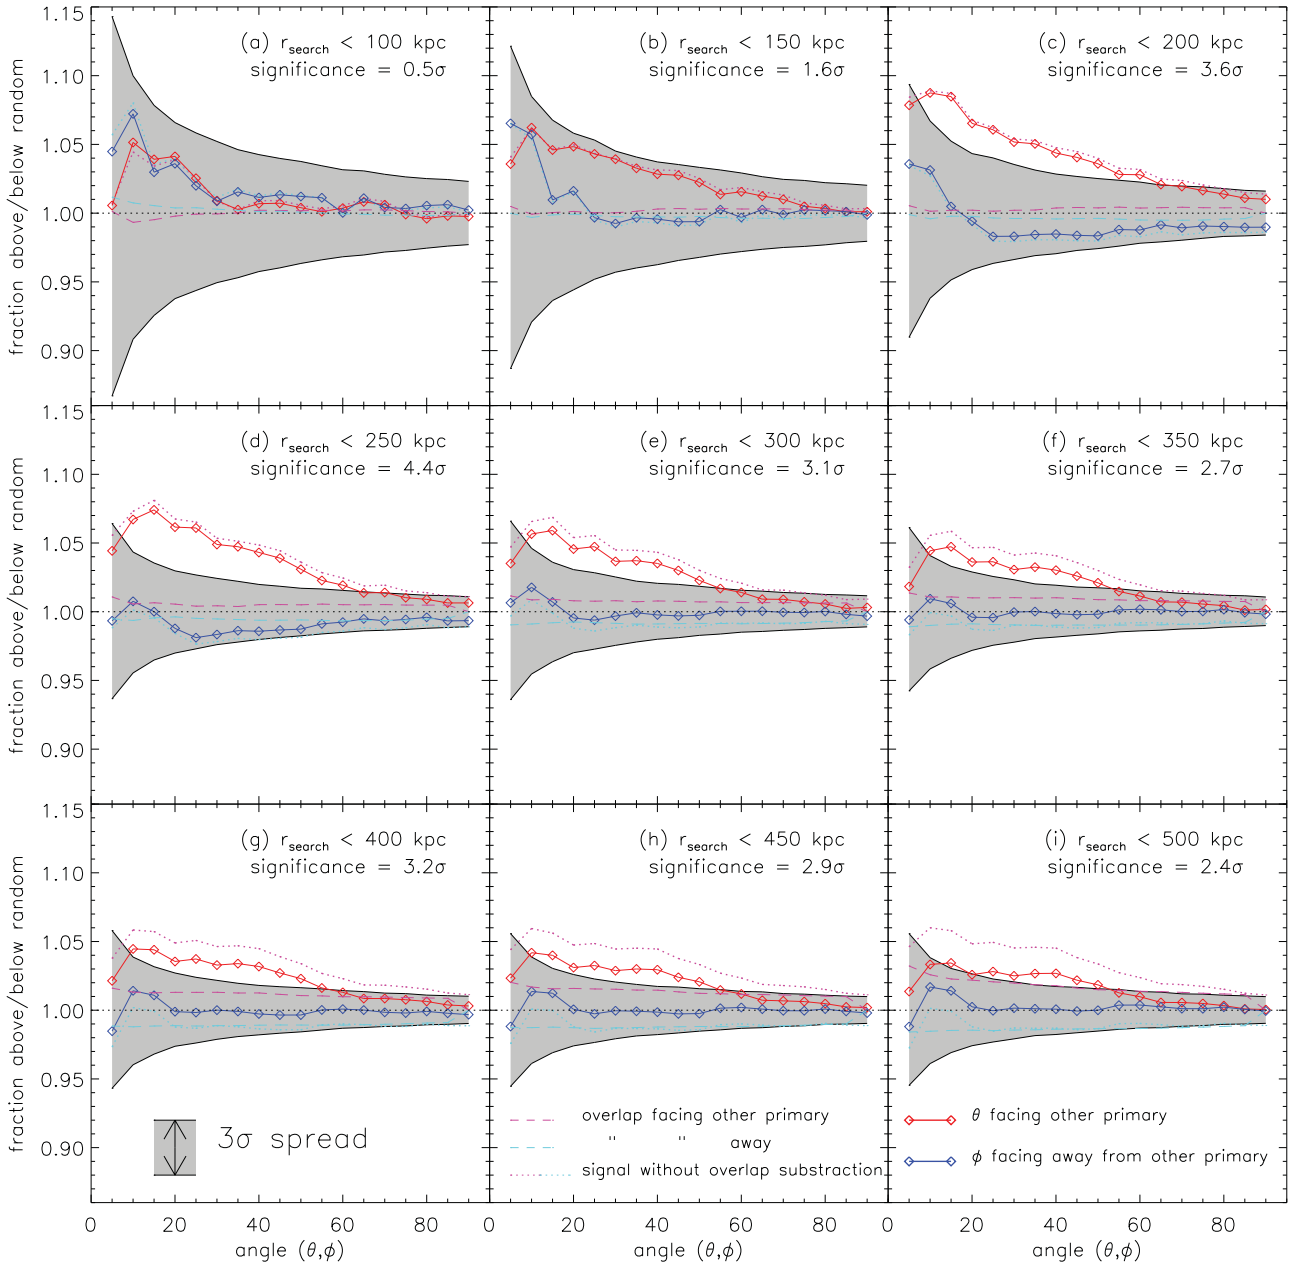

Figure 10. The signal varies as a function of the search radius $r_{\rm search}$ examined. For each satellite, the distance to each pair member is computed, and it is assigned to the host to which it is closest. We show the signal for all satellites that fall within 100, 150, 200, 250, 300, 350, 400, 450, and 500 kpc in (a)–(i) as the dotted magenta and cyan line (for the signal θ facing toward and ϕ facing away from the other pair member). The expected bias due to overlapping satellite distributions, computed by artificially placing isolated galaxies and their satellites at the same separation as the pair members, is shown by the dashed (cyan and magenta) lines. The full signal is computed as the measured signal (dotted lines) divided by the overlap bias (dashed lines) and is shown as the red and blue solid lines. The reader will note that as $r_{\rm search}$ increases, so too does the overlap bias. The 3σ variance expected from a uniform distribution of the same size is shown as the gray shaded region. At each (θ, ϕ) we may compute the statistical significance of the signal and express it in units of σ, the 68% variance. The average value of the significance is indicated in each panel.

The satellites of each of these isolated primaries are then randomized many times while keeping their radial distance from their host fixed, such that the effect of overlapping contamination can be gauged. Our overlap sample quantifies this overlap effect as being consistent with random at well below the 3σ level. The overlap bias can be subtracted from the observed signal to reveal the "true signal," which is quantified by its average departure from uniformity, in units of σ, the standard deviation expected from random trials.

The possibility that some of the galaxy pairs studied here are simply isolated galaxies being observed along a cosmic filament is also considered. By using the SDSS filament catalog published by Tempel et al. (2014b), the nature of satellite distribution among pairs within and not in filaments was scrutinized. Filamentary environment is found to have a negative effect, namely, the dynamically active filamentary environment destroys the lopsided signal. Similarly when pairs that are widely separated along the line of sight (namely, in Δz) are selected, the lopsided effect weakens, hinting toward a gravitational origin.

In order to ascertain whether any of the specific choices that define the fiducial sample examined here were driving the lopsided signal, a systematic study was undertaken, wherein the lopsided effect was examined as a function of host magnitude

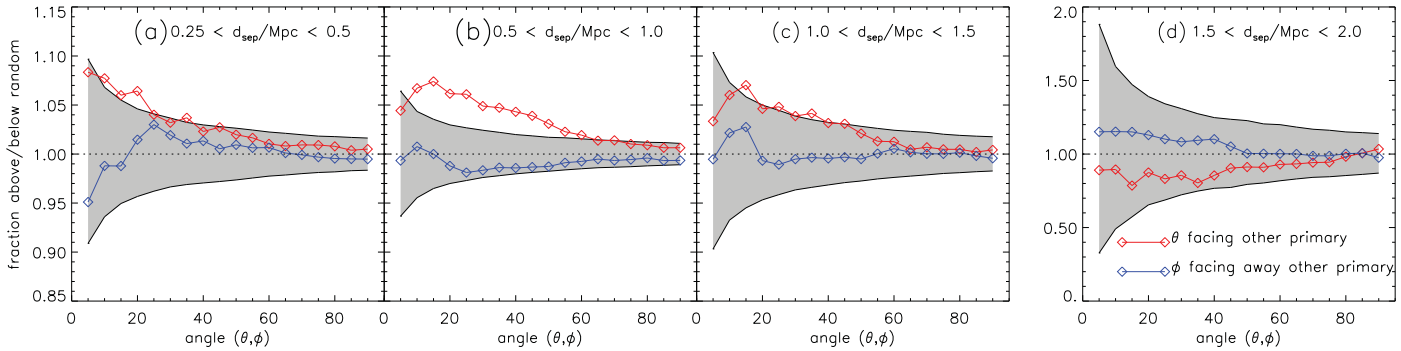

Figure 11. The signal is examined as a function of separation distance. Pairs where the hosts are separated by 0.25 Mpc $< d_{\rm sep} <$ 0.5 Mpc; 0.5 Mpc $< d_{\rm sep} <$ 1.0 Mpc; 1.0 Mpc $< d_{\rm sep} <$ 1.5 Mpc; and 1.5 Mpc $< d_{\rm sep} <$ 2.0 Mpc; are shown in (a)–(d). The number of hosts that make the cut drops significantly in the most widely separated bin (d). For the bin (0.5 Mpc $< d_{\rm sep} <$ 1.0 Mpc) used in our observed sample, the number of pairs and satellites is the largest.

and pair separation. In these cases, often the subsampling results in poor statistics and no credible signal. In the case where the search radius is increased, statistically viable lopsided signals are seen out to large radii.

Since our fiducial sample completely ignores the photometric redshift of the satellites, it is full of interlopers, and the lopsided signal should thus be viewed as a lower limit. When the photometric redshifts are included, and satellites whose redshifts are clearly inconsistent with the host are excluded, the signal is significantly strengthened.

To some extent the alignment of satellites with the geometry of their parents is not wholly unexpected for two reasons. One is that the systems examined here are by construction members of galaxy pairs and thus not necessarily relaxed halos in equilibrium, where spherical symmetry is assumed. Instead, these may be merging, dynamically active pairs. Indeed, the axial symmetries implicit in both the paradigmatic one-dimensional NFW density profile (Navarro et al. 1997) and the extended triaxial profile suggested by Jing & Suto (2002) assume different forms of isolation. It is unlikely that these conditions are met in the samples considered here (or in the Local Group for that matter). In these binary cases, we show that axial symmetry is clearly a poor assumption and that the amount of bulging—at least in galaxy pairs—can boost the number of satellites by up to around 10%.

The second reason why the position of satellite galaxies with respect to the geometry of their parents is not expected to be symmetric is related to the formation of halos within the filamentary network known as the cosmic web. It is known that galaxy pairs and satellites are aligned with the filament they are embedded in (Libeskind et al. 2011; Paz et al. 2011; Lee et al. 2014; Forero-Romero & González 2015; Lee & Choi 2015; Libeskind et al. 2015a; Singh et al. 2015; Tempel et al. 2015; Tempel & Tamm 2015), and it is expected that these filaments channel satellites toward their hosts (Knebe et al. 2004; Libeskind et al. 2014, 2015a; Gonzalez & Padilla 2016). That said, these two empirical findings do not immediately and obviously imply axial asymmetry. In fact, such a lopsidedness may be related to exchange of so-called "renegade" satellite galaxies among such pairs, as seen in simulations such as those of Knebe et al. (2011).

Our results indicate that triaxial modeling of the Milky Way's halo shape based, for example, on the geometry of stellar streams (Law & Majewski 2010) is likely only an approximation in galaxy pairs such as the Local Group, as is mass modeling based on satellite proper motions (or Local Group timing arguments), since these model halos as symmetric or axial-symmetric objects. Future studies will quantify how severe the lopsidedness in our own halo is and thus how inaccurate the halo shape and halo mass measurements are.

Simulations of Local-Group-like pairs are the obvious place to look to see how such asymmetries are formed and how they evolve. Such a study would require analyzing a large cosmological volume (such that the number of galaxy pairs is large) with high resolution (in order to resolve substructures). It remains to be seen whether simulations can shed light on this effect, an avenue of investigation we are currently following.

The axial asymmetry reported here is potentially related to anisotropic satellite galaxy planes. For one, the plane of satellites discovered in Andromeda (Ibata et al. 2013) points to and is lopsided toward the Milky Way. It remains to be seen whether other planes of satellites show similar alignments. The fact that just 4 out of the brightest 11 Milky Way satellites point toward M31 could be indicative of a large Milky Way mass wherein its potential is deeper and therefore more resilient to tidal effects from M31.

In sum, our results show that for galaxy pairs the axial-symmetric modeling of halos may be incorrect. Such a complication undoubtedly limits the applicability of studies that ignore it, for example, in attempts to measure a halo mass or shape based on dynamical tracers. This is often applied to the Milky Way using her stellar streams or satellites. Second, we expect satellite planes in galaxy pairs not to be symmetric about their host but to be biased toward the other partner, as seen in M31. The binary nature of such systems can thus not be ignored when considering satellite distributions. The Milky Way and M31's satellite system should therefore not be viewed in isolation, but must be considered as a single dynamical system.

The authors would like to thank and acknowledge Marcel Pawlowski and Rain Kipper for useful discussions. E.T. acknowledges the support by the Estonian Research Council grants IUT26-2 and IUT40-2 and by the European Regional Development Fund (TK133).

REFERENCES

Agustsson, I., & Brainerd, T. G. 2010, ApJ, 709, 1321
Ahn, C. P., Alexandroff, R., Allende Prieto, C., et al. 2014, ApJS, 211, 17

Allgood, B., Flores, R. A., Primack, J. R., et al. 2006, MNRAS, 367, 1781
Azzaro, M., Patiri, S. G., Prada, F., & Zentner, A. R. 2007, MNRAS, 376, L43
Bardeen, J. M., Bond, J. R., Kaiser, N., & Szalay, A. S. 1986, ApJ, 304, 15
Bond, J. R., Kofman, L., & Pogosyan, D. 1996, Natur, 380, 603
Boylan-Kolchin, M., Bullock, J. S., & Kaplinghat, M. 2012, MNRAS, 422, 1203
Brainerd, T. G. 2005, ApJL, 628, L101
Buck, T., Macciò, A. V., & Dutton, A. A. 2015, ApJ, 809, 49
Cautun, M., Bose, S., Frenk, C. S., et al. 2015, MNRAS, 452, 3838
Conn, A. R., Lewis, G. F., Ibata, R. A., et al. 2013, ApJ, 766, 120
Faltenbacher, A., Jing, Y. P., Li, C., et al. 2008, ApJ, 675, 146
Forero-Romero, J. E., & González, R. 2015, ApJ, 799, 45
Gillet, N., Ocvirk, P., Aubert, D., et al. 2015, ApJ, 800, 34
Gonzalez, R. E., & Padilla, N. D. 2016, ApJ, 829, 58
Guo, Q., Tempel, E., & Libeskind, N. I. 2015, ApJ, 800, 112
Ibata, R. A., Lewis, G. F., Conn, A. R., et al. 2013, Natur, 493, 62
Jing, Y. P., & Suto, Y. 2002, ApJ, 574, 538
Klypin, A., Kravtsov, A. V., Valenzuela, O., & Prada, F. 1999, ApJ, 522, 82
Knebe, A., Gill, S. P. D., Gibson, B. K., et al. 2004, ApJ, 603, 7
Knebe, A., Libeskind, N. I., Doumler, T., et al. 2011, MNRAS, 417, L56
Kroupa, P., Theis, C., & Boily, C. M. 2005, A&A, 431, 517
Law, D. R., & Majewski, S. R. 2010, ApJ, 714, 229
Lee, J., & Choi, Y.-Y. 2015, ApJ, 799, 212
Lee, J., Rey, S. C., & Kim, S. 2014, ApJ, 791, 15
Libeskind, N. I., Frenk, C. S., Cole, S., Jenkins, A., & Helly, J. C. 2009, MNRAS, 399, 550
Libeskind, N. I., Frenk, C. S., Cole, S., et al. 2005, MNRAS, 363, 146
Libeskind, N. I., Hoffman, Y., Tully, R. B., et al. 2015a, MNRAS, 452, 1052
Libeskind, N. I., Knebe, A., Hoffman, Y., & Gottlöber, S. 2014, MNRAS, 443, 1274
Libeskind, N. I., Knebe, A., Hoffman, Y., et al. 2011, MNRAS, 411, 1525
Libeskind, N. I., Tempel, E., Hoffman, Y., Tully, R. B., & Courtois, H. 2015b, MNRAS, 453, L108
Liu, L., Gerke, B. F., Wechsler, R. H., Behroozi, P. S., & Busha, M. T. 2011, ApJ, 733, 62
Lynden-Bell, D. 1976, MNRAS, 174, 695
Majewski, S. R. 1994, ApJL, 431, L17
Metz, M., Kroupa, P., & Libeskind, N. I. 2008, ApJ, 680, 287
Navarro, J. F., Frenk, C. S., & White, S. D. M. 1997, ApJ, 490, 493
Pahwa, I., Libeskind, N. I., Tempel, E., et al. 2016, MNRAS, 457, 695
Pawlowski, M. S., Famaey, B., Jerjen, H., et al. 2014, MNRAS, 442, 2362
Pawlowski, M. S., & Kroupa, P. 2013, MNRAS, 435, 2116
Paz, D. J., Sgró, M. A., Merchán, M., & Padilla, N. 2011, MNRAS, 414, 2029
Phillips, J. I., Cooper, M. C., Bullock, J. S., & Boylan-Kolchin, M. 2015, MNRAS, 453, 3839
Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2014, A&A, 571, A16
Sales, L., & Lambas, D. G. 2004, MNRAS, 348, 1236
Sales, L., & Lambas, D. G. 2009, MNRAS, 395, 1184
Sawala, T., Frenk, C. S., Fattahi, A., et al. 2016, MNRAS, 457, 1931
Singh, S., Mandelbaum, R., & More, S. 2015, MNRAS, 450, 2195
Tempel, E., Guo, Q., Kipper, R., & Libeskind, N. I. 2015, MNRAS, 450, 2727
Tempel, E., & Libeskind, N. I. 2013, ApJL, 775, L42
Tempel, E., Libeskind, N. I., Hoffman, Y., Liivamägi, L. J., & Tamm, A. 2014a, MNRAS, 437, L11
Tempel, E., Stoica, R. S., Kipper, R., & Saar, E. 2016, A&C, 16, 17
Tempel, E., Stoica, R. S., Martínez, V. J., et al. 2014b, MNRAS, 438, 3465
Tempel, E., & Tamm, A. 2015, A&A, 576, L5
Tempel, E., Tamm, A., Gramann, M., et al. 2014c, A&A, 566, A1
Tully, R. B., Libeskind, N. I., Karachentsev, I. D., et al. 2015, ApJL, 802, L25
Wang, J., Frenk, C. S., & Cooper, A. P. 2013, MNRAS, 429, 1502
Wang, W., Sales, L. V., Henriques, B. M. B., & White, S. D. M. 2014, MNRAS, 442, 1363
Wang, Y., Yang, X., Mo, H. J., et al. 2008, MNRAS, 385, 1511
Yang, X., van den Bosch, F. C., Mo, H. J., et al. 2006, MNRAS, 369, 1293
York, D. G., Adelman, J., Anderson, J. E., Jr., et al. 2000, AJ, 120, 1579
Zaritsky, D., Smith, R., Frenk, C. S., & White, S. D. M. 1997, ApJL, 478, L53
Zentner, A. R., Kravtsov, A. V., Gnedin, O. Y., & Klypin, A. A. 2005, ApJ, 629, 219