

A sharper view of Pal 5’s tails: discovery of stream perturbations with a novel non-parametric technique

Denis Erkal,^{1★} Sergey E. Koposov^{1,2★} and Vasily Belokurov^{1★}

¹*Institute of Astronomy, Madingley Road, Cambridge CB3 0HA, UK*

²*McWilliams Center for Cosmology, Department of Physics, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA*

Accepted 2017 May 15. Received 2017 May 15; in original form 2016 September 6

ABSTRACT

Only in the Milky Way is it possible to conduct an experiment that uses stellar streams to detect low-mass dark matter subhaloes. In smooth and static host potentials, tidal tails of disrupting satellites appear highly symmetric. However, perturbations from dark subhaloes, as well as from GMCs and the Milky Way bar, can induce density fluctuations that destroy this symmetry. Motivated by the recent release of unprecedentedly deep and wide imaging data around the Pal 5 stellar stream, we develop a new probabilistic, adaptive and non-parametric technique that allows us to bring the cluster’s tidal tails into clear focus. Strikingly, we uncover a stream whose density exhibits visible changes on a variety of angular scales. We detect significant bumps and dips, both narrow and broad: two peaks on either side of the progenitor, each only a fraction of a degree across, and two gaps, $\sim 2^\circ$ and $\sim 9^\circ$ wide, the latter accompanied by a gargantuan lump of debris. This largest density feature results in a pronounced intertail asymmetry which cannot be made consistent with an unperturbed stream according to a suite of simulations we have produced. We conjecture that the sharp peaks around Pal 5 are epicyclic overdensities, while the two dips are consistent with impacts by subhaloes. Assuming an age of 3.4 Gyr for Pal 5, these two gaps would correspond to the characteristic size of gaps created by subhaloes in the mass range of 10^6 – $10^7 M_\odot$ and 10^7 – $10^8 M_\odot$, respectively. In addition to dark substructure, we find that the bar of the Milky Way can plausibly produce the asymmetric density seen in Pal 5 and that GMCs could cause the smaller gap.

Key words: Galaxy: fundamental parameters – globular clusters: individual (Palomar 5) – Galaxy: structure – dark matter.

1 INTRODUCTION

Palomar 5 is so diffuse that it was once mistaken for a low surface brightness galaxy by Wilson (1955) who ‘rediscovered the globular and called it the Serpens Dwarf, the name that appears surprisingly fitting today, after the detection of the conspicuous S-shaped tails attached to the cluster (Odenkirchen et al. 2001). Naturally, both the low stellar density of the satellite and the prominence of its associated stellar stream are tell-tale signs of the ongoing disruption by the Galactic tides. Over the years, Pal 5’s tidal tails grew and are currently traced across several tens of degrees on the sky (see e.g. Rockosi et al. 2002; Odenkirchen et al. 2003; Grillmair & Dionatos 2006). Thus, Pal 5 has quickly become a poster child for Milky Way accretion. To date, the cluster’s role as a possible powerful and precise Galactic accelerometer has been emphasized by many (see e.g. Ibata, Lewis & Martin 2016), but it remains modelled by few (e.g. Dehnen et al. 2004; Küpper et al. 2015).

Dehnen et al. (2004) who presented the very first – but none the less impressively comprehensive – study of the Pal 5 disruption, established with certainty several key aspects of the satellite’s accretion: the cluster’s orbit, its mass and size, and the importance of disc shocks in driving the mass-loss. However, while getting many observables right, such as the shape of the stream track and the overall behaviour of the debris density along the tails, no model in the Dehnen et al. (2004) suite could match the level of asymmetry between the star counts in the leading and the trailing tail of the cluster, as displayed in e.g. their figs 16 and 4 of Odenkirchen et al. (2003). It was then concluded that the observed asymmetry ought to be due to the processes not captured by the simulations. The authors point out that the most likely phenomenon – not included in their numerical set-up – which could produce such a small-scale density enhancement in one of the tails is the interaction of the stream with a low-mass substructure. They offer three examples of such perturbers: giant molecular clouds (GMCs), spiral arms and dark matter (DM) subhaloes.

The irrefutable detection of small-scale density perturbations in the Pal 5 tails as presented by Odenkirchen et al. (2003) and emphasized by Dehnen et al. (2004) called for an explanation.

* E-mail: derkal@ast.cam.ac.uk (DE); skoposov@cmu.edu (SEK); vasily@ast.cam.ac.uk (VB)

This inspired Capuzzo Dolcetta, Di Matteo & Miocchi (2005) to revisit the numerical experiments of Combes, Leon & Meylan (1999) who had predicted that globular cluster tails ought to contain low-level stellar clumps. In their simulations, Capuzzo Dolcetta et al. (2005) not only confirmed the presence of ubiquitous small-scale substructure in tidal tails, but also provided an intuitive justification of their existence: the stars in the clumps move slightly slower compared to the rest of the surrounding debris in the tail. The deceleration of the stars in the clumps was the clue which helped Küpper, MacLeod & Heggie (2008) establish the genesis of the overdensities: the orbits of the stripped stars in the reference frame of the progenitor are oscillatory, composed of the guiding centre circular orbit and the epicyclic ellipse. The two vertices of the ellipse are crossed at the lowest speed, thus leading to episodic bunching of stars along the tidal tail. Although the presence of these epicyclic overdensities was initially shown for progenitors on circular orbits, further efforts showed that it is also present on orbits with a wide range of eccentricities (e.g. Küpper et al. 2010).

Now, could the clump in the Pal 5’s trailing tail simply be an example of an epicyclic overdensity as described above? This seems unlikely as the simulations of Dehnen et al. (2004) actually produce epicyclic ‘feathering’ of the cluster’s tails as is obvious from e.g. fig. 6: from panel to panel, the bunching appears either enhanced or reduced depending on the orbital phase of the progenitor considered. Remarkably, however, the epicyclic clumping is almost undetectable in the snapshot corresponding to Pal 5’s current location as shown in their fig. 18, thus begging for a conclusion that the apocentre – near which Pal 5 is currently situated – is not the optimal location for the detection of epicyclic overdensities. Most importantly, however, the overdensities produced are always at a very similar level in the leading and the trailing tails, especially so close to the progenitor. This is exactly the point highlighted by Dehnen et al. (2004): whether the epicycles are strong or not, the density profile of the two tails should be symmetric (see their fig. 16). Note also that such a (approximate) symmetry of the tidal tails’ density profiles might simply be a direct consequence of the symmetry of the Hill surface (see e.g. Binney & Tremaine 2008).

Thus, given that the epicyclic clumps – a generic feature of the globular cluster tidal streams – cannot seem to explain the sizeable overdensity in the trailing tail of Pal 5, it is prudent to attempt to establish the actual mechanism behind the apparent asymmetry. Our work is motivated not only by the fact that this conundrum has remained unsolved for more than 10 yr, but also by the two recent advances in the tidal tail studies. First, the Pal 5 stream has been mapped by Ibata et al. (2016) to an unprecedented depth allowing for a robust determination of the minute details of the debris distribution. Armed with this remarkable data set, we use a powerful new and robust non-parametric algorithm to determine the stream track and the associated density variation. We confirm with very high confidence, both the detection of a non-monotonic star count evolution along the trailing tail and the asymmetry between the leading and trailing portions of the stream. In addition, we also find a gap-like feature in the leading tail and evidence of two epicyclic overdensities near the progenitor. At first glance, these findings appear to be in tension with Ibata et al. (2016) who analysed the same data set and found that there are no statistically significant gaps in Pal 5. However, their search was performed on scales smaller than 1° while the search for gaps in this work is focused on larger scales, guided by predictions of the expected distribution of gap sizes from subhaloes in Erkal et al. (2016b) which brings us to the second advance in studies of tidal streams.

There is now a much better understanding of the impact of the massive perturber fly-bys on the structure of the tidal tails (see e.g. Carlberg 2009; Yoon, Johnston & Hogg 2011; Carlberg 2012, 2013; Erkal & Belokurov 2015a,b; Erkal et al. 2016b; Sanders, Bovy & Erkal 2016; Bovy, Erkal & Sanders 2017) – the primary mechanism put forward by Dehnen et al. (2004) to explain the unruly star counts in the Pal 5 tails. While it had been known that a fly-by leads to a density depletion around the projected impact point (see e.g. Ibata et al. 2002; Johnston, Spergel & Haydn 2002), it has now been established that the induced stream gap is always accompanied by density hikes on either side (see e.g. Carlberg 2012; Erkal & Belokurov 2015a). Moreover, stream gaps go hand in hand with stream wobbles: small-scale perturbations visible in *all* phase-space projections of the debris track (see Erkal & Belokurov 2015a,b; Bovy et al. 2017). Finally, notwithstanding the degeneracy between the age of the gap, the mass of the perturber and its speed, there exists a distinct characteristic gap size for subhaloes with different masses as shown in Erkal et al. (2016b), with smaller subhaloes tearing smaller holes. However, there also exists a lower bound to the size of the gap. This minimum size emerges because lighter subhaloes on average impart a smaller velocity kick on to the stream stars. Therefore, it takes longer for the density in the gap to drop to detectable levels hence widening it to sizes comparable to those of the gaps induced by more massive subhaloes. For example, according to Erkal et al. (2016b), it is not feasible to expect DM subhaloes with masses of $10^7 M_\odot$ to produce deep gaps less than 5° wide in a Pal 5-like stream, with most detectable gaps produced by these subhaloes opening to $\sim 10^\circ$.

Interestingly, the fly-by of dark subhaloes is not the only conceivable mechanism that can produce small-scale perturbations in the stream. Naturally, exactly the same generic features described above are also expected from the gaps torn by GMCs (Amorisco et al. 2016). In addition, in Hattori, Erkal & Sanders (2016) it was shown that the rotating bar of the Milky Way can reshape the stream drastically since different portions of the debris approach their pericentre at different times and hence experience a different force from the bar. The influence of the bar was also studied in Price-Whelan et al. (2016b) in terms of the chaos it can create. Sending some of the tidal debris on chaotic orbits can dramatically affect the stream appearance, leading to substantial perturbations of the stream track (e.g. Pearson et al. 2015), as well as stream fanning (Price-Whelan et al. 2016a).

In this work, through a series of numerical experiments involving N -body simulations of the Pal 5-like cluster disruption as well as the approximate stream models based on modified Lagrange Cloud Stripping (mLCS; Gibbons, Belokurov & Evans 2014), we will demonstrate that the observed small-scale disturbances of the Pal 5 tails are consistent with an impact by two low-mass substructures. If the stream features are indeed caused by the passage of dark subhaloes, we argue that the features in the leading and trailing tails are most likely caused by subhaloes in the mass range 10^6 – $10^7 M_\odot$ and 10^7 – $10^8 M_\odot$, respectively. Such subhaloes have long been predicted in Λ cold dark matter (Λ CDM) and their detection would represent a stunning confirmation of the theory. However, unfortunately, with the data in hand, we cannot distinguish the higher mass fly-by from the influence of the Milky Way bar. In addition, we cannot distinguish the lower mass fly-by from a GMC fly-by. The complications introduced by the bar and GMCs also suggest that searches for gaps which focus on streams at larger radii will yield detections that can be interpreted more straightforwardly and hence may be more fruitful.

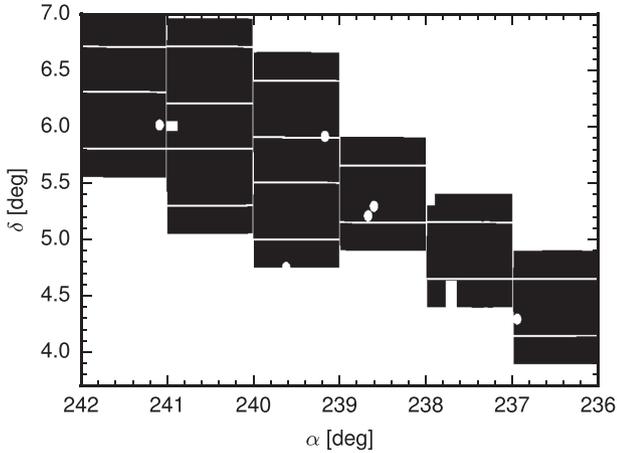


Figure 1. A portion of the footprint derived for the CFHT data set presented in Ibata et al. (2016). The footprint shows noticeable gaps between the CCD chips as well as holes due to masked bright stars.

This Paper is organized as follows. In Section 2, we discuss the deep Canada–France–Hawaii Telescope (CFHT) photometry published by Ibata et al. (2016). In Section 3, we present a novel non-parametric model that we use to extract the stream track and to measure the stellar density variation and the evolution of its width. In Section 4, we give a review of how tidal streams form in static and smooth potentials to highlight the discrepancy with the observed features. Next, in Section 5, we study the effect of an impact by substructure and that of the Milky Way’s rotating bar and show that they can both contribute to the features seen in Pal 5. In Section 6, we discuss several other mechanisms such as the internal Pal 5 rotation, chaos, and perturbations by other globular clusters, as well as how these can be distinguished. The results are compared against expectations in Section 7. Finally, we conclude in Section 8.

2 DATA

In this study, we use the catalogue produced from CFHT observations by Ibata et al. (2016) which is publicly available through the VizieR web site.¹ For details of the observations and the data reduction we refer the reader to the original paper. However, before we can proceed with the analysis, several additional processing steps are required. The most important one is the determination of the survey footprint, i.e. the area of the sky containing data of sufficiently good quality. This step is essential as the CFHT catalogues do not yield a continuous coverage of the stream. To establish the footprint, we consider all individual telescope pointings and all CCDs in the mosaic used for the catalogue creation, while removing all objects located closer than 20 pixels from the edges of each CCD. The remaining objects from all CCDs are then used to construct the combined footprint. To make sure that the data quality is as uniform as possible throughout the footprint, all regions of the data within 3 arcmin of the known bright stars (i.e. those brighter than $V_T = 8$ in Tycho-2) are removed. Additionally, several regions where the source density is noticeably lower compared to the typical levels are masked out. Fig. 1 illustrates a portion of the footprint constructed. Note the multiple apparent CCD chip gaps and holes caused by bright stars. The footprint of the entire data set was set up on a HEALPix (Górski et al. 2005) grid at a high resolution of

$N_{\text{side}} = 65536$ (pixel size of $3\text{arcsec} \times 3\text{arcsec}$). Finally, in addition to excising some of the problematic portions of the data as described above, one group of pointings at $\alpha \sim 245^\circ$ which is disconnected from the rest of the survey was also excluded.

With the footprint in hand we can construct a detailed density map of the stars in and around the Pal 5 stream. Although we would ideally prefer to work with the unbinned data, in practice when dealing with a complex footprint, rather than modelling the set of unbinned positions of the objects on the sky, it is more convenient to describe the Poisson number counts, H_j , in HEALPix pixels ($\alpha_{\text{pix},j}$, $\delta_{\text{pix},j}$). If the pixels are small enough relative to the relevant length scales in the problem then the information content is the same as in the unbinned stellar distribution. Specifically in this study, we use stellar densities calculated inside pixels with $N_{\text{side}} = 4096$ which have a size of $\sim 50 \text{ arcsec} \times 50 \text{ arcsec}$, which is approximately seven times smaller than the width of the Pal 5 stream on the sky.

2.1 Colour–magnitude mask

Matched filters have been used with great success to maximize the signal to noise when constructing density maps of low surface brightness substructure in the stellar halo (see e.g. Rockosi et al. 2002; Odenkirchen et al. 2003; Grillmair 2009). The standard matched filter approach proceeds by weighting each star by the ratio of the target (typically a single stellar population) stellar density in colour–magnitude space, $P_{\text{str}}(g-r, r)$, to the density of the background stars, $P_{\text{bg}}(g-r, r)$. While this weighting scheme produces the optimal signal to noise, its disadvantage is that the weighted densities are not Poisson distributed. Cumbersomely, the distribution of the stellar weights is typically very asymmetric as some stars (e.g. those along the red giant branch) have very high weights, but their incidence on the sky is extremely low. Therefore, to mitigate the above drawbacks of the matched filter approach we use a simpler, more convenient method. Namely, in this paper, candidate Pal 5 stars are selected using boolean masks based on the ratio of $P_{\text{str}}(g-r, r)/P_{\text{bg}}(g-r, r)$, i.e. weights equal to 1 when $P_{\text{str}}(g-r, r)/P_{\text{bg}}(g-r, r)$ is greater than some threshold, and 0 below the threshold. This approach yields signal-to-noise levels that are only marginally lower compared to those obtained with a matched filter, while preserving the Poisson distribution of the stellar densities. The exact value of the threshold which maximizes the signal to noise for $P_{\text{str}}(g-r, r)/P_{\text{bg}}(g-r, r)$ can be easily found from Monte Carlo simulations.

Fig. 2 shows the extinction-corrected² density of stars in the colour–magnitude space (Hess diagram), with stars inside the optimal colour–magnitude mask shown in the right-hand panel of the figure. Note that the mask goes down to a magnitude limit of $g \sim 23.5$ and $r \sim 23$. While the main analysis of this paper was carried out using a $g < 23.5$ magnitude cut, we have verified that all conclusions of this paper remain unchanged if a more conservative cut of $g < 23$ is chosen. We also note that Ibata et al. (2016) measured a small but detectable distance gradient along the stream of $0.009 \text{ mag deg}^{-1}$. Here, we choose to ignore the possible distance variation along the stream. Our colour–magnitude diagram (CMD) mask does not include the red giant branch and thus is not very sensitive to the small shifts along the magnitude direction. However, this should not impact our analysis since the maximum offset in

¹ <http://vizier.u-strasbg.fr/viz-bin/VizieR-3?-source=J/ApJ/819/1>

² Please note that the magnitudes listed in the CFHT catalogue are already corrected for the Galactic dust reddening. For details, see Ibata et al. (2016).

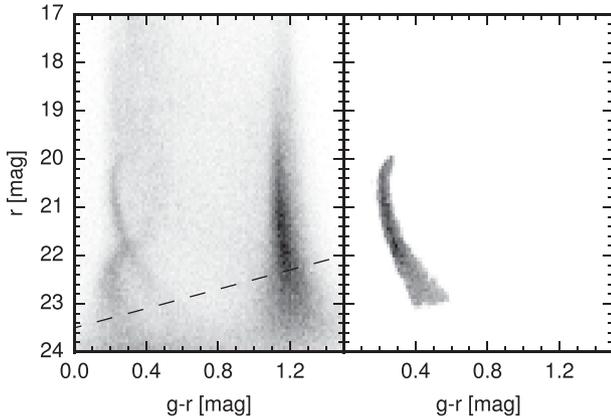


Figure 2. Extinction-corrected colour–magnitude density distribution (a.k.a. Hess diagram) of stars in the CFHT Pal 5 data set published by Ibata et al. (2016). Left-hand panel shows all objects in the catalogue within 10° of the Pal 5 globular cluster, while the right-hand panel gives the subset of stars inside the optimal mask which will be used in this work. Note that the Hess diagram on the left shows not only the stellar main sequence of the Pal 5 (with a turn-off at $r \sim 20.5$) but also that of the leading tail of Sagittarius stream with a turn-off at $r \sim 23$. The dashed line shows the $g = 23.5$ mag limit adopted for this work.

magnitude at the edges of our data set would be ~ 0.13 mag, which is significantly lower than the width of our mask in r magnitude.

2.2 Stream coordinate system

Throughout the paper, similarly to e.g. Koposov, Rix & Hogg (2010), we use a rotated coordinate system, (ϕ_1, ϕ_2) , which is approximately aligned with the stream: the ϕ_1 axis is along the stream, and the ϕ_2 axis is perpendicular to the stream. The pole of this coordinate system is at $(\alpha_p, \delta_p) = (138^\circ 95', 53^\circ 78')$. The zero-point of the coordinate system, $(\phi_1, \phi_2) = (0^\circ, 0^\circ)$, lies at the crossing of the above great circle and the $\alpha = 229^\circ$ great circle. The position of Pal 5 in this coordinate system is $(\phi_1, \phi_2) \sim (-0^\circ 07', -0^\circ 13')$. For convenience, the transformation matrix is also provided in Appendix A.

3 NON-PARAMETRIC STREAM MODELLING

When measuring the properties of tidal tails in large area photometric data sets, the familiar challenge is the absence of an appropriate analytical model which would adequately describe the behaviour of the stellar density in the stream. Stellar streams are more than just Gaussian tubes of stars following great circles on the sky. The stream track deviates from a naive great circles due to the Galactic parallax (see e.g. Eyre 2010) or the precession in the aspherical gravitational potential of the host (see e.g. Ibata et al. 2001; Helmi 2004; Johnston, Law & Majewski 2005; Belokurov et al. 2014; Erkal, Sanders & Belokurov 2016a). The star counts along the stream can have small-scale bumps and dips (see e.g. Küpper et al. 2010; Yoon et al. 2011; Mastrobuono-Battisti et al. 2012), while the across-stream profile can become broader or narrower due to a range of effects, including distance variation, epicyclic feathering, interactions with small sub-structure and differential orbital plane precession (Koposov et al. 2010; Amorisco 2015; Erkal & Belokurov 2015a; Amorisco et al. 2016; Erkal et al. 2016a).

Motivated by the complex picture of stellar streams which has recently emerged from the literature, we develop here a novel, automated, flexible, non-parametric/semi-parametric model. Our stream

model is based on cubic splines that are used to extract the stream’s trajectory, the density variation along its track and the across-stream profile evolution. The cubic splines are also used to model the background stellar density. The main innovative features of our algorithm are as follows. First, the model complexity is set independently for each stream dimension. Secondly, the flexibility of the model is driven by the data and is determined automatically within a probabilistic framework. Finally, because such a model by construction would require a significant number of free parameters, we place a particular emphasis on making sure that the model has computable derivatives. This ensures that the likelihood gradient can be obtained by automatic differentiation software, which allows us to use efficient gradient-based methods for both optimization and sampling of the posterior.

3.1 Cubic splines

The key ingredient of our algorithm, the part of the model that guarantees its flexibility, is the cubic splines. Here, we give a brief summary of the spline properties that are relevant to our particular case. The cubic splines are piecewise polynomial functions described on each interval of the real line as

$$F(x|a_j, b_j, c_j, d_j, x_j) = \sum_{0 \leq j \leq n-1} \mathbf{I}_j(x) (a_j + b_j(x - x_j) + c_j(x - x_j)^2 + d_j(x - x_j)^3), \quad (1)$$

where x_j are the locations of the spline nodes and $\mathbf{I}_j(x)$ are the indicator functions of $x_j < x < x_{j+1}$. At the nodes, the shape of the spline model changes, but with a high degree of smoothness. The polynomial coefficients in the spline, $\{a_j, b_j, c_j, d_j\}$, are set in such a way that the function $F(x)$ and its first and second derivatives are continuous across the whole interval considered. With this condition and the requirement that the second derivatives of $F(x)$ at the edges of interval are zero, $F''(x_0) = 0$ and $F''(x_{n-1}) = 0$ (the so-called natural splines), the coefficients of the polynomials are fully determined by the values of the spline at the nodes, $\{y_j = F(x_j)\}$. Importantly, the polynomial coefficients can be obtained by a simple tri-diagonal matrix inversion operation from the values of x_j and y_j (see e.g. Bronshtein et al. 2007; Press et al. 2007). Thus, with the location of the nodes x_j fixed, the function $F(x|y_j)$ has an easily computable derivative with respect to the values at the nodes y_j .

3.2 Constructing a flexible model of Pal 5’s tails

In this section we demonstrate how, by using cubic splines as a building block, it is possible to assemble a flexible Pal 5 stream model that describes the data with high fidelity. In what follows, it is assumed that the Pal 5 stream has a Gaussian cross-section with a variable width along the stream. The stream density is also variable and the track is allowed to deviate from a great circle. The Gaussianity of the stream profile is clearly an oversimplification, as has been demonstrated in numerical experiments. Simulated tidal tails typically show more complicated behaviour with wisp-like features known as epicyclic feathers (see e.g. Combes et al. 1999; Capuzzo Dolcetta et al. 2005; Küpper et al. 2010; Amorisco 2015). None the less, as demonstrated below, we find that our model provides a satisfactory description of the data in hand, devoid of any noticeable inconsistencies.

As described in Section 2, the data consists of the set of HEALPix pixels on the sky at locations $(\phi_{1,i}, \phi_{2,i})$ and the number counts of the CMD-masked stars in those pixels, H_i (the average number of

stars in the pixel is $\langle H_i \rangle \sim 0.1$). The number counts in each pixel are Poisson distributed, with the expected number of stars in the pixel given by the density function $\lambda(\phi_1, \phi_2)$:

$$H_i \sim \text{Poisson}(\lambda(\phi_1, \phi_2)).$$

The expected number of stars per pixel, $\lambda(\phi_1, \phi_2)$, is the sum of the contributions of the Galactic background/foreground stellar density, $B(\phi_1, \phi_2)$, and the stellar density of the stream stars, $S(\phi_1, \phi_2)$:

$$\lambda(\phi_1, \phi_2) = B(\phi_1, \phi_2) + S(\phi_1, \phi_2).$$

Fortunately, in our case, the model for the background distribution can be made extremely simple as the structure we are studying is very elongated in the ϕ_1 direction while being very narrow in ϕ_2 . Furthermore, the CFHT footprint does not extend beyond several degrees from the stream track. Therefore, the following background model is adopted such that it changes linearly in ϕ_2 :

$$\log B(\phi_1, \phi_2) = B_0(\phi_1) + B_1(\phi_1)\phi_2.$$

However, the average dependence of the background on ϕ_1 , $B_0(\phi_1)$, and the background slope in ϕ_2 , $B_1(\phi_1)$, are modelled non-parametrically by cubic splines with nodes at $\phi_{b0,j}$, the node values $b_{0,j}$ and $\phi_{b1,j}$, $b_{1,j}$, respectively. As noted in Ibata et al. (2016), the Sagittarius stream crosses the leading tail of Pal 5 ($\phi_1 < 0^\circ$) and its presence can also be seen in the CMD shown in Fig. 2. The background model can account for this contamination since the Sagittarius stream is broad ($\sim 10^\circ$, Majewski et al. 2003) compared to the width of Pal 5. Furthermore, since our model only fits stars within the CMD mask shown in the right-hand panel of Fig. 2, this excludes the vast majority of stars from the Sagittarius stream so we do not expect a large contamination.

The model for the surface brightness of the stream is postulated to have a Gaussian cross-section with the width and the density varying along the stream:

$$S(\phi_1, \phi_2) = I(\phi_1) \exp\left(-\frac{1}{2} \left(\frac{\phi_2 - \Phi_2(\phi_1)}{\Sigma(\phi_1)}\right)^2\right).$$

The track of the stream on the sky, $\Phi_2(\phi_1)$, is represented by a cubic spline with nodes at $\phi_{\Phi,j}$ and values at the nodes, $\Phi_{2,j}$. Likewise, the logarithm of the central surface brightness of the stream, $\log I(\phi_1)$, is represented by a cubic spline with nodes at $\phi_{I,j}$ and values at the nodes of $\log I_j$, while the logarithm of the width of the stream, $\log \Sigma(\phi_1)$, is represented by a spline with nodes at $\phi_{\Sigma,j}$ and values at the nodes, $\log \Sigma_j$. We note that the width and surface brightness of the stream, as well as the background density, are parametrized in log-space in order to enforce non-negativity.

As mentioned earlier, the locations and the number of the spline nodes is set *independently* for each component of the model. In other words, the background density, the surface brightness of the stream, its width and the track on the sky can all have a different total number of interpolating nodes placed at different ϕ_1 locations (see Section 3.3 for more details).

The model as described above has been coded in PYTHON using the THEANO (Theano Development Team 2016) module. This module generates and compiles the C++ code from the computation graph, thus providing a high performance likelihood evaluation. Moreover, importantly, the module has a symbolic differentiation functionality, and therefore also delivers the gradient of the likelihood function with respect to the model parameters.

3.3 Automated spline node selection

The main reason for choosing the spline parametrization for the stream modelling is its flexibility compared to e.g. a simple polynomial. Importantly, the spline framework provides a straightforward functionality to increase the model complexity through the addition of extra nodes. Accordingly, there needs to be an algorithm that dictates how to choose the number and the location of the interpolation nodes. For example, one possible solution is to place the nodes in ϕ_1 as densely as possible – that would produce the most flexible, albeit overfitting, model. Instead, we seek a method that can deliver a model whose complexity is data driven, i.e. we would like to select the simplest model that describes the data well, in line with the *Occam's razor* reasoning.

One way to implement such a principle is to start from an initial node placement, $\{x_i\}$, evaluate the likelihood of the data given the node positions, $\mathcal{P}(\mathcal{D}|\{x_i\})$, add a trial new node x_* , and then evaluate the likelihood of the data with the candidate node added $\mathcal{P}(\mathcal{D}|\{x_i, x_*\})$. Ideally, the decision whether to include the new node x_* would be based on the set of criteria that can test the predictive performance of the model, e.g. the cross-validation (Gelman, Hwang & Vehtari 2014). However, for computational reasons, in this work we choose to use the Akaike Information Criteria (AIC; Akaike 1974). Specifically, we start the procedure with four interpolation nodes defined for each of the stream track, width and stream density models (these interpolation nodes are located at the edges of our data, $\phi_1 = -7:1$ and $\phi_1 = 16:3$, and on either side of Pal 5, $\phi_1 = 0:2$ and $\phi_1 = -0:4$), and three interpolation nodes at $\phi_1 = -7:1$, 0° and $16:3$ for the background model. Then, new interpolation nodes are trialled one by one at the locations between the existing nodes. Namely, the candidate nodes are placed at 25 per cent, 50 per cent and 75 per cent of the distance between the previously chosen positions. At the next step, the maximum likelihood is calculated together with the AIC of the model. The new candidate node giving the highest likelihood is chosen if the AIC of the corresponding model is better than the AIC without the new node. If the AIC of the best candidate new node is higher compared to the previous iteration, the node search procedure is stopped.

When testing this procedure on artificial data sets, we have found it to perform well, with the exception of when the data exhibit deviations of borderline statistical significance (i.e. 4σ – 5σ) with a characteristic size much smaller than the distance between existing nodes. In this case, the procedure described above could miss such a feature. From one point of view, this is a desired quality of the algorithm as it provides a natural regularization for the model, suppressing small-scale overfitting. However, because we know that stream densities in particular can have low-amplitude small angular scale fluctuations (i.e. due to epicyclic effects, uneven mass-loss or interaction with perturbers), it may be appropriate for the algorithm to explore small scales more efficiently. Therefore, for the stream density part of the model only [i.e. $I(\phi_1)$], we extend the node placement algorithm with three additional steps. First, we split all node intervals that are longer than 2° , i.e. we insert $\lfloor \frac{x_{j+1} - x_j}{2} \rfloor$ equidistant nodes inside every interval (x_j, x_{j+1}) that is longer than 2° , irrespective of the AIC change. This ensures that the small-scale density deviations are explored appropriately. Then, the node placement algorithm is run again to insert new nodes based on the AIC changes. As a final step, we iteratively remove the (excessive) nodes that do not decrease AIC. We have found that with this modification, the algorithm is noticeably more sensitive to small-scale structures down to $\sim 0:1$. Appendix C provides an illustration of the algorithm's performance on simulated stellar streams. According to these tests,

Table 1. The list of spline interpolation nodes used in the Pal 5 model.

| | |
|--|---|
| Background nodes $\phi_{b0,j}$ (deg) | −7.1, 0., 8.2, 16.3 |
| Background slope nodes $\phi_{b1,j}$ (deg) | −7.1, 0., 16.3 |
| Stream track nodes $\phi_{\Phi,j}$ (deg) | −7.1, −3.75, −0.4, −0.1, 0.05, 0.2, 4.23, 8.25, 16.3 |
| Stream density nodes $\phi_{I,j}$ (deg) | −7.1, −5.425, −4.59, −3.75, −2.91, −1.66, −1.03, −0.82, −0.4 −0.1, 0.2, 0.7, 0.96, 3.23, 4.23, 14.96, 16.3 |
| Stream width nodes $\phi_{\Sigma,j}$ (deg) | −7.1, −3.75, −0.4, 0.2, 8.25, 16.3 |

Table 2. Stream track measurements at the interpolation nodes (medians) together with 16 per cent, 84 per cent percentiles.

| | | | | | | | | | |
|------------------------|--------|--------|--------|--------|-------|-------|-------|--------|-------|
| ϕ_1 (deg) | −7.10 | −3.75 | −0.40 | −0.10 | 0.05 | 0.20 | 4.23 | 8.25 | 16.30 |
| Φ_j (deg) | −0.154 | −0.115 | −0.231 | −0.115 | 0.063 | 0.098 | 0.085 | −0.075 | 0.062 |
| Φ_j (16 per cent) | −0.203 | −0.135 | −0.254 | −0.138 | 0.043 | 0.079 | 0.074 | −0.090 | 0.006 |
| Φ_j (84 per cent) | −0.106 | −0.096 | −0.208 | −0.091 | 0.083 | 0.117 | 0.095 | −0.060 | 0.118 |

the method correctly recovers the debris density evolution including small-scale behaviour such as epicyclic overdensities.

For the CFHT Pal 5 data set, starting with 18 initial interpolation nodes, the procedure described above leads to a model with 39 nodes in total, as specified in Table 1. Note that for some of the parameters such as the stream track and the density along the stream, the number of nodes is significantly higher than for the width of the stream or for the Galactic background, reflecting the higher information content of the data for these stream dimensions. We also remark that even if the total number of nodes assigned by the algorithm is higher than required to describe the data (i.e. at the onset of overfitting), but the full covariance information from the posterior samples between the values at the nodes is preserved, we expect any further inference based on our measurements, such as e.g. constraints on the gravitational potential from the stream track fitting to be unbiased.

3.4 Fitting the data

Here, we describe the choice of the model parameter priors adopted as well as the details of the posterior sampling. A uniform prior is chosen for the stream track, $\phi_{\Phi,j} \sim U(-3, 3)$, while the background parameters, $b_{0,j}$, $b_{1,j}$, and the logarithm of the stream density, $\log I_j$, have improper uniform priors. The prior on the stream width is log-uniform, such that it would lie within the range of 0:01 and 0:5:

$$\log \Sigma_j \sim U(\log(0.01), \log(0.5)),$$

where in practice, instead of the uniform probability distribution, we use its smooth approximation with two logistic functions.

With the priors defined and the likelihood function specified in Section 3.2, we proceed to fitting the model and sampling the posterior. The posterior is first optimized to find the maximum a posteriori parameters using a Quasi-Newton L-BFGS-B algorithm (Zhu et al. 1997). Because of high dimensionality of the problem, multiple fits are always run from a number of different, overdispersed sets of starting points to ensure that the solution is not trapped in a local minimum. Finally, in order to properly explore the covariance between different parameters, the posterior is sampled using the Markov chain Monte Carlo (MCMC). Due to the high number of the model parameters (39) as well as the substantial time required for a single likelihood evaluation, it was crucial to use the Hamiltonian Monte Carlo (HMC) sampling rather than e.g. Metropolis–Hastings sampler. The former boasts much better scaling with the number of problem dimensions d ($d^{5/4}$ versus d^2 ; see Neal 2011) by avoiding random-walk behaviour.

Table 3. Stream width measurements at the interpolation nodes (medians) together with 16 per cent, 84 per cent percentiles.

| | | | | | | |
|--------------------------------|--------|--------|--------|--------|--------|--------|
| ϕ_1 (deg) | −7.10 | −3.75 | −0.40 | 0.20 | 8.25 | 16.30 |
| $\log[\Sigma_j/1 \text{ deg}]$ | −2.073 | −1.764 | −2.172 | −2.103 | −2.230 | −1.836 |
| $\log \Sigma_j$ (16 per cent) | −2.377 | −1.878 | −2.248 | −2.163 | −2.366 | −2.355 |
| $\log \Sigma_j$ (84 per cent) | −1.756 | −1.651 | −2.096 | −2.041 | −2.094 | −1.358 |

The HMC sampler was implemented by the authors in PYTHON, following the recommendations given in Neal (2011). In our realization, each HMC fit starts from the tuning step, where we determine the relative leapfrog step sizes for different parameters, ϵ_i , in order to achieve acceptance rates in the range of ~ 50 per cent–90 per cent. The total number of leapfrog steps in one HMC trajectory was always fixed to 100. This is informed by the test runs that showed that this number of steps produces chains with short autocorrelation times ($\lesssim 1-3$). In the final run, each chain is advanced for 7500 iterations with first 2500 iterations thrown away for the burn-in. To assess the convergence we ran many chains in parallel, ensuring that the Gelman & Rubin (1992) \hat{R} statistic is below 1.05 for all of the model parameters.

The estimates of the model parameters from the MCMC runs, namely the medians and the 16 per cent and 84 per cent percentiles from 1D posteriors are given in Tables 2–4. We have also made the MCMC chains for all parameters publicly available³ as these are required to determine other useful statistics, such as covariance matrices between parameters.

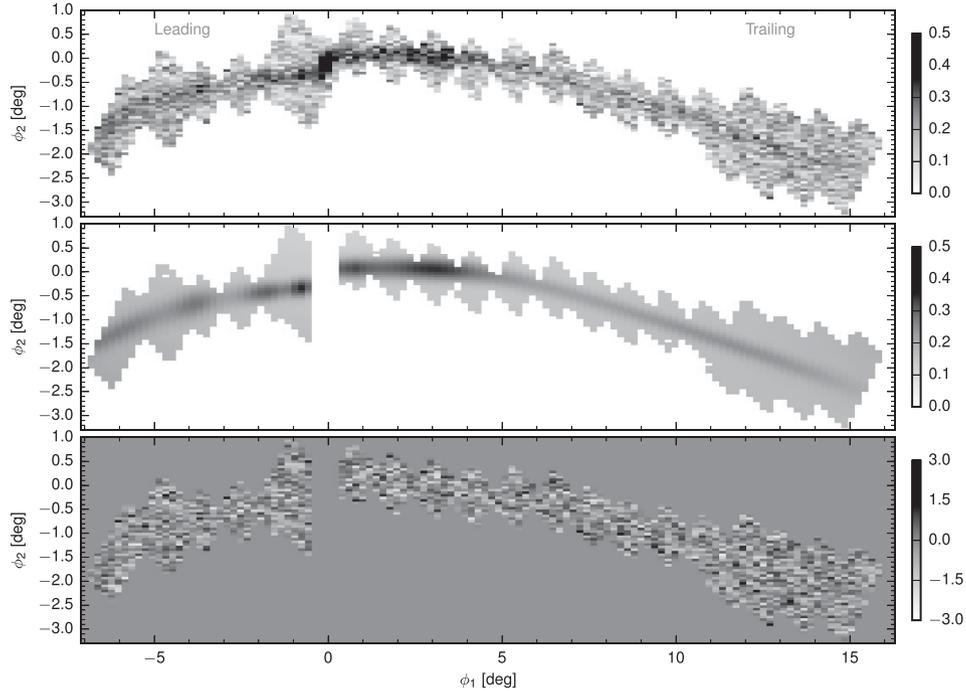
3.5 Results: a sharper view of the Pal 5 tails

This section presents the results of modelling the Pal 5 stream data. First, we focus on the 2D stellar density maps showing the area in and around the stream, comparing the observed features and the behaviour of our model. We also discuss the measurements of the individual stream parameters, such as the debris track on the sky, its width and its density. The overall quality of the stream reconstruction can be gleaned from Fig. 3. The top panel of the figure shows the 2D stellar density distribution in the stream-aligned coordinate system, (ϕ_1, ϕ_2) . The middle panel gives the reconstructed 2D maximum likelihood model of the stream and the Galactic background. In the bottom panel of the figure, the map of the best-fitting model residuals is presented. As the figure clearly demonstrates, the model does indeed correctly reproduce the vast majority of the observed features of the Pal 5 stream leaving no significant residuals.

³ MCMC chains are available at <https://zenodo.org/record/151912>

Table 4. Stream density measurements at the interpolation nodes (medians) together with 16 per cent, 84 per cent percentiles.

| | | | | | | | | | | | | | | | | | |
|--------------------------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| ϕ_1 (deg) | -7.10 | -5.42 | -4.59 | -3.75 | -2.91 | -1.66 | -1.03 | -0.82 | -0.40 | -0.10 | 0.20 | 0.70 | 0.96 | 3.23 | 4.23 | 14.96 | 16.30 |
| $\log I_j$ | -2.74 | -2.46 | -2.67 | -2.07 | -2.92 | -1.76 | -1.83 | -1.25 | -1.70 | 0.38 | -1.12 | -1.48 | -1.57 | -1.35 | -1.91 | -3.73 | -4.92 |
| $\log I_j$ (16 per cent) | -3.36 | -2.63 | -2.86 | -2.22 | -3.13 | -1.88 | -1.98 | -1.41 | -2.03 | -0.75 | -1.41 | -1.59 | -1.67 | -1.42 | -1.98 | -4.10 | -5.59 |
| $\log I_j$ (84 per cent) | -2.19 | -2.31 | -2.49 | -1.93 | -2.73 | -1.64 | -1.69 | -1.11 | -1.41 | 1.47 | -0.86 | -1.37 | -1.48 | -1.29 | -1.85 | -3.41 | -4.35 |

**Figure 3.** *Top:* 2D (in stream-aligned coordinates, ϕ_1 and ϕ_2) on-sky density distribution of the Pal 5 stream candidate stars selected to lie within the optimal colour mask (see Fig. 2). The density is given in units of the number of stars per square arcminute. *Middle:* density distribution of stars predicted by our best-fitting (maximum likelihood) model. The region near the progenitor, $-0.4 < \phi_1 < 0.2$, has been excised as the data from this region were not used in the fits. *Bottom:* 2D map of the stellar density residuals with respect to our model scaled by the expected standard deviation in each pixel.

A more nuanced view of the Pal 5 stream and the model performance can be obtained by glancing at Fig. 4. Here, the observed density distribution and its model representation are given in coordinates corrected for the curvature of the average stream track. The transformed coordinates are $(\phi_1, \hat{\phi}_2)$, where $\hat{\phi}_2 = \phi_2 - \Phi_2(\phi_1)$, so that the stream would be located at the $\hat{\phi}_2 = 0^\circ$ line in the case of a perfect fit. According to the figure, the stream appears to follow the $\hat{\phi}_2 = 0^\circ$ line without deviations. Additionally, the plot highlights some of the important stream features discussed below. In particular, a dramatic density variation along the trailing tail is visible, as well as a significant width evolution along the stream. More precisely, the debris distribution appears much broader in the shorter leading tail compared to the long and narrow trailing tail. Finally, the bottom panel of the figure gives the 1D histogram of the model residuals within 0.2 of the stream centre. Impressively, no significant small-scale overdensities or underdensities are discernible, confirming that the density behaviour of the stream is correctly reconstructed by the model.

The summary of the behaviour of the stream model parameters as a function of the longitude ϕ_1 is presented in Fig. 5. Extracted from the posterior samples and shown here as medians, the 16 per cent and the 84 per cent percentiles are as follows: the run of the stream track on the sky [$\Phi_2(\phi_1)$; second from the top], the stream width [$\Sigma(\phi_1)$; middle panel], the stream central surface brightness [$I(\phi_1)$; second from the bottom panel], as well as the linear stellar density

evolution [$\sqrt{2\pi}I(\phi_1)\Sigma(\phi_1)$; bottom panel]. As evidenced in the figure, the above properties of the Pal 5 tails are measured with unprecedented precision, thus enabling us to make the following key inferences.

(i) The stream track displayed in the top two panels of Fig. 5 exhibits a significant large-scale deviation from the great circle as measured with unprecedentedly high precision: the median ϕ_2 error is spectacularly low at $0:014$. The stream itself seems to flow extremely smoothly without major small-scale irregularities. The only two broad features visible in the track is the expected misalignment between the leading and the trailing tails and a significant change in the stream curvature at $\phi_1 \sim 4^\circ$. We also note that, expectedly, the stream track uncertainties are truly minuscule in the central, most populated, parts of the stream and start to expand closer to either end at high values of $|\phi_1|$, where the tails are traced with fewer stars.

(ii) The recovered stream width profile demonstrates the power of the adaptive non-parametric approach employed. While the longer trailing tail seems to have an almost constant width of $0:12$, the shorter leading tail possesses a very quick and significant width increase, from $0:12$ to $\sim 0:18$ on moving away from the progenitor. Given the CFHT coverage available, it is not clear whether the leading tail width stays high further out at $\phi_1 < -5^\circ$ or returns back to the original narrow values of $\sim 0:12$. We also note that the width

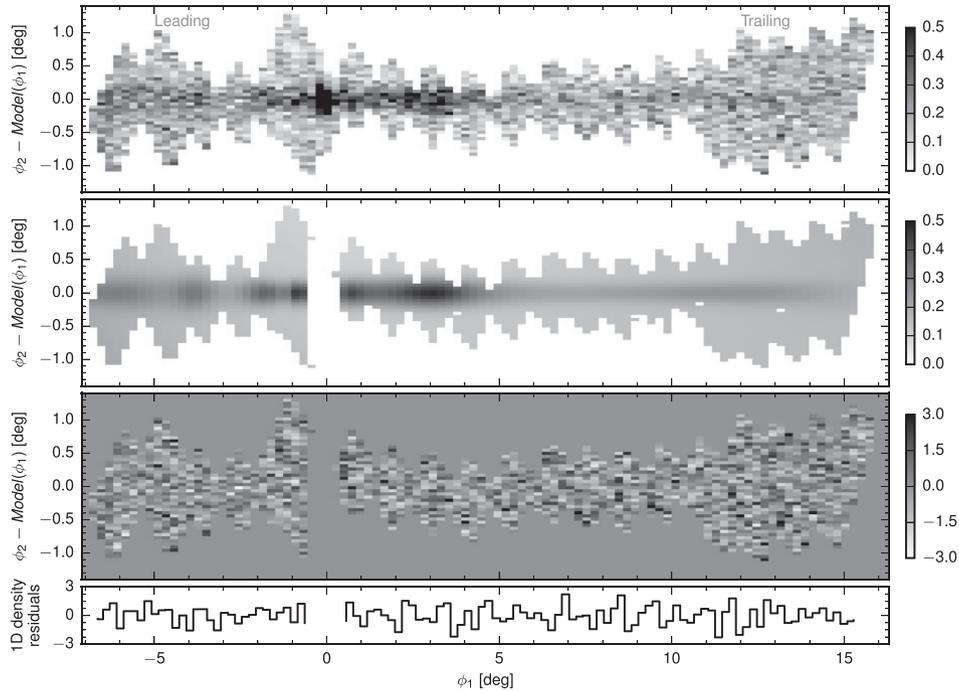


Figure 4. *Top, second and third panels:* observed, modelled and residual density distributions of the Pal 5 candidate stars, similar to Fig. 3, but in a coordinate system where the best-fitting stream track, $\hat{\phi}_2 = \phi_2 - \Phi_2(\phi_1)$, has been subtracted, i.e. the stream should follow a horizontal straight line $\hat{\phi}_2 = 0$. *Bottom:* 1D distribution of density residuals for the data within 0.2 of the stream track, where the residuals have been normalized by the standard deviation.

of the stream at the very edge of the trailing tail at $\phi_1 \sim 16^\circ$ appears to grow up to somewhat higher values, i.e. nearly 0.15 . However, the stream width error bars also increase significantly due to the low stream density far away from the progenitor (see the bottom panel of the figure). Therefore, we believe that there is no strong evidence for the width increase in the trailing tail.

(iii) The most interesting and surprising measurement is perhaps that of the Pal 5 stream surface brightness (shown on the second from the bottom panel of Fig. 5). First, we would like to draw attention to the stream’s average surface brightness levels. Even though the CFHT data reach an impressive depth of $r \sim 23.5$, and in spite of the fact that the Pal 5’s debris is the most prominent globular cluster stream in the Milky Way halo, typically, there is still only one stream star per 10 arcmin^2 . This means that despite the application of a matched filter, the Galactic foreground/background density is at a similar or even higher level (as illustrated by the red band in the second panel from the bottom) across the lion’s share of the surveyed area. This truly emphasizes the extreme challenges associated with the analysis of the low surface brightness halo substructures.

(iv) Notwithstanding the rather humble star counts in the stream, the evolution of the debris density is captured with high precision and is shown to exhibit prominent variations. Let us start by looking at the trailing tail ($\phi_1 > 0^\circ$), where the most noticeable feature is located, namely the very clear surface brightness peak at $\phi_1 \sim 3^\circ$ [labelled (*d*)]. The surface brightness in this debris pile-up is more than three times that of the typical level in the rest of the tail. Most importantly, this striking trailing tail feature does not seem to have a detectable counterpart on the leading side of the stream. Curiously, beyond the bump, i.e. at higher ϕ_1 , the stream density decreases but then briefly recovers at $\phi_1 \sim 12^\circ$ [(*e*) label in the figure], where another, albeit lower amplitude bump can be seen. In fact, the bumpy behaviour of the trailing tail can already be glimpsed from the 2D density maps shown in Fig. 3 and particularly Fig. 4. This strong

density variation present in the trailing tail and missing from the leading tail will be the main focus of the rest of the paper. Farther, past the second shallow density peak ($\phi_1 \gtrsim 12^\circ$), the star counts in the stream appear to drop. It is unclear whether the stream stops, i.e. the density decays to zero, or simply reaches the levels below our detection limit.

(v) Pal 5’s leading tail ($\phi_1 < 0^\circ$) also boasts a number of interesting features in its density profile. In order to quantify the significance of these density fluctuations, we introduce the following statistic. The density excursion parameter, S , is simply the ratio of the density level at the location ϕ_1 to the density linearly interpolated between two fiducial background points ($\phi_{1,l}$, $\phi_{1,r}$):

$$S = \frac{(\phi_{1,r} - \phi_{1,l})D(\phi_1)}{(\phi_{1,r} - \phi_1)D(\phi_{1,l}) + (\phi_1 - \phi_{1,l})D(\phi_{1,r})}.$$

Therefore, S approximates the amplitude of the density drop or enhancement at given ϕ_1 . If the density changes linearly the S statistic will be equal to 1, but if there is an x per cent density drop, its value will change to $1 - x/100$. Samples from the posterior (see Section 3.4) are used to evaluate the uncertainty in the S statistic and thus assign significance to a possible depletion or overdensity in the stream star counts by looking at the tail probabilities $\mathcal{P}(S < 1|\mathcal{D})$ or $\mathcal{P}(S > 1|\mathcal{D})$, respectively. Note that for depletions, the fiducial points are chosen to be the peaks closest to the gap, while for overdensities, the fiducial points are chosen to be the nearest troughs.⁴ Equipped with the S statistic, we examine the structures visible in the leading tail. Closest to the progenitor is the sharp peak at $\phi_1 \sim -0.8$ [labelled as (*b*) in the figure]. This density spike is

⁴ We note that the tail probabilities based on the S statistic measure the local significance, so the Bonferroni correction (Gross & Vitells 2010) with the number of tests equal to ~ 15 (the number of parameters of the density model) may need to be applied to evaluate the global significance.

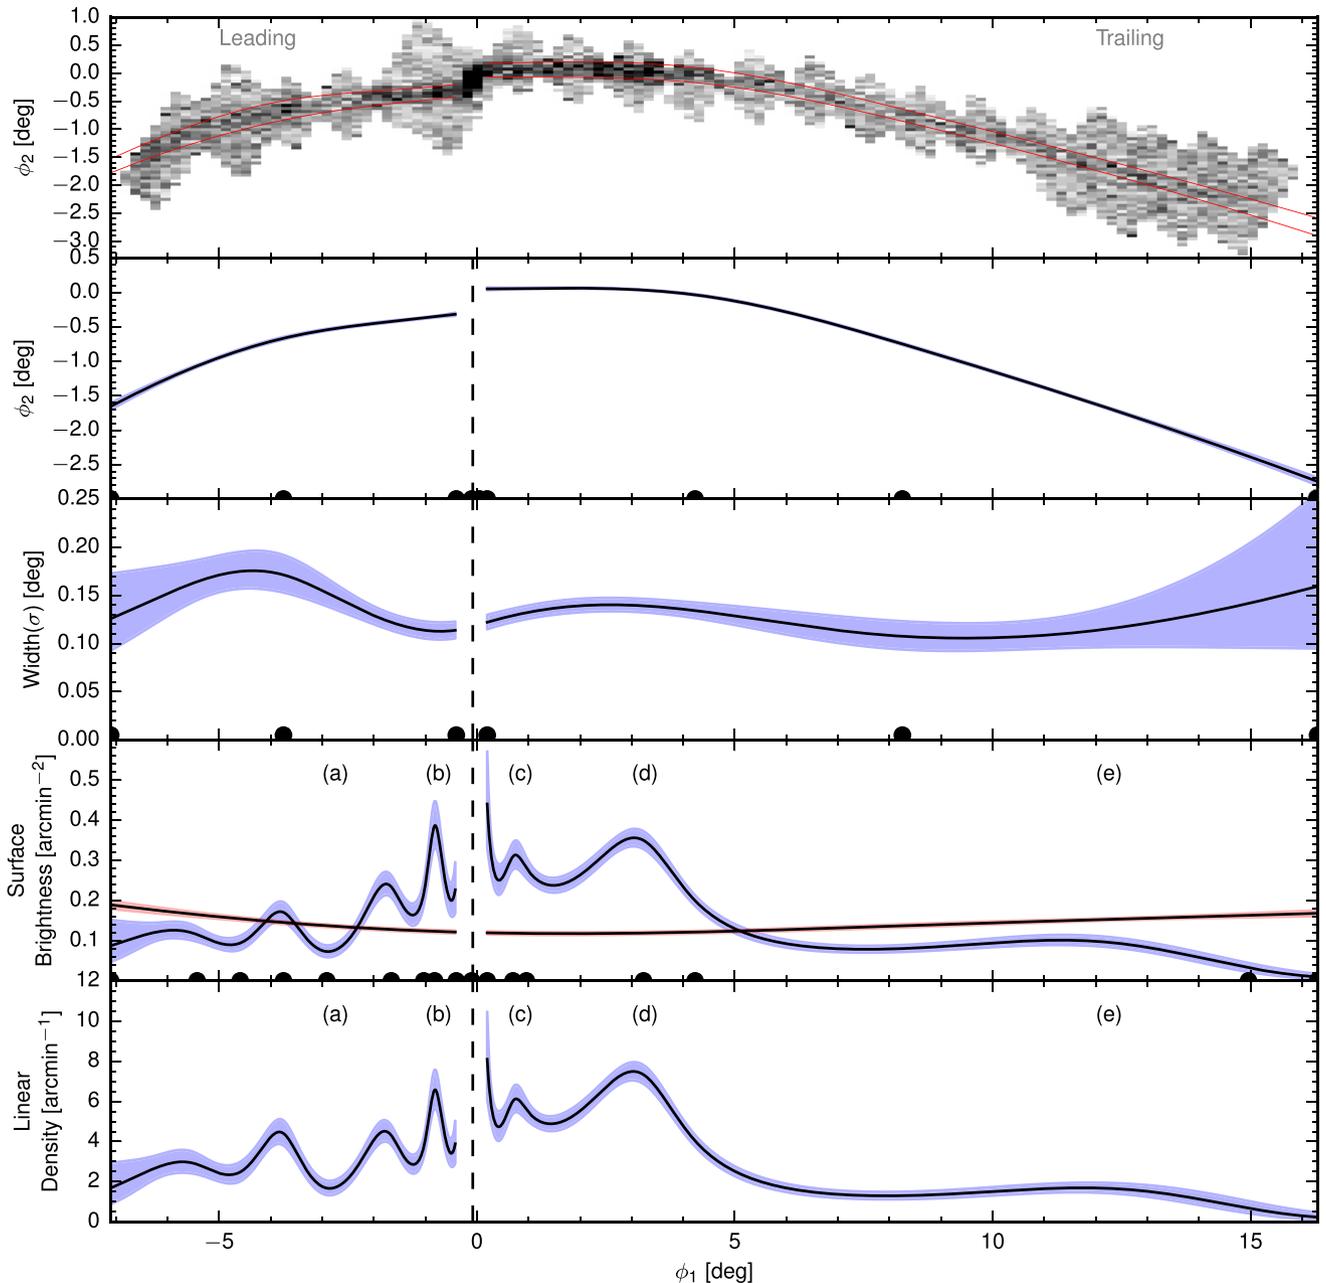


Figure 5. Summary of the measured Pal 5 stream properties, each shown as a function of the along-stream longitude ϕ_1 . The dashed line indicates the location of the Pal 5 cluster. *Top:* 2D density map of the stream stars with the red lines delineating the stream track \pm stream width. *Second panel:* distribution of the posterior samples of the stream track. In this and the rest of the panels below, the black line shows the median of the parameter in consideration, while the blue band gives the 16 per cent and 84 per cent credible intervals. Filled black circles at the bottom of each panel mark the locations of the interpolation nodes chosen. *Third panel:* stream width evolution. Note that in the trailing tail, the width of the debris distribution remains approximately constant. This is in stark contrast with the leading tail, which broadens significantly at $\phi_1 \sim -3^\circ$. *Fourth panel:* stream central surface brightness. For comparison, the surface brightness of foreground/background stars is shown in red. Note several prominent features in the trailing arm: (i) the rise of the surface brightness near the progenitor, (ii) a striking overdensity at $\phi_1 \sim 3^\circ$ [labelled by (d)], (iii) the density depletion near $\phi_1 \sim 8^\circ$ and (iv) another lower amplitude overdensity at $\phi_1 \sim 12^\circ$ [labelled by (e)] followed by apparent complete or almost complete density drop off at $\phi_1 \gtrsim 15^\circ$. Similarly, in the leading tail, there are several notable features as well: (i) a rise of the star counts towards the progenitor, (ii) a narrow peak at $\phi_1 \sim -0.8$ [labelled by (b)] and (iii) a pronounced dip at $\phi_1 \sim -3^\circ$ [labelled by (a)] surrounded by low-significance bumps on either side. *Fifth (or bottom) panel:* linear density along the stream. Since the stream width is quite constant with the exception of the edge of the leading stream, the linear density mimics closely the surface brightness profile along the stream.

quite narrow (width of ~ 0.2) and has a respectable significance of $\sim 3.3\sigma$ using fiducial points at -1.2 and -0.5 . The likely explanation for this compact overdensity located right next to Pal 5 itself is the so-called epicyclic ‘bunching’, a phenomenon com-

monly observed in simulations of globular cluster disruption (see e.g. Combes et al. 1999; Capuzzo Dolcetta et al. 2005; Küpper et al. 2010; Amorisco 2015). Indeed, in N -body simulations presented in Section 4, an epicyclic overdensity can be seen at almost exactly the

same location in the leading arm, i.e. at $\phi_1 = -0.7$. Interestingly, on the other side of the progenitor, the trailing tail also seems to have a small peak at $\phi_1 \sim 0.9$ [labelled (c)]. We can speculate that this feature is also related to the epicyclic bunching, however, in the current data set it is not very statistically significant ($< 2\sigma$).

(vi) Additionally, the leading tail shows a considerable density decrease at $\phi_1 = -3^\circ$ [labelled (a) in the figure]. The significance of this drop is higher than 5σ using fiducial points at -4.0 and -1.8 (the true significance cannot be computed as none of our posterior samples has $S > 1$). The S statistic value is 0.3 ± 0.1 indicating that the density decrease in this feature is about 70 per cent. In fact, this strong debris depletion can also be spotted in the 2D density maps of Fig. 4. We will discuss the possible cause of this feature in latter sections. As a note of caution on this and other features in the leading tail, we emphasize that while the detection of the density variations at $-4^\circ < \phi_1 < -0.5$ is unambiguous, the classification of the observed features as overdensities and underdensities is somewhat model-dependent and as such is open to interpretation.

(vii) Finally, the stream density in the leading tail has a small-scale density drop visible at $\phi_1 \sim -5^\circ$. However, since its significance is only $\sim 2.5\sigma$ using fiducial points at -5.7 and -3.9 , its nature remains uncertain. Furthermore, by examining the extinction maps in this area, we noticed that a filament of increased reddening [$E(B - V) \sim 0.16$] crossing the stream roughly at the location of this feature. In other words, it is possible that this particular density feature is spurious and is related to the effects of obscuration by the interstellar dust. To demonstrate this point quantitatively, we show the average reddening within one stream width of Pal 5 in Fig. D1 which shows there is only significant reddening near $\phi_1 \sim -5^\circ$.

The amount of structure visible in the Pal 5 debris distribution is remarkable. However, while hints of some of the features had been seen previously, others are revealed here for the first time. Therefore, to verify the robustness of the stream measuring machinery, several consistency checks have been performed. The results of two of these checks are presented in the Appendix.

The first check (see Appendix B) compares the model constrained on the CFHT data to the stream density profile obtained with the data from the DECam Legacy Survey (DECaLS; Blum et al. 2016). Although DECaLS data clearly suffers from similar Poisson sampling errors, the instrumental effects are expected to be different. Re-assuringly, as shown in Fig. B1, the independent DECaLS data set appears to exhibit very similar density fluctuations to the ones extracted from CFHT data. The second consistency check is reported in Appendix C. Here, the same adaptive non-parametric density measurement algorithm is used on simulated stream data. The streams with and without the expected Λ CDM substructure are analysed. Importantly, in all cases, the reconstructed density profile matches well the underlying density, correctly identifying the presence (or lack) of the density features over a large range of angular scales. Finally, we have also checked that the features are not due to a varying detection efficiency of stars across the field of view (as seen in fig. 6 of Ibata et al. 2016). More precisely, we mask out the bottom half of the bottom right CCD and the top left corner of the top left CCD and repeat the analysis. Reassuringly, this gives an almost identical stream track, width and density, as well as depth of each stream overdensity/underdensity (i.e. the S statistic of the most significant density drop at $\phi_1 = -3^\circ$ is $S = 0.35 \pm 0.1$, indistinguishable from the measurement without masking of $S = 0.3 \pm 0.1$).

Now let us briefly compare our results with the analyses of the Pal 5 stream reported previously. For example, in one of the very first studies of the Pal 5 tails, Odenkirchen et al. (2003) show the stream star count behaviour in their fig. 4. There, one can already get a glimpse of a strong density peak at $\phi_1 \sim 3^\circ$ and the resulting leading/trailing tail asymmetry. The later work by Carlberg, Grillmair & Hetherington (2012) presents the results of a search for gaps in the trailing tail of Pal 5 using the Sloan Digital Sky Survey (SDSS) data. One of the most significant gaps they found was at the distance of 8.45 from the progenitor, with a size of 7.7 which agrees quite well with the gap-like feature visible in the bottom panel of Fig. 5. This is an underdensity that continues from $\phi_1 \sim 3^\circ$ to $\phi_1 \sim 12^\circ$ [between peaks labelled as (d) and (e) on the figure]. In their analysis of the SDSS data, Küpper et al. (2015) found many possible overdensities, some of which could be matched to the stream features reported here. For example, their most prominent overdensity (labelled T4) can be matched to the large density peak that we observe at $\phi_1 \sim 3^\circ$. Furthermore, their overdensity near the progenitor labelled L1 could be identified with the epicyclic overdensity that we see at $\phi_1 \sim -0.8$.

Lastly, we can compare our results with those of Ibata et al. (2016). In the Pal 5 panorama presented in their fig. 7, one can observe a density peak at $\phi_1 \sim 3^\circ$ (located at $\xi \sim 2.5$ in their coordinate system), as well as a hint of the gap at $\phi_1 \sim -3$ (at $\xi \sim -2^\circ$), best seen in the bottom two left panels of their fig. 8. One of the main differences between this work and Ibata et al. (2016) is that they limited their search to features below 1° in size while we searched for features on a range of scales motivated by the gap size predictions of Erkal et al. (2016b) and by the small-scale features expected from epicyclic overdensities. Indeed, in the left-hand panel of their fig. 8, where they show the results of their match filtering technique on different scales, the gap at $\xi \sim -2^\circ$ is growing in significance as the search is performed on successively larger scales up to 1° . If their search had continued to larger scales, $\sim 2^\circ$, it is likely that they would have found a statistically significant gap. We note that Ibata et al. (2016) caution that background subtraction can introduce gaps and overdensities in the stream due to inhomogeneities in the survey. As a check that this is not the case for our claimed features, we show the background-subtracted density from DECaLS in Appendix B which shows precisely the same features as seen with the CFHT data.

4 STREAMS IN A SMOOTH AND STATIC POTENTIAL

In order to understand the significance of the features observed in the tidal tails of Pal 5, let us first review the mechanism of the stellar stream formation in the simplest case, namely in a smooth and static potential.

4.1 Mechanics of tidal disruption

The basics of the process that leads to the emergence of narrow tidal tails around in-falling satellites are now well established (e.g. Johnston 1998; Helmi & White 1999; Eyre & Binney 2011). The studies above and references therein have presented a simple picture where stars escape from the Lagrange points of the disrupting progenitor. The stripped stars have slightly different energies and angular momenta compared to the parent host, and hence the orbits of the debris and the satellite will diverge in time. More precisely, stars ejected from the outer (at larger Galactocentric distance) Lagrange point have higher energies than the progenitor and hence will have a

longer orbital period and form the trailing tail. Likewise, the leading tail comprises of lower energy stars ejected from the inner (lower Galactocentric distance) Lagrange point. The Lagrange points are typically taken to be equidistant from the progenitor's centre, which is justified if the progenitor is significantly smaller than the scale over which the host gravitational potential changes. For globular clusters, in particular, it is difficult to imagine a situation where this assumption can be broken. Since the Lagrange points are symmetric and assuming the progenitor is roughly spherical, it follows that a similar number of stars should be leaving each Lagrange point, thus yielding comparable number of stars pumped into each tail.

To compare stellar densities at increasing distance intervals along the tails, we have to consider the rate at which the stripped stars move away from the progenitor. This rate is governed by the offsets in energy and angular momentum of the debris which are equal in amplitude and opposite in sign (e.g. Johnston 1998; Helmi & White 1999; Eyre & Binney 2011). Thus, we expect the density in the leading and trailing arms to be roughly symmetric about the progenitor. Note that the exposition above neglects the stretching and compressing of the stream at pericentre and apocentre, respectively. Therefore, the assertion of a symmetric density distribution along the tails only holds when the leading and trailing arm have roughly the same Galactocentric distances. This result was also found in Niederste-Ostholt, Belokurov & Evans (2012) where they used simulations to show that the leading and trailing tails of the Sagittarius dwarf should be symmetric at apocentre and pericentre.

The picture based on Lagrange point stripping not only predicts the stream density, but also provides an expectation for the stream track on the sky and the stream width. If the progenitor is orbiting in a static and spherical potential, the tidal debris will remain in the same plane and an observer in the centre of the galaxy would see the stream which is confined to a great circle on the sky. For a heliocentric observer, the stream track should deviate smoothly from a great circle due to the Galactic parallax. In an aspherical potential, angular momentum is no longer conserved and the orbits of the tidal debris will precess and nutate (e.g. Ibata et al. 2001; Helmi 2004; Johnston et al. 2005; Belokurov et al. 2014; Erkal et al. 2016a). The effect of the aspherical host potential on the tidal debris distribution is two-fold. First, it will cause a further divergence of the stream track from a great circle albeit producing a gradual and smooth deviation. Moreover, given slightly different initial conditions of the stripped stars, they will experience differential orbital plane precession, leading to an increase in the width of the stream on the sky. Thus, in a representative Galaxy, the stream track is expected to deviate slowly from a great circle but have no small-scale features. In addition, the stream width is not expected to stay constant and must evolve smoothly with distance from the progenitor.

4.2 Orbit of Pal 5

In what follows, we are interested in discerning between various mechanisms capable of producing a significant asymmetry in the Pal 5 tails. For the two of these, namely the rotating bar and the GMCs, the exact value of the cluster's pericentre plays an important role. Therefore, while not striving to produce the absolute best fit to the stream observables, we would like to identify the closest match to the cluster's orbit in a realistic host potential.

The radial velocity⁵ of Pal 5 itself was measured in Odenkirchen et al. (2002) to be $v_r = -58.7 \pm 0.2 \text{ km s}^{-1}$. Similar results have been found in Kuzma et al. (2015) with $v_r = -57.4 \pm 0.3 \text{ km s}^{-1}$ (and a systematic uncertainty of 0.8 km s^{-1} in the radial velocity zero-point), as well as Ishigaki et al. (2016) who measured $v_r = -58.1 \pm 0.7 \text{ km s}^{-1}$. The proper motion of the Pal 5 cluster was recently measured in Fritz & Kallivayalil (2015) who give $\mu_\alpha = -2.296 \pm 0.186 \text{ mas yr}^{-1}$ and $\mu_\delta = -2.257 \pm 0.181 \text{ mas yr}^{-1}$. Kuzma et al. (2015) also measured the radial velocity of 17 stars in the leading arm and 30 stars in the trailing arm. Finally, Dotter, Sarajedini & Anderson (2011) measured the distance to Pal 5 using isochrone fitting and found a value of $23.6 \pm 0.9 \text{ kpc}$. Equipped with these measurements, as well as the stream track measured in the previous section and shown in Fig. 5, we can explore the range of pericentric distances within a given potential.

In particular, we consider two choices of the gravitational potential for the Milky Way host. The first one is the `MWPotential2014` model from Bovy (2015) which consists of a spherical NFW halo, a Miyamoto–Nagai disc (Miyamoto & Nagai 1975) and a power-law density bulge with an exponential truncation. Specifically, the NFW halo has a mass of $M_{\text{vir}} = 8 \times 10^{11} M_\odot$, $c = 15.3$ and a scale radius of 16 kpc, the Miyamoto–Nagai disc has a mass of $M = 6.8 \times 10^{10} M_\odot$, $a = 3 \text{ kpc}$ and $b = 280 \text{ pc}$, and the bulge has a mass of $M = 5 \times 10^9 M_\odot$, a power-law exponent of $\alpha = -1.8$ and an exponential truncation radius of 1.9 kpc. The second model is the spherical halo potential described in Pearson et al. (2015), which consists of a spherical logarithmic halo, a Miyamoto–Nagai disc and a Hernquist bulge. The logarithmic halo has a circular velocity of 172.39 km s^{-1} and a scale radius of 12 kpc, the Miyamoto–Nagai disc has a mass of $M = 10^{11} M_\odot$, $a = 6.5 \text{ kpc}$ and $b = 260 \text{ pc}$, and the Hernquist bulge has a mass of $M = 3.4 \times 10^{10} M_\odot$ and a scale radius of 700 pc. For our coordinates, we have the X-axis pointing towards the Galactic centre, Y is aligned with the Galactic rotation and Z pointing towards the Northern Galactic pole. Following a combination of Schönrich, Binney & Dehnen (2010) and Bovy et al. (2012), we take the Sun's velocity relative to the local circular velocity to be $(U_\odot, V_\odot, W_\odot) = (11.1, 26, 7.3) \text{ km s}^{-1}$ and place the Sun at $(-8.3, 0, 0) \text{ kpc}$. Within each potential, we then compute the tangential motion of the Sun by adding V_\odot to the circular velocity, V_c , at the Sun's location. In the `MWPotential2014` potential from Bovy (2015) we find $V_c = 219.0 \text{ km s}^{-1}$, and in the spherical halo potential from Pearson et al. (2015) we find $V_c = 220.8 \text{ km s}^{-1}$.

To zoom-in on to the most appropriate orbit in these potentials, we sample the proper motion of the Pal 5 cluster from the allowed range of the observed values and generate model streams using the mLCS method described in Gibbons et al. (2014). The progenitor is modelled as a $2 \times 10^4 M_\odot$ Plummer sphere with a scale radius of 15 pc. Particles are released from the Lagrange points with a velocity dispersion given by the Plummer profile. The particles are stripped near each pericentre following a Gaussian stripping rate with $\sigma = 10 \text{ Myr}$. At the present epoch, the progenitor is fixed to be at a distance of 23.6 kpc and a radial velocity of -57.4 km s^{-1} . Given the chosen proper motion value, we then rewind the cluster's orbit for 5 Gyr and produce a stream. With the suite of stream simulations in hand, we explore how the goodness of the model fit depends on Pal 5's pericentric distance.

In order to assess how well each realization fits the data, we define a likelihood for the stream track and the run of radial velocities. For

⁵ Note that here and throughout the rest of this paper, the radial velocities are always measured in the heliocentric frame.

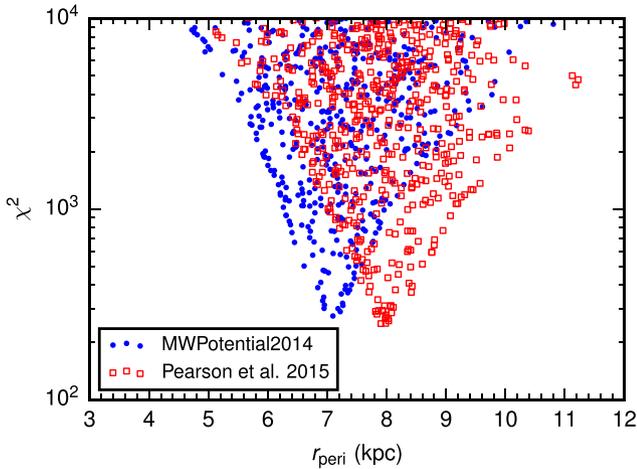


Figure 6. Goodness-of-fit χ^2 as a function of r_{peri} for Pal 5 stream models with the cluster proper motions sampled near the best fit within each potential. The filled blue circles show the result for the `MWPotential2014` potential from Bovy (2015). The empty red squares show the result for the potential from Pearson et al. (2015) with a spherical halo. As evidenced from this plot, there is a substantial uncertainty in the pericentric radius of the Pal 5 cluster due to the uncertainty in the MW potential.

the stream track, we perform a linear fit to the simulated stream in bins of 1° in ϕ_1 and determine the value and uncertainty of ϕ_2 in the centre of each bin. For the data, we take the mean track with uncertainties from Section 3.5. The likelihood for the track is then defined by

$$\mathcal{L} = \prod_i \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{(d_i - m_i)^2}{2\sigma_i^2}\right), \quad (2)$$

where the index i runs over the ϕ_1 bins, d_i is the measured ϕ_2 from the data, m_i is the mean of ϕ_2 for the model and σ_i is the sum in quadrature of the observational and model error. The likelihood is similarly defined for the radial velocity except the index i now runs over the radial velocity data points. For each data point, a Gaussian fit is performed for the model points within 0.5 in ϕ_1 to get the mean and error on the mean of the radial velocity. The likelihoods for the stream track and radial velocity are then multiplied to get the total likelihood.

For the `MWPotential2014` model from Bovy (2015), we find a best-fitting proper motion of $\mu_\alpha = -2.23 \text{ mas yr}^{-1}$ and $\mu_\delta = -2.22 \text{ mas yr}^{-1}$. This proper motion gives a pericentre of 7.1 kpc. For the model from Pearson et al. (2015), we find a best-fitting proper motion of $\mu_\alpha = -2.30 \text{ mas yr}^{-1}$ and $\mu_\delta = -2.29 \text{ mas yr}^{-1}$ which gives a pericentre of 8.0 kpc. Note that this is slightly different than the best-fitting proper motion reported in Pearson et al. (2015), in part due to our use of a different cluster’s radial velocity and solar motion, and in part due to our updated stream track. In order to showcase the range of possible pericentric distances in these potentials, the proper motion is sampled 1000 times around the best-fitting values with a spread of 0.1 mas yr^{-1} . Fig. 6 displays the χ^2 of these models as a function of the pericentric distance. We see that within a given potential, the range of pericentres consistent with the stream data available is quite small. However, the systematic error, i.e. the uncertainty due to the choice of the potential, is large. Namely, the allowed range of pericentres for Pal 5 is from 7 to 8 kpc. Further modelling of the Pal 5 disruption with a more flexible potential is needed to produce a more robust measurement of the cluster’s pericentre.

4.3 Model of Pal 5 in a static and smooth potential

Our fiducial model of the Pal 5 disruption is shown in Fig. 7. This simulation is run using the N -body part of `GADGET-3` which is similar to `GADGET-2` (Springel 2005). For the potential, we use `MWPotential2014` from Bovy (2015) which satisfies a wide range of constraints, as opposed to the potential model of Pearson et al. (2015) which was selected specifically to reproduce the Pal 5 stream. In this potential, the globular is represented with a King profile with a mass of $2 \times 10^4 M_\odot$, a core radius of 15 pc and $w = 2$. It is modelled with 10^5 equal mass particles and a softening length of 1 pc. The final phase-space coordinates of the progenitor are as follows: a proper motion of $(\mu_\alpha, \mu_\delta) = (-2.235, -2.228) \text{ mas yr}^{-1}$, a radial velocity of $v_r = -57.4 \text{ km s}^{-1}$ and a distance of 23.6 kpc. From the position and velocity of the cluster today, the orbit is rewound for 5 Gyr and the disruption is initiated.

Fig. 7 is a clear demonstration of the expectations of a typical stream behaviour as described in Section 4.1. The figure shows the same set of stream properties as in Fig. 5 as seen by a heliocentric observer, as well as the run of the stream’s line-of-sight velocity and Galactocentric distance with ϕ_1 . As expected, we see that the debris density is broadly symmetric about the progenitor with no significant small-scale features except for the epicycles, in stark contrast to the data from Pal 5 which is overplotted. As demonstrated in the figure, the centroid of the debris smoothly deviates from a great circle in a manner very similar to that of the observed stream track although the match is not perfect. Reassuringly, the simulated stream has a width similar to that observed in Pal 5. The fourth panel of the figure (second from the bottom) shows the radial velocity variation along the stream in comparison to the measurements provided by Kuzma et al. (2015). Finally, the fifth (bottom) panel gives the Galactocentric distance evolution along the tails. Note that, in the vicinity of the progenitor, the leading and trailing arms have very similar distances, thus demonstrating that the stretching and compressing of the stream cannot be responsible for the asymmetry discussed in Section 3.5. This is in agreement with Ibata et al. (2016) who also find similar distances to the leading and trailing arms, with a small variation of ~ 3 kpc along the observed portion of the stream. To conclude, overall, the fiducial model faithfully reproduces the behaviour of the centroid of each of the tails in the phase space, but not the details of the stellar density along the stream.

4.4 Epicyclic overdensities

The simple picture of tidal disruption described in Section 4.1 glosses over the small-scale details of the stellar stream formation, namely the density variations due to epicyclic motion (Küpper et al. 2008; Just et al. 2009; Küpper et al. 2010). As mentioned before, these overdensities arise because the initial conditions of the stripped stars are similar and they complete the motion around the epicyclic ellipse at comparable times. The similarity of the perturbed motion of the debris (compared to the stream track) is emphasized at the vertices of the ellipse where the stars spend the most time. In the fiducial stream displayed in Fig. 7, the epicycles are included self-consistently. The locations of the two most prominent epicyclic overdensities are marked with arrows and as discussed in Küpper et al. (2008), Just et al. (2009) and Küpper et al. (2010), appear equally spaced from the progenitor.

Epicyclic feathering is unambiguously abundant in the numerical simulations of the globular cluster disruption referenced above. Is it surprising that the overdensities appear so underwhelming in Fig. 7?

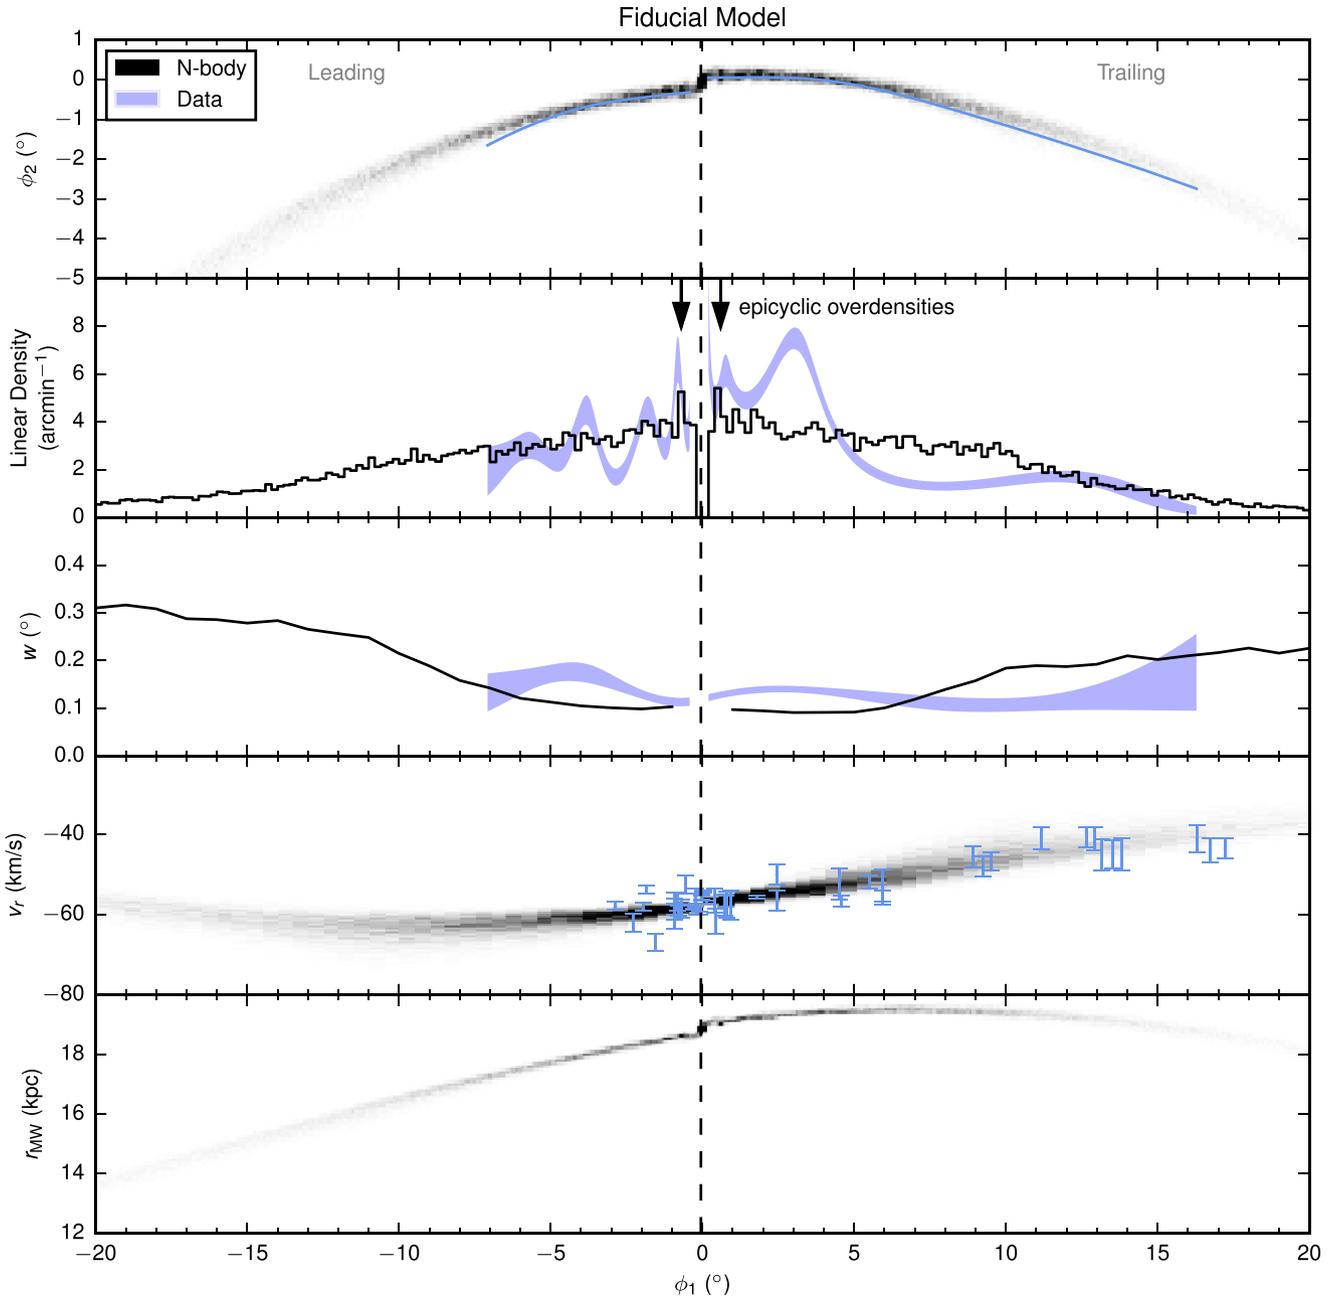


Figure 7. Comparison between measurements of the Pal 5 stream and the fiducial (unperturbed) N -body model of the cluster disruption. Similar to Fig. 5, this shows the evolution of the stream centroid (top), density of particles (second panel) and the width of the debris distribution (third panel). Additionally, the stream’s radial velocity (fourth panel) and the Galactocentric radius (fifth panel) are shown. In all panels, the black histograms show the results of our N -body simulations. In the top panel, the solid blue line shows the median from Fig. 5. In the second and third panels, the filled blue region shows the 16–84 per cent credible interval from Fig. 5. In the fourth panel, the observed radial velocities from Kuzma et al. (2015) are shown as blue error bars. The vertical dashed line shows the location of the progenitor at $\phi_1 = -0.05$ which separates the leading (left) and trailing (right) tails of the stream. While the simulated stream track is remarkably close to the observed track, the simulated density is broadly symmetric, as expected, and in stark contrast to the observed Pal 5 density. The two arrows in the second panel mark the approximate locations of the two most prevalent epicycles at $\phi_1 = -0.7, +0.6$. This slight asymmetry is due to the fact that the simulated progenitor is located at $\phi_1 = -0.05$. The epicycle in the leading arm also lines up well an overdensity in the data at roughly the same location, $\phi_1 = -0.8$. There is also a reasonable match in the trailing arm albeit the measured overdensity is not as significant. Note that the second panel, the linear density of particles in the simulation has been scaled down by a factor of 8.3 to approximately match the observed linear density in the leading arm.

To answer this question, let us follow the evolution of the epicyclic overdensities as a function of the satellite’s orbital phase. Fig. 8 presents the debris density at five different times, starting from the previous pericentre and ending at the present. Note that this plot shows the stream star counts as viewed from the Galactic centre,

hence we denote the angle along the stream as θ_{GC} rather than ϕ_1 . As evidenced in the figure, the clumping is most visible at pericentre and is barely detectable at apocentre. Some of this behaviour might be due to the stretching and compressing of the stream as it goes from pericentre to apocentre. However, if the stream is thought of

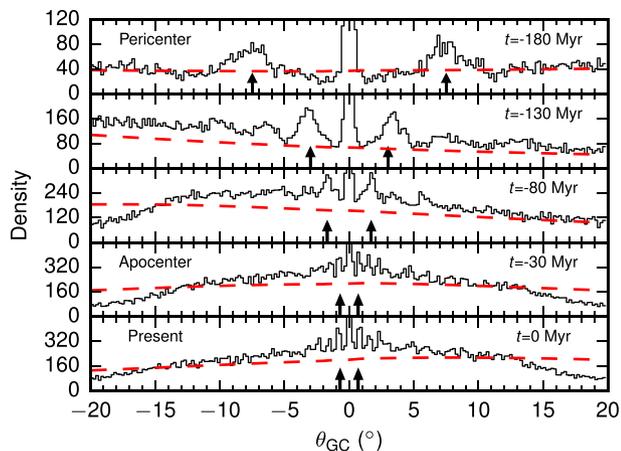


Figure 8. Evolution of the stream density in the fiducial (unperturbed) Pal 5 simulation at various progenitor’s orbital phases for an observer at the Galactic centre. The solid black histogram shows the linear density computed in 0.2° bins. The dashed red curve is proportional to r^2 and has a maximum equal to the median of the density. At each time, the stream plane is defined by the angular momentum of the progenitor. The most pronounced epicyclic overdensities can be seen at the pericentre (top panel) where they appear to be highly symmetric. The second and third panels show the density at intermediate times between pericentre and apocentre. The density is asymmetric at these epochs due to the differing radii of the leading and trailing arms. This asymmetry roughly follows the expected scaling of r^2 . Note that the epicyclic overdensities become less pronounced as the progenitor approaches apocentre. The fourth panel shows the stream at apocentre where it is broadly symmetric and the epicyclic overdensities are barely visible. The final panel shows the stream at the present time where the tails are slightly asymmetric and there are two epicyclic overdensities visible at $\theta_{\text{GC}} = \pm 0.7^\circ$.

as a train of particles following the same orbit, then conservation of the angular momentum dictates that $r^2 \frac{d\theta}{dt}$ is constant in time, where $d\theta$ and dt are, respectively, the angle and the time delay between the two particles. The angle between two particles trailing each other will thus vary as $d\theta \propto r^{-2}$. This implies that if the growth of the stream is ignored, the debris density goes like r^2 . Thus, the angular distances between the stars in the epicyclic clumps are expected to compress by a factor of $(r_{\text{peri}}/r_{\text{apo}})^2$ on going from pericentre to apocentre. Similarly, we expect their density to go up by a factor of $(r_{\text{apo}}/r_{\text{peri}})^2$. Perhaps, the strength of the epicyclic overdensities at pericentre is simply due to the fact that most of the stripping actually happens at pericentre. Given that the clumps form at integer multiples of the radial period after they are stripped, their amplitudes are pronounced near peri crossings. This is illustrated in the top panel of Fig. 8 where the strong epicyclic bunchings are composed of the particles stripped during the previous pericentric passage.

Note that some of the panels in Fig. 8 exhibit a small asymmetry in the density between the leading and trailing arm. Also included in the figure is a dashed red curve that is proportional to r^2 as measured in 1° bins along the stream. The slope of the curve reveals that the asymmetry seen in the second and third panels of Fig. 8 is mostly due to the difference in the Galactocentric radius along the stream. Most importantly, we see that at the present day (bottom panel), the debris density near the progenitor is roughly symmetric since the two tails are at approximately the same distance from the Galactic centre. The small asymmetry discussed above agrees well with the findings of Just et al. (2009) and Küpper et al. (2010) who both

note that the amplitudes of the epicyclic overdensities could differ slightly. We note, however, that this difference is nowhere near the dramatic mismatch between the trailing and leading tail densities as measured here. Finally, Zotos (2015) found some evidence that the rate at which the stars leave the inner and outer Lagrange points can be different. Yet, again, the reported asymmetry, if real, is at much lower level, i.e. some 10 per cent rather than a factor of 2 as detected in the Pal 5 stream.

5 MECHANISMS TO PRODUCE OBSERVED FEATURES

Having established the expected shape and density behaviour of a Pal 5 stream in a smooth and static potential, this section considers two distinct mechanisms that can plausibly produce the observed discrepant features. Namely, a fly-by of substructure, either in the form of a DM subhalo or a GMC, and the effect of the Milky Way bar.

5.1 Interaction with subhaloes

DM subhaloes affect streams locally (e.g. Ibata et al. 2002; Johnston et al. 2002) and hence ought to be able to induce the features discussed above. In fact, models of subhalo fly-bys generically predict an underdense region (around the point of closest approach) surrounded by overdensities, (e.g. Carlberg 2012; Erkal & Belokurov 2015a,b), similar to what can be seen in the trailing tail of Pal 5. Thus, interactions with dark perturbers are an obvious candidate mechanism, which is considered in this subsection.

To demonstrate that perturbations caused by DM subhaloes can look very similar to the features detected, we run a suite of semi-analytic simulations of subhalo fly-bys. We start with the fiducial Pal 5 stream described in Section 4.3 and shown in Fig. 7. Using the effective N -body model described in Erkal et al. (2016b), we can rapidly simulate the effect of a fly-by by taking a snapshot of the simulation at an earlier time and then perturbing the particles using the form of the impulse approximation of Sanders et al. (2016). As in Erkal et al. (2016b), the subhaloes are assumed to be Plummer spheres. The perturbed particles are then evolved using a kick-drift-kick leapfrog integrator to the end of the simulation where they are combined with the unperturbed particles.

With a small amount of trial and error, we found a two fly-by set-up whose effect roughly matches the observed stream density. The first fly-by, which creates the feature in the trailing tail, occurs 1.4 Gyr in the past with a mass of $5 \times 10^7 M_\odot$, a scale radius of 1.15 pc and a velocity of 200 km s^{-1} relative to the stream. The fly-by occurs 5.7 kpc from the progenitor, roughly 45 per cent of the way along the trailing stream at that time, and has an impact parameter of 1 kpc. The second fly-by, which creates the feature in the leading tail, occurs 500 Myr in the past with a mass of $10^6 M_\odot$, a scale radius of 162 pc and a velocity of 100 km s^{-1} relative to the stream. The fly-by occurs 2.0 kpc from the progenitor, roughly 17 per cent of the way along the leading stream at that time, and is a direct impact. Interestingly, almost all of the simulated particles within the gap at $\phi_1 \sim -3^\circ$ were stripped within the last 2 Gyr indicating that the fly-by must have occurred within that time period.

Fig. 9 compares the realization of Pal 5 which has interacted with both subhaloes to the stream data and shows a qualitative match to the density profile. In addition, the width in the trailing arm is now almost constant until $\phi_1 \sim 10^\circ$. This provides a better match to the observed width, unlike the unperturbed stream (see Fig. 7), which has fanned appreciably by then. Our search for a

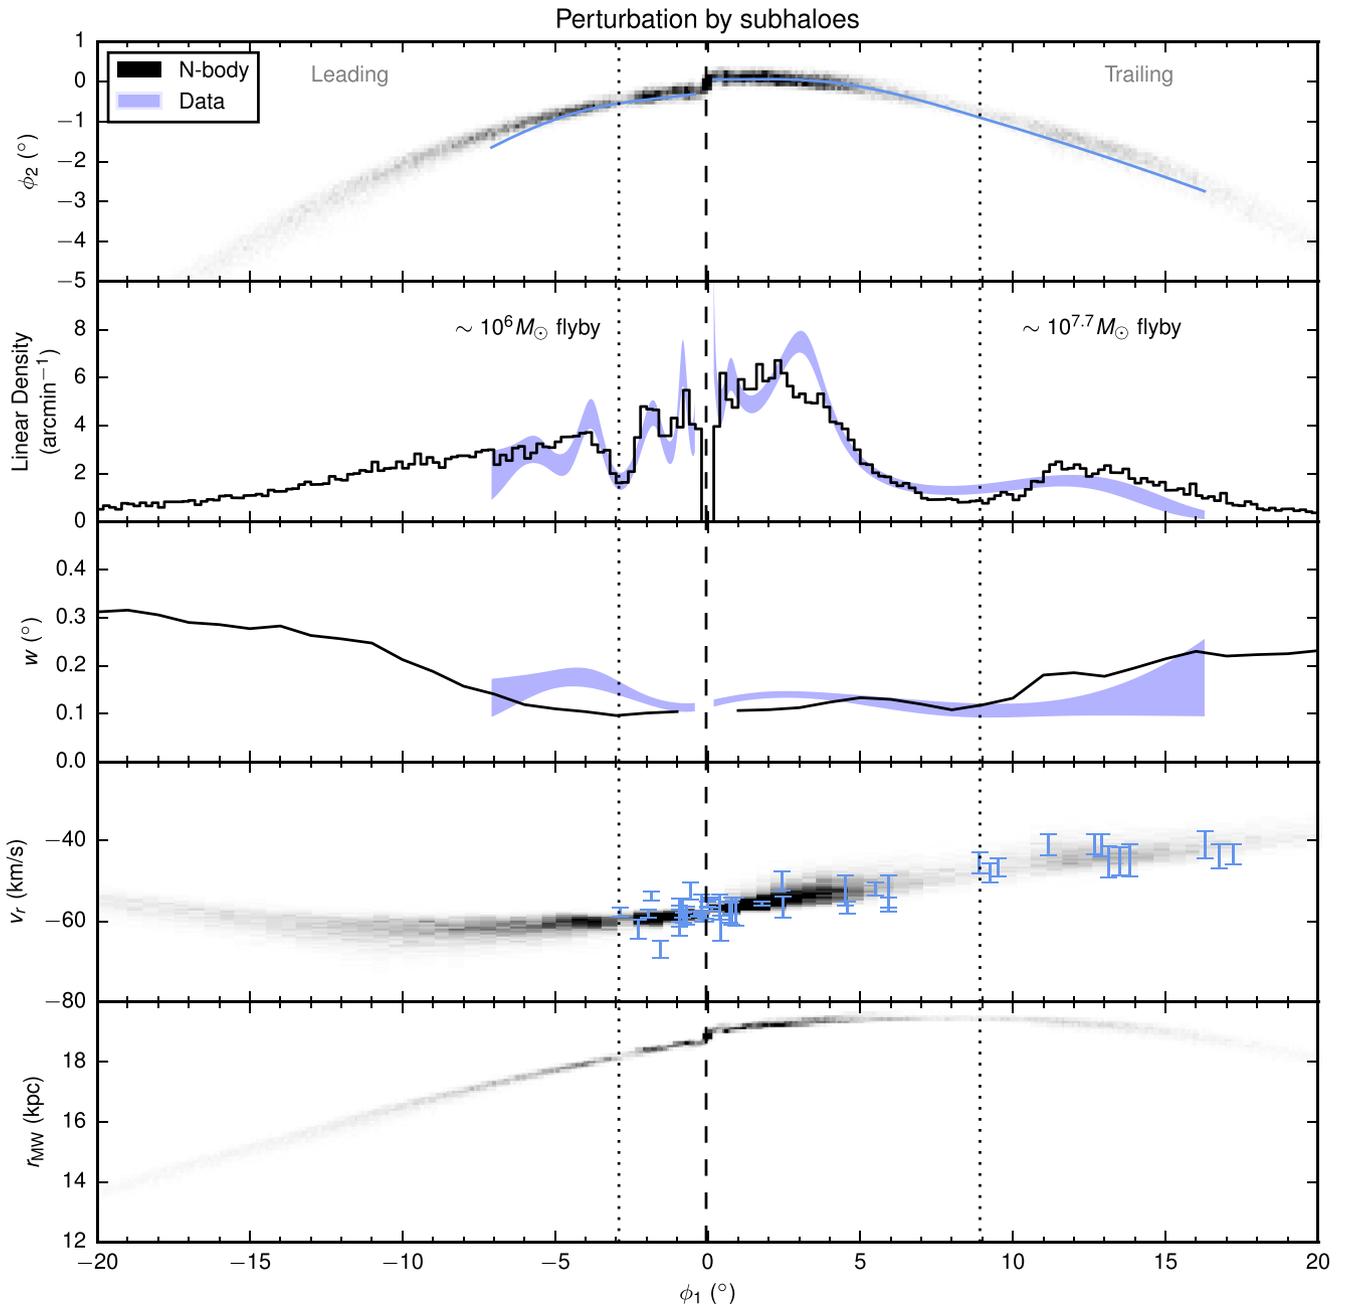


Figure 9. Comparison between the measured stream properties and the model of the Pal 5 stream in the MWPotential2014 potential from Bovy (2015) perturbed by two subhalo fly-bys. As in Fig. 7, the stream track, debris density, stream width, radial velocity and Galactocentric radius are shown. Details of the panels are described in the caption of Fig. 7. The vertical dotted lines show the approximate locations of the centres of the two gaps produced as a result of interaction with DM subhaloes. The gap in the leading arm (left) is created by a $10^6 M_{\odot}$ subhalo while the gap in the trailing arm (right) is created by a $5 \times 10^7 M_{\odot}$ subhalo. In an apparent contrast to Fig. 7, there is a good match between the measured and modelled stream density. In addition, the trailing tail now has a roughly constant width until $\phi_1 \sim 10^\circ$, providing a better match to the observations.

convincing DM fly-by configuration has been far from exhaustive. The quick success of the parameter exploration exercise is easy to understand. As discussed in Erkal & Belokurov (2015b), there is a large degeneracy in going from the shape of the gap to the subhalo properties and hence a wide range of fly-bys can give rise to the same density profile. Note, however, that this degeneracy is almost entirely broken by looking at other observables, such as the stream track and the radial velocity profile. Yet, before a more comprehensive modelling of the Pal 5 stream is carried out, we must

stress that the subhalo properties presented above are merely chosen by hand to generate an approximately similar density profile.

Because the debris gap properties are solely controlled by the velocity kick imparted on the stream stars, similar sized kicks can be produced by a slowly moving low-mass perturber nearby and a fast-flying massive satellite farther away (see Erkal & Belokurov 2015a). Therefore, it is prudent to consider the two most obvious candidate perturbers in the Milky Way: the Sagittarius dwarf galaxy and the Large Magellanic Cloud. Both have in-fall masses in excess

of $10^{10} M_{\odot}$ (see e.g. Peñarrubia et al. 2016; Gibbons, Belokurov & Evans 2017) and therefore may have affected a number of objects throughout the Galaxy. However, the size of the gap induced in the stream is also related to the impact parameter (see Erkal & Belokurov 2015a), so a distant fly-by, $\gtrsim 10$ kpc, will produce gaps significantly larger than even the 9° feature seen in Fig. 5. Thus, unless Sagittarius had a passage significantly closer than this, or the gap in Pal 5 is actually much larger than $\sim 9^{\circ}$, it is unlikely that either could produce the gap-like features seen in Pal 5. As an additional argument, Erkal et al. (2016b) studied the gaps created in a Pal 5-like stream by subhaloes in the range 10^5 – $10^9 M_{\odot}$ and found that subhaloes in the range 10^8 – $10^9 M_{\odot}$ made a significantly lower contribution to the number of gaps than subhaloes in the range 10^6 – $10^8 M_{\odot}$ due to their low rate of encountering the stream. We note that while these massive perturbers are unlikely to have caused the features seen in Pal 5, the streams of DM stripped from the massive satellites can also produce gaps as argued in Bovy (2016).

In addition to creating features in the density, the perturbations from subhalo fly-bys also cause the stream track to oscillate about the unperturbed track (Erkal & Belokurov 2015a,b). This oscillation takes place over roughly the same scale as the gap size with an amplitude governed by the properties of the fly-by (e.g. Erkal & Belokurov 2015b). For the fly-by near the trailing arm shown in Fig. 9, the amplitude of the oscillation is 0.5 and, by chance, the oscillation happens to be close to zero at the present time. Thus, the lack of a stream track variation does not necessarily imply there was no subhalo perturbation. Furthermore, in order to detect such a small oscillation of ~ 0.5 over the gap size, $\sim 10^{\circ}$, one would need an accurate model of the Milky Way and the unperturbed stream track. Thus, the lack of an obvious stream track oscillation in the data does not rule out a subhalo fly-by.

5.2 Interaction with giant molecular clouds

Recent work by Amorisco et al. (2016) has also found that the Pal 5 stream could be affected by GMCs. They found that since Pal 5 is on a prograde orbit, these gaps are somewhat enhanced compared to those produced by subhaloes due to the smaller relative velocity between the stream and the GMCs. However, due to the lower mass of the GMCs, $M < 10^7 M_{\odot}$, the gaps they produce span a smaller range of sizes. The distribution of gap sizes is explored in fig. 4 of Amorisco et al. (2016), where they find that deep gaps have sizes between 0.4 and 3 kpc. At a distance of 23 kpc, this would correspond to 1° – 7.5° . We also note that since Amorisco et al. (2016) did not require their Pal 5-analogues to be near apocentre at the present, the gap sizes they predict are likely an overestimate. Thus, while the 2° feature in the leading arm at $\phi_1 \sim -3^{\circ}$ is consistent with a GMC perturbation, the larger feature at $\phi_1 \sim 8^{\circ}$ with a size of 9° is unlikely to have been caused by a GMC. This also agrees with the result of Section 5.1 where we found that a subhalo mass of $10^6 M_{\odot}$ could explain the feature in the leading arm.

While we cannot distinguish between the effects of a DM subhalo and a GMC with the current data, in Erkal & Belokurov (2015b) it was shown that the properties of the fly-by, notably the time since impact, can be extracted from the shape and the density of the stream perturbation. Thus, with additional measurements of the radial velocity, and, perhaps, proper motion, one should be able to reconstruct the fly-by and identify the likely perturber. This should be possible because interactions with GMCs only occur in the disc, unlike those with DM subhaloes that can occur anywhere in the halo.

5.3 Rotating bar

An interaction with an intervening satellite provides a short time-scale change of the local gravitational potential and thus can affect only a portion of a long stellar stream. Similarly, a rotating Milky Way bar adds a varying component to the force field of the Galaxy. Hattori et al. (2016) studied the consequences of the presence of a bar on the Ophiuchus stream (Bernard et al. 2014) and found a dramatic effect where the bar can induce different changes in the energy to different sections of a stream. Thus, since these sections will now orbit with different periods, the bar can cause variations in the debris distribution, resulting in both under and overdensities. The effect of the bar on streams was also considered in Price-Whelan et al. (2016b) where they focused on the importance of orbital chaos, a topic that will be briefly discussed in Section 6.

Here, since we expect that the bar can only induce relatively large-scale features compared to the \sim few degree scale features possible from substructure, we focus solely on the asymmetry in the density near the progenitor to ascertain whether the bar can plausibly create this. The Galaxy is represented with the NFW halo and Miyamoto–Nagai disc of `MWPotential2014` from Bovy (2015) and the bar is the prolate bar model of Long & Murali (1992) described in Hattori et al. (2016). Specifically, we use a mass of $M = 5 \times 10^9 M_{\odot}$, a half-mass size of 3 kpc and a Plummer softening of 1 kpc. These values give a bar that is broadly consistent with the mass constraints reported in Portail et al. (2015). The current bar angle is taken to be -30° to match observed constraints (e.g. López-Corredoira, Cabrera-Lavers & Gerhard 2005). In order to explore the broad enough range of bar effects, we consider three different pattern speeds and a non-rotating bar. We use pattern speeds of $\Omega_{\text{bar}} = -30, -50, -70 \text{ km s}^{-1} \text{ kpc}^{-1}$ which span the gamut consistent with the Milky Way’s bar observations (see Gerhard 2011, for a review). In this potential, 1000 model streams are evolved using the mLCS technique described in Section 4.2. As in Section 4.2, the current distance and radial velocity of Pal 5 are fixed at 23.6 kpc and -57.4 km s^{-1} . The proper motions are sampled from a normal distribution centred on $(\mu_{\alpha}, \mu_{\delta}) = (-2.23, -2.22) \text{ mas yr}^{-1}$ with a spread of 0.1 mas yr^{-1} in each component. Each simulated stream is evolved for 5 Gyr and the minimum pericentric distance during this time is recorded. To quantify the bar’s influence on the stream, the asymmetry statistic, A , describing the density difference between the leading and trailing arm is calculated. The asymmetry statistic is defined as follows:

$$A = \sum_i (\rho_i^{\text{lead}} - \rho_i^{\text{trail}})^2, \quad (3)$$

where the index i runs over bins in ϕ_1 between 1° and 5° in steps of 0.5° , and $\rho_i^{\text{lead}}, \rho_i^{\text{trail}}$ are the normalized densities in the leading and trailing arm, respectively. For the observed Pal 5 density, a value of $A = 0.059$ is measured which can now be compared against the asymmetry seen in each model realization of the stream. Of course, this asymmetry statistic may overestimate the impact of the bar alone since the small-scale feature in the data near $\phi_1 \sim -3^{\circ}$ will contribute to the asymmetry but was likely created by a different mechanism.

Fig. 10 shows this asymmetry versus the pericentric distance for three different bar pattern speeds and for a non-rotating bar. As evidenced in the figure, a rotating bar can create a substantial asymmetry in the stream even if the pericentre is sufficiently large, i.e. ~ 7 – 8 kpc. The exact location of the maximal asymmetry depends on the pattern speed of the bar. The figure also shows that a non-rotating bar can only create symmetric streams. Thus, it is the

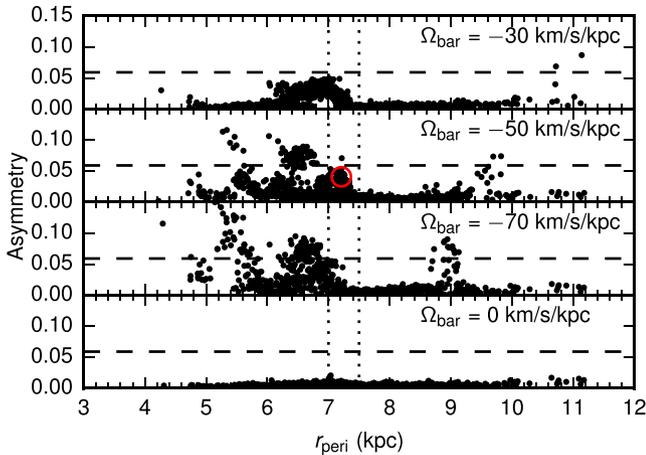


Figure 10. Asymmetry statistic, A (see main text for details), for a sample of streams with varying proper motion in the presence of a rotating bar with three different pattern speeds as a function of Pal 5’s pericentre. The horizontal dashed line shows the observed asymmetry in Pal 5 of $A = 0.059$. Interestingly, the bar can produce a similar amount of asymmetry as in the observed Pal 5’s tails even if its orbit has a pericentre larger than 7 kpc. The red circle in the top panel shows the stream we have resimulated with GADGET-3 (see Fig. 11). This particular point was chosen since the stream track and radial velocity also provide a good match to that of Pal 5. The dotted vertical lines show the range of pericentric radii where we get the best fits to the Pal 5 stream.

rotation of the bar that is crucial, as expected from the earlier results of Hattori et al. (2016). As a test, we also ran the streams in a bar with a present-day angle of -20° and found a similar number of asymmetric streams. Note that, importantly, while a reasonable fraction of the model streams studied here exhibit a substantial asymmetry, many of these do not provide a good match to the stream track and the radial velocity profile of Pal 5.

As a further piece of evidence, one of the streams shown in Fig. 10 (the red circle) is re-simulated with an N -body disruption. This particular stream was chosen since it has a significant asymmetry, and the stream track and radial velocity run are a relatively good fit to Pal 5. The simulation set-up is the same as in Section 4.3, except the bulge has been replaced by the rotating bar described above. The resulting stream properties are shown in Fig. 11. The density has an asymmetry of $A = 0.041$, ~ 30 per cent less than that in Pal 5 and the perturbation is slightly off-set along ϕ_1 compared to what is observed. However, overall, the size and amplitude of the feature produced are similar to what is seen in the trailing arm. During the progenitor’s orbit, its minimum pericentric distance is 7.1 kpc, confirming the result of Fig. 10 that the bar is important as such large distances. Finally, we note that this asymmetry is extremely sensitive to the pattern speed. For this particular set of final velocities, the asymmetry is erased if the pattern speed is changed by just $0.5 \text{ km s}^{-1} \text{ kpc}^{-1}$. This is likely because the asymmetry depends on a precise alignment of the bar and the stream’s pericentre at an earlier time, as was the case in Hattori et al. (2016). Thus, while we have demonstrated that the bar can in principle create the density asymmetry seen in the Pal 5 tails, the uncertainty in both the pattern speed and the orbit of Pal 5 means we cannot be sure that the bar is the culprit for the feature in the trailing arm. Future work is needed to determine whether the effect of the bar can produce the precise features seen in Pal 5.

Note that only a single bar model is examined in detail in this work, i.e. that described in Long & Murali (1992). As noted in

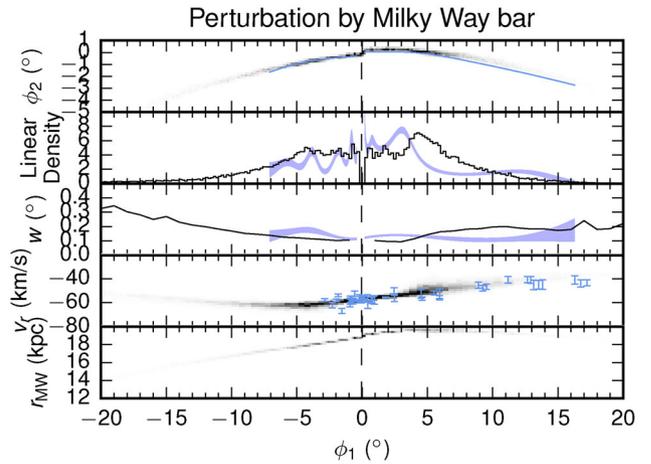


Figure 11. Influence of a rotating bar on Pal 5’s tidal debris. The panels, respectively, show the stream track, density, width, radial velocity and Galactocentric radius of a Pal 5 realization in a potential with $\Omega_{\text{bar}} = -50 \text{ km s}^{-1} \text{ kpc}^{-1}$. See the caption of Fig. 7 for more details on the panels. This realization has an asymmetry of $A = 0.041$ and demonstrates that the bar can produce significant asymmetries in the stream density. This particular realization corresponds to the red circle in Fig. 10 with a pericentre of 7.1 kpc. We note that the asymmetry is extremely sensitive to the pattern speed and a change of only $0.5 \text{ km s}^{-1} \text{ kpc}^{-1}$ will yield a different stream perturbation with a significantly different value of A .

Hattori et al. (2016), this bar is slightly lighter than the constraints in Portail et al. (2015). Thus, it is possible that we have somewhat underestimated the effect of the bar. In addition, it is possible that bars obeying different density laws may produce different features in the stream. Furthermore, note that the effect of the bar on the Pal 5 cluster itself was considered in Allen, Moreno & Pichardo (2006) where they studied a single realization of Pal 5’s orbit and found that their version of the rotating bar did not substantially affect it. However, as was shown above and in Hattori et al. (2016), the presence of a bar induces slight changes in the orbital periods of stars in a stream and hence can create structure while not dramatically altering the orbit of the progenitor.

6 OTHER MECHANISMS

This section considers other mechanisms that could lead to the features seen in the density of the Pal 5 tidal tails, some of which can be ruled out. In addition, several routes are suggested to help distinguish between the plausible mechanisms.

6.1 Rotating Pal 5

Observations of globular clusters around the Milky Way suggest that a large fraction of them may possess internal rotation (e.g. Bellazzini et al. 2012; Fabricius et al. 2014). Non-zero net angular momentum will affect the velocities with which stars leave the globular cluster and therefore can naturally affect the stream properties. The rate at which particles move away from the progenitor is controlled by their energy and hence mainly by the component of their velocity aligned with the progenitor’s systemic velocity. As long as the progenitor remains roughly spherical, the stars at the Lagrange points will receive opposite but equal boosts of their velocity from the rotation. Thus, they will move away from the progenitor at the same rate, maintaining a symmetric stream density. In addition, this mechanism should not create significant small-scale features in

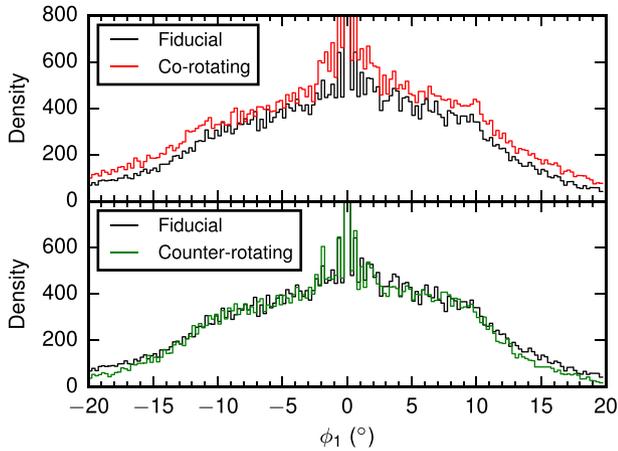


Figure 12. Stream density for rotating progenitors. Comparison of the tidal tail density for a non-rotating (black) Pal 5 with a corotating (red, top) and counter-rotating (green, bottom) Pal 5. As expected, the non-rotating Pal 5 has an intermediate extent with the corotating (counter-rotating) being more (less) extended. More importantly, however, the rotation does not lead to any significant asymmetry in the debris distribution or significant small-scale density features.

the stream. According to this picture, if the progenitor’s rotation is aligned with the orbital angular momentum, the debris will spread out at a faster rate. Likewise, if the progenitor’s rotation is anti-aligned, we should expect a shorter stream.

In order to study this effect, we ran a simulation of a rotating globular cluster based on the unperturbed Pal 5 model described in Section 4.3 and presented in Fig. 7. The particle initial conditions in the fiducial simulation depend only on the magnitude of the velocity relative to the cluster’s centre of mass. Thus, we can flip the sign of any individual velocity and still have a stable system. In order to have a net rotation around a given direction, we compute the component of the angular momentum along this axis and require each particle to have a positive angular momentum in this direction. If the angular momentum is negative, we simply flip the sign of the particle’s velocity. This results in a spherical and stable progenitor with a net rotation. Fig. 12 presents the stream density at the present day for a non-rotating Pal 5, a Pal 5 whose internal rotation is aligned with the initial orbital angular momentum (corotating), and a Pal 5 whose internal rotation is anti-aligned with the initial orbital angular momentum (counter-rotating). As expected, the corotating realization produces the most extended stream and the counter-rotating realization produces the shortest debris distribution. Critically, however, the cluster rotation does not introduce a significant density mismatch at levels similar to those measured in Pal 5 tails or any significant small-scale features. Thus, it appears that rotation cannot be responsible for the observed features.

6.2 Chaos

Another important mechanism worth considering is chaos. Pearson et al. (2015) simulated Pal 5 in the triaxial potential of Law & Majewski (2010) and found that the cluster’s tidal debris experienced a large amount of dispersal, which resulted in a dramatic drop in surface density along the stream. The exact nature of this stream fanning is uncertain. However, the same potential was studied in more detail in Price-Whelan et al. (2016a) who found the presence of weak chaos in the orbital behaviour and concluded that even weak chaos could lead to substantial stream fanning. Note that

at least part of this fanning is due to differential orbital plane precession that occurs in axisymmetric and triaxial potentials (Erkal et al. 2016a). However, this effect treats each tail equally and hence cannot create asymmetries. We also note that the most dispersed stream models, e.g. fig. 5 of Pearson et al. (2015), look very different from the observed Pal 5 tails. Namely, the centroid of the simulated debris distribution is dramatically misaligned with respect to the observed stream track and the tails display a rapidly increasing width (inconsistent with our measurement presented above). Moreover, the simulation in the triaxial potential cannot reproduce the run of the radial velocity along the stream. It remains unclear if a potential that was mildly chaotic could ameliorate these discrepancies and also exhibit an asymmetry in the density. We conclude that further investigation is needed to understand the importance of chaos for the Pal 5 stream although it does not appear to be a likely mechanism.

6.3 Other baryonic effects

As far as the localized stream perturbations are concerned, in addition to DM subhaloes and GMCs, there are several other baryonic substructures that could potential re-shape the Pal 5 tails. Three obvious candidates are globular clusters, the disc of the Milky Way and the spiral arms of the Milky Way.

In order to estimate the importance of interactions with (other) globular clusters, we can compare their numbers and masses to those of subhaloes and use the formalism of Erkal et al. (2016b). Using the table of globular cluster properties around the Milky Way from Gnedin & Ostriker (1997), we find there are three globular clusters in the mass range $10^6 M_{\odot} < M < 10^7 M_{\odot}$ with Galactocentric radii between 7 and 20 kpc (e.g. spanning the range of Pal 5’s orbit), with the most massive being $1.45 \times 10^6 M_{\odot}$. In contrast, on average there are expected to be roughly 10 subhaloes in the same mass and radial range for Milky Way mass hosts (e.g. Erkal et al. 2016b, and references therein). If we assume that the three globular clusters seen at present are representative of the average number of globular clusters in this radial range, there will be roughly three times as many impacts by subhaloes as by globular clusters in this mass range. Furthermore, many of the subhaloes will have substantially higher masses so we expect the subhaloes to produce significantly more prominent gaps. Thus, we conclude that subhaloes are more likely to have created gaps in Pal 5. However, future orbital analysis of the most massive globular clusters will also shed light on whether they interacted with Pal 5.

The disc itself cannot create an asymmetry since each tail will pass through nearly the same region of the disc and will experience the same forces. This was also demonstrated in Section 5.3 where we showed that only a rotating bar can create an asymmetry in the stream. If the bar is static, each part of the stream receives almost the same perturbation and the stream remains symmetric. In addition, the fiducial model presented in Fig. 7 included the effect of a disc and produced a symmetric stream with no significant small-scale features. However, as we will discuss in the next subsection, disc shocks can vary the stripping rate of the progenitor and introduce small-scale, symmetric structure in the stream. We leave the importance of spiral arms for future work.

6.4 Variable stripping rate

Another mechanism that can produce structure in the density is a variable stripping rate. For globular cluster streams, the majority of the stripping occurs at pericentre, especially those corresponding to disc crossings (e.g. Johnston, Sigurdsson & Hernquist 1999;

Dehnen et al. 2004), where the tidal forces are the highest. Each stripping episode sends a packet of material into the stream which then broadens due to the debris energy dispersion. Depending on the amount stripped in each subsequent passage and the rate at which each packet broadens, this can introduce structure into the stream density which will appear symmetric near the progenitor. Indeed, this effect is naturally included in our N -body simulations and we saw evidence of this in Fig. 8 where we presented the density of our fiducial model at various times. Although the density variations at the present time are not remotely as significant as what is seen in the data (e.g. Fig. 7), it is possible that the effect can be enhanced in a different potential with, for example, a significant flattening in the halo. This will result in a larger variation in the tidal force at each pericentric passage which can cause a more substantial variation in the amount of stripped material. This could in principle explain the density minimum seen at $\phi_1 \sim -3^\circ$ although not the intertail asymmetry. However, estimates of the halo flattening based on recent modelling efforts suggest it is not significantly flattened (e.g. Koposov et al. 2010; Bowden, Belokurov & Evans 2015; Küpper et al. 2015) so we do not think this is a likely explanation. Of course, the stripping rate also depends on the properties of the progenitor (e.g. Dehnen et al. 2004). Clearly, additional work is needed to explore the range of density variations produced in different potentials for progenitors consistent with the observations of Pal 5 and its tails.

6.5 Distinguishing the various mechanisms

The stream-fanning due to weak orbital chaos could in principle cause density variations along the tidal tails. However, the most obvious feature of the debris dispersal, the fast stream width growth, is inconsistent with the results of our analysis. Similarly, although the variable stripping rate could in principle produce substantial density variations along the stream, our N -body streams evolved in a realistic potential exhibit no such features (e.g. Fig. 7), and we argued that a highly flattened halo that could enhance the effect is not supported by observations. While we cannot completely rule these effects out, we argue that there are two main plausible mechanisms that could create the features seen in Pal 5: namely, an impact by substructure (e.g. Fig. 9) and the effect of the Milky Way bar (e.g. Fig. 10). In Erkal & Belokurov (2015b), it was shown that the fly-by of a subhalo produces an almost unique signature that can be used to infer the entire set of subhalo properties. This signature is imprinted in the 6D phase-space structure of the stream and only by combining observations (e.g. the stream density, radial velocity profile and stream track) is it possible to make the inference. Presumably, the effect of the bar also produces its own unique signatures that can be distinguished from those caused by substructure. Additional efforts are clearly needed to map out these signatures. Going forward, the large catalogues of radial velocities expected from WEAVE (Dalton et al. 2012), 4MOST (de Jong et al. 2012) and DESI (Levi et al. 2013), as well as the proper motions from *Gaia* (Perryman et al. 2001), will help distinguish these mechanisms or perhaps even show that they are working in concert.

In this vein, we propose an additional diagnostic of the asymmetry mechanisms. Fig. 13 compares the observed cumulative number of stars along the leading and trailing Pal 5 tails against our fiducial N -body simulation, as well as the perturbed simulations including a subhalo fly-by and the effect of the Milky Way bar. The observed cumulative number (top panel) shows that there is substantially more material in the studied section of the trailing arm compared to the measured portion of the leading arm. The fiducial model (second panel) re-iterates that an unperturbed stream would

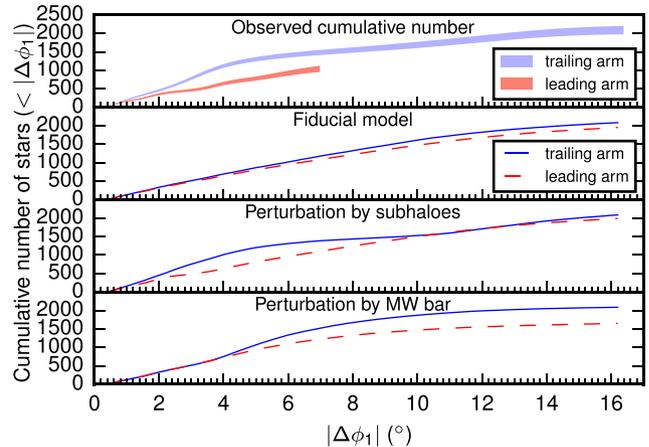


Figure 13. Comparison of the cumulative number of stars for the leading and trailing arms for the observed Pal 5 stream and three different stream models. Since the progenitor of Pal 5 is slightly offset from the origin, we show the cumulative number in terms of $\Delta\phi_1 = \phi_1 - \phi_{1, \text{Pal 5}}$. *Top panel:* 16–84 per cent confidence interval of the observed cumulative number distribution. This demonstrates that there is significantly less material in the observed section of the leading arm compared to the observed section of the trailing arm. *Second panel:* fiducial (unperturbed) N -body simulation of the Pal 5 disruption. As expected, the arms appear almost completely symmetric. *Third panel:* example of a two subhalo fly-by from Section 5.1 which looks similar to the observations by construction. The leading and trailing arms have a similar cumulative number of stars for $|\Delta\phi_1| > 12^\circ$. *Fourth panel:* example of a perturbation by the Milky Way bar from Section 5.3. In contrast to the subhalo interaction, here the difference in the debris density persists beyond the region shown here, eventually becoming symmetric only at very large distances from the progenitor. Thus, it is not immediately apparent what observations of the leading arm will reveal. We note that the cumulative numbers of stars in the simulations are scaled to match the observed cumulative number of the trailing arm in the right most bin.

possess an almost symmetric cumulative number of stars. The slight difference between the two arms here is due to the fact that leading arm is heading towards pericentre and is being stretched out. The third and fourth panels show the cumulative number behaviour for a stream perturbed by a subhalo fly-by and by the Milky Way bar, respectively. In the case of the subhalo fly-by, the asymmetry only continues until $|\Delta\phi_1| \sim 12^\circ$ after which the cumulative numbers are symmetric. In contrast, the perturbation by the Milky Way bar produces an asymmetry which persists beyond the region shown here, although it does become symmetric sufficiently far from the progenitor. Further observations of the leading arm beyond what is measured to date could help distinguish between these cases although significant additional modelling efforts are also needed to understand the features the bar (and other mechanisms) can produce.

7 DISCUSSION

7.1 Width of Pal 5

In this work, we have focused on the stream’s centroid and the debris density. However, there is also an information that can be gleaned from the stream width. Ignoring the perturbations we have discussed above, the change in the stream width is due to the stream fanning out in a non-spherical potential (e.g. Ibata et al. 2001; Helmi 2004; Johnston et al. 2005; Erkal et al. 2016a). A constant width

indicates a spherical potential and a rapidly increasing width indicates a flattened (or a triaxial) potential. In this light, we can re-visit Fig. 7 that compares the observed properties of Pal 5 with that of an N -body simulation. The N -body simulation has widths that are broadly symmetric near the progenitor with both tails exhibiting a modest fanning out. In contrast, the observed width is nearly constant for the trailing arm ($\phi_1 > 0^\circ$) but with a more rapid change in the leading arm ($\phi_1 < 0^\circ$). As with the density, such an asymmetry is likely a hallmark of a perturbation to Pal 5's stellar debris which affects each tail differently. If the leading tail is unperturbed, the large change in the width of the leading arm could indicate that the potential in the region inhabited by Pal 5 is flatter than suggested by `MWPotential2014`. Alternatively, if the trailing tail is unperturbed, the constant width would be an indication that the potential is in fact more spherical. In order to extract information from the width, we would first need to determine which tail has been perturbed. Intriguingly, in Fig. 9 which shows a simulated stream impacted by two subhaloes, we see that the subhalo impact on the trailing arm makes the width constant until $\phi_1 \sim 10^\circ$, suggesting that perhaps the relatively constant observed width is due to perturbations. Interestingly, in the presence of a bar the stream width broadens rapidly (see Fig. 11) suggesting that the width may also help distinguish a subhalo impact from the effect of the bar.

7.2 Constraining the subhalo properties

If the asymmetry seen in Pal 5 is indeed due to DM substructure, we can ask what mass range subhalo was responsible for the perturbation. In Erkal & Belokurov (2015b), a simplified model of a stream on a circular orbit was used to show that given three observables of the stream, e.g. the stream track, the stream density and the radial velocity along the stream, there is enough information to constrain the properties of the subhalo fly-by down to a degeneracy between the subhalo mass and the fly-by velocity. In Fig. 9, we compare the observed properties of Pal 5 with those of an N -body simulation including the effect of two fly-bys. We include two dotted vertical lines that show the locations of the gaps. If the Milky Way's gravitational potential was well understood there should in principle be enough information to infer the subhalo properties. This inference would also give the time since impact, allowing us to determine whether either interaction occurred near or far from the disc. This would allow us to more confidently attribute the effect to a subhalo or a GMC. However, given the current level of uncertainty in the Milky Way's mass distribution, (as well as the influence of the bar/chaos, and the properties of the Pal 5 progenitor), we will leave this for future work.

To provide a preliminary ball-park estimate, we can use the results of Erkal et al. (2016b) which computed the characteristic gap size produced by several different mass subhaloes for the Pal 5 stream. From Fig. 8 of that work, we see that subhaloes with a mass of 10^6 , 10^7 and $10^8 M_\odot$ would produce gaps with characteristic sizes of $\sim 3^\circ$, 5° and 30° in Pal 5, respectively, assuming that Pal 5 has an age of 3.4 Gyr. While the two possible gaps in Pal 5 shown in Fig. 9 have sizes of $\sim 2^\circ$ and $\sim 9^\circ$, the analysis in Erkal et al. (2016b) neglects the stretching and contraction of the stream gap due to the eccentric orbit of Pal 5. Since Pal 5 is currently near apocentre, the gaps will be contracted. This suggests that the subhalo mass for the gap in the leading arm would be in the 10^6 – $10^7 M_\odot$ range while the subhalo mass for the gap in the trailing arm would be in the 10^7 – $10^8 M_\odot$ range.

7.3 Comparison with expected number of gaps

There are arguably two non-epicyclic features in the Pal 5 stream, namely an underdensity between $\phi_1 \sim 3^\circ$ – 12° and a 2° underdensity at $\phi_1 \sim -3^\circ$. Although we have argued that there are other possible explanations for these features, we can ask if they are consistent with expectations from substructure. In Erkal et al. (2016b), it was argued that the expected background of subhaloes should produce 0.7 gaps deeper than 75 per cent in the Pal 5 stream, where the depth is relative to the unperturbed stream. In terms of the S statistic discussed in Section 3.5, this would approximately correspond to S smaller than 0.75. Note that we cannot measure the depth of the gaps observed in Pal 5 this way since we do not know the unperturbed stream density. However, we argue that a gap deeper than 75 per cent should be relatively easy to detect so the prediction should be thought of as 0.7 gaps that are detectable. The gaps from subhaloes can span a wide range of sizes from a few degrees up to a few tens of degrees with a characteristic size of roughly 8° (Erkal et al. 2016b). Note, however, that the prediction from Erkal et al. (2016b) overestimates the gap size since it assumes that the stream is on a circular orbit while Pal 5 is near apocentre which compresses the stream and the gap. Thus, both features are consistent with the expected number and size of gaps from subhaloes.

In Amorisco et al. (2016), the expected number of gaps from GMCs was studied and they found 0.65 gaps deeper than ~ 71 per cent. As discussed in Section 5.2, due to their lower masses, the GMCs create smaller gaps. So, while a perturbation from a GMC could explain the small gap seen near $\phi_1 \sim -3^\circ$, it is very unlikely that a GMC could produce the feature seen at $\phi_1 \sim 8^\circ$. Thus, the number and sizes of the features appear consistent with the combined effect of both subhaloes and GMCs. We note that Amorisco et al. (2016) assumed a pericentre of 8 kpc for Pal 5. However, as we saw in Section 4.2, the pericentre is uncertain due to underlying uncertainty about the Milky Way potential. Since the number density of GMCs depends on distance from the Milky Way's centre (e.g. Roman-Duval et al. 2010), the predicted number of gaps can change if the pericentre of Pal 5 is substantially different from 8 kpc.

Finally, in Bovy et al. (2017) the authors showed that the power spectrum and bi-spectrum of stream observables are sensitive to the amount of substructure. Using the stream density reported in Ibata et al. (2016) they found that the density fluctuations in Pal 5 are consistent with a population of subhaloes 1.5–9 times more numerous than expected in Λ CDM (accounting for a factor of 3 depletion of the subhaloes by the Milky Way disc, e.g. D'Onghia et al. 2010). This is in agreement with the direct gap counting where we expected 0.7 gaps from Erkal et al. (2016b) but found 2 gaps which suggests ~ 3 times the population of subhaloes. However, since the uncertainties are large both estimates are still consistent with Λ CDM. The density and stream track in this work can also be used with the technique of Bovy et al. (2017); however, care should be taken to apply the same cubic spline technique on the simulated streams since the measured density and stream track have been somewhat smoothed over and hence will lack power on small scales.

7.4 Putting the puzzle together

Thanks to the staggering progress in the last 10 yr, there exist clear predictions as to what might have produced the density variations we have detected in the Pal 5 stream. Interestingly, a different phenomenon (or even a set of phenomena) is involved in the production of density fluctuations at each angular scale. For example,

the sharp, small-scale density peak at $\phi_1 \sim -0.8$ (as well as its less significant counterpart in the trailing tail) is likely the result of the epicyclic bunching. The epicyclic overdensities have been studied extensively in the literature (see e.g. Combes et al. 1999; Capuzzo Dolcetta et al. 2005; Küpper et al. 2008, 2010; Mastrobuono-Battisti et al. 2012; Amorisco 2015) and are the result of the similarity of the orbits of the unbound stars. As further confirmation, these observed peaks are closely aligned with the peaks seen in our simulation at $\phi_1 = -0.7, +0.6$ (Fig. 7). Additionally, we have also studied the prominence of the epicyclic overdensities as a function of the progenitor’s orbital phase and found that near the apocentre, i.e. close to the current position of the cluster, the epicycles in the N -body simulations are strongly suppressed. Moreover, when transforming from counts of particles in the simulations to the observable number of stars, the small-amplitude epicyclic bunches are further reduced in significance, as demonstrated recently by Thomas et al. (2016). It is therefore unsurprising that even in data of such depth, only the epicycles nearest to the progenitor are detected.

While the small-scale bunching due to the epicyclic feathering is a generic feature of the stream production in any gravitational potential, the other larger scale density perturbations detected here, and in particular the asymmetry between the trailing and leading tails are highly unlikely in a smooth and static potential.

DM subhaloes have long been predicted to cause damage to the stellar streams (see e.g. Ibata et al. 2002; Johnston et al. 2002). More recently, the expectation as to the shape and size of stream perturbations caused by subhaloes has finally crystallized. First, it is now quite clear that the fly-by does not only produce a depletion in the stream density around the impact point, but also causes stars to pile up at either side of the gap (see e.g. Carlberg 2009; Erkal & Belokurov 2015a). Typically, in the numerical experiments, the perturbed parts of the stream are chosen to be sufficiently far away from the progenitor. These sections of tidal tails are characterized by relatively uniform debris densities. Therefore, the density bumps on either side of the gap have broadly similar strength with slight asymmetries in peak height possible for fly-bys with an impact parameter (e.g. Erkal & Belokurov 2015a). Closer to the progenitor, however, non-zero density gradients are expected, with the number of stars normally decaying away with distance from the Lagrange points. Accordingly, for stream perturbations in the vicinity of the parent satellite, significantly more stars might gather on the side of the gap which is nearest to the progenitor. Secondly, there exists a lower bound to the size of the stream gap a subhalo can tear. It takes time for the affected stars to move away from the impact point, and by the time the stream density drops to detectable levels, the gap has already grown substantially. Erkal et al. (2016b) carefully considered the combination of factors affecting the gap evolution and gave gap sizes expected for perturbers of different mass. In their picture, the characteristic gap size created by DM subhaloes of $10^7 M_\odot$ is $\sim 10^\circ$.

Guided by these expectations, we put forward a hypothesis that the dramatic density asymmetry between the trailing and leading debris is the result of a fly-by of a DM subhalo. More precisely, the projection of the perturber’s impact point lies somewhere between 8° and 10° , which corresponds to the centre of the density depletion of an appreciable extent. The gap is accompanied by a substantial pile-up of stars at $\phi_1 = 3^\circ$ on the side of the progenitor. We thus estimate that the gap size is approximately 9° , which is supported by the fact that the stream density recovers to normal levels at around $\phi_1 = 12^\circ$. To further demonstrate the feasibility of our conjecture, we run a series of numerical simulations of the disruption of a Pal 5-like cluster. The unperturbed density of a stream

produced in a static and smooth potential is always symmetric (see Fig. 7), thus emphasizing both the peak at $\phi_1 = 3^\circ$ and the dip between $\phi_1 = 5^\circ$ and $\phi_1 = 11^\circ$. Additional confirmation is provided by a simulation of a DM fly-by which appears to re-shape the stream density profile into one closely resembling the observations (see Fig. 9). Based on the analytic predictions of Erkal et al. (2016b) and the numerical experiments reported here, we believe that the mass of the perturber interaction which could have produced the observed density fluctuations must be in the range of 10^7 – $10^8 M_\odot$.

We considered several other mechanisms that could possibly produce the observed Pal 5 stream asymmetry. For example, we studied the effect of the Milky Way’s rotating bar whose importance for stellar streams has been highlighted in Hattori et al. (2016) and Price-Whelan et al. (2016b). In particular, we simulated the disruption of a large number of Pal 5 realizations by sampling the cluster’s proper motion and found that the presence of a bar can indeed create significant asymmetries in the stream (e.g. Fig. 11). We have established that this effect depends sensitively on the pattern speed of the bar and on the orbit of Pal 5, similar to what was found in Hattori et al. (2016). Thus, better modelling of the bar and the Milky Way potential is needed to conclusively determine if the features in Pal 5 are due to the bar. Another important effect is that of chaos that can arise both in a smooth static potential and the rotating bar (e.g. Price-Whelan et al. 2016b). Pearson et al. (2015) evolved a Pal 5 analogue in a triaxial potential and found that a seemingly mild chaos could give rise to a significantly perturbed stream. However, the particular model they considered did not match Pal 5. Most importantly, the best fit to the Pal 5 stream data available at the time was found in an axisymmetric potential which exhibited no chaos. Thus, it is unclear whether chaos could create the observed density variations or asymmetry while maintaining a thin stream. To conclude, based on the evidence in hand, the conspicuous features in the trailing tail of Pal 5 could easily be produced by a fly-by of a DM subhalo or by the Galaxy’s bar. Less likely, albeit impossible to rule out at the present, is the possibility that the stream asymmetry was caused by chaos in the Milky Way’s gravitational potential.

We stress that this substructure is not necessarily a dark subhalo, but could also be a giant molecular cloud as suggested in Amorisco et al. (2016). However, we also note that the gap sizes found in that work are smaller than the proposed gap in the Pal 5’s trailing tail. More precisely, given that the typical masses of the GMCs are less than $10^6 M_\odot$, the characteristic size of the gaps they tend to produce are less than 3° . Interestingly, this is a good match to the size of the density depletion we detect in the leading tail at $\phi_1 = -3^\circ$. This 70 per cent dip in the star counts (as measured from trough to peak) is clearly less prominent than the spectacular ripple in the trailing arm, but, none the less, carries a significance of at least 5σ . We conjecture that either a GMC or a DM subhalo, both with a mass around 10^6 – $10^7 M_\odot$ could be responsible for this gap. However, unfortunately, a straightforward interpretation of the nature of this $\sim 2^\circ$ wide gap is not possible at the moment. This is because an alternative theory can be put forward to explain it. Namely, a substantial change in the Pal 5’s pericentric distance could cause the stripping rate to vary significantly between the bouts of disruption. The varying stripping rate would induce features in the debris distribution that might look similar to the feature discussed. Of course, the debris density waves due to variable stripping efficiency must be symmetric with respect to the progenitor’s position, which is not observed. However, the presence of a large density fluctuation in the trailing tail prevents us from testing whether there exists a counterpart feature.

8 CONCLUSIONS

In this paper we have analysed new, high-quality photometry of the Palomar 5 stellar stream published recently by Ibata et al. (2016). In order to fully take advantage of this superb data, we have developed a novel non-parametric method to extract the stream properties such as the density along the tails, the centroid track of the debris distribution on the sky, as well as the stream width. Our probabilistic method is adaptive in the sense that the model complexity is not fixed a priori, but rather is driven by the data in hand. The combination of the quality of the data and the power of the modelling technique yields an exquisite determination of the stream properties.

For the first time, we measure significant changes in the stream width and show that the debris cross-section varies differently along each tail. We also detect dramatic stream density fluctuations on a variety of angular scales. First, on the scale of a fraction of a degree, density spikes are detected very close to the progenitor. Secondly, further away from the Pal 5 cluster, at $\phi_1 = -3^\circ$ in the leading tail, a density depletion approximately 2° across is measured. It is accompanied by two low-level bumps on either side. Finally, the trailing tail exhibits a prominent density enhancement at a distance of $\sim 3^\circ$ from the progenitor, followed by a smooth drop in star counts, observable for some 8° along the tail, and then a mild bump at $\phi_1 \sim 12^\circ$. As we demonstrate with utmost clarity, the remarkable rise and fall of the trailing debris density do not have counterparts in the leading tail.

We interpret the small-scale debris pile-ups in the vicinity of the progenitor as epicyclic overdensities, and conjecture that the two larger scale density perturbations are induced by interactions with small substructure. If DM subhaloes were the cause of these stream wrinkles, then their masses are of the order of 10^6 – 10^7 and 10^7 – $10^8 M_\odot$. Impressively, the size of the larger gap discovered here agrees well with the characteristic gap scale expected in the presence of Λ CDM substructure with masses between 10^5 and $10^9 M_\odot$ as predicted in Erkal et al. (2016b). It is not easy to overemphasize the importance of this discovery if the substructure that wrought havoc in the stream was non-baryonic. In fact, a subhalo in the 10^6 – $10^7 M_\odot$ range would increase the lower bound on the warm DM particle mass to >9 – 18 keV (Viel et al. 2013).

Note, however, that currently we cannot rule out other plausible explanations, such as the effect of the rotating bar, impacts by GMCs, or other complexities in the gravitational potential of the Galaxy, that would lead to mild orbital chaos or would induce substantial variations in the stripping rate of the cluster. In order to distinguish between these mechanisms, we must predict the features each of these can create in the Pal 5 tails. For example, as was shown in Erkal & Belokurov (2015b), the fly-by of a subhalo produces an almost unique signature that can be used to infer the subhalo properties. The precise signatures of the bar, chaos and the pericentre wobble are likely very different, so with additional data, especially the improved radial velocity measurements expected from WEAVE, 4MOST and DESI, it should be possible to determine the culprit.

ACKNOWLEDGEMENTS

We are grateful to Ibata et al. (2016) for making their catalogues publicly available. We thank the Streams group at Cambridge for valuable discussions. The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013)/ERC Grant Agreement no. 308024. SK thanks the United Kingdom Science and Technology Council (STFC) for the award

of Ernest Rutherford fellowship (grant number ST/N004493/1). Finally, we thank the referee for a constructive report and useful suggestions.

REFERENCES

- Akaike H., 1974, *IEEE Trans. Autom. Control*, 19, 716
 Allen C., Moreno E., Pichardo B., 2006, *ApJ*, 652, 1150
 Amorisco N. C., 2015, *MNRAS*, 450, 575
 Amorisco N. C., Gómez F. A., Vegetti S., White S. D. M., 2016, *MNRAS*, 463, L17
 Bellazzini M., Bragaglia A., Carretta E., Gratton R. G., Lucatello S., Catanzaro G., Leone F., 2012, *A&A*, 538, A18
 Belokurov V. et al., 2014, *MNRAS*, 437, 116
 Bernard E. J. et al., 2014, *MNRAS*, 443, L84
 Binney J., Tremaine S., 2008, *Galactic Dynamics: Second Edition*. Princeton Univ. Press, Princeton, NJ
 Blum R. D. et al., 2016, *Am. Astron. Soc. Meeting Abstr.* p. 317.01
 Bovy J., 2015, *ApJS*, 216, 29
 Bovy J., 2016, *Phys. Rev. Lett.*, 116, 121301
 Bovy J. et al., 2012, *ApJ*, 759, 131
 Bovy J., Erkal D., Sanders J. L., 2017, *MNRAS*, 466, 628
 Bowden A., Belokurov V., Evans N. W., 2015, *MNRAS*, 449, 1391
 Bronshtein I., Semendyayev K., Musiol G., Mühlhig H., 2007, *Handbook of Mathematics*. Springer-Verlag, Berlin
 Capuzzo Dolcetta R., Di Matteo P., Miocchi P., 2005, *AJ*, 129, 1906
 Carlberg R. G., 2009, *ApJ*, 705, L223
 Carlberg R. G., 2012, *ApJ*, 748, 20
 Carlberg R. G., 2013, *ApJ*, 775, 90
 Carlberg R. G., Grillmair C. J., Hetherington N., 2012, *ApJ*, 760, 75
 Combes F., Leon S., Meylan G., 1999, *A&A*, 352, 149
 D'Onghia E., Springel V., Hernquist L., Keres D., 2010, *ApJ*, 709, 1138
 de Jong R. S. et al., 2012, in McLean I. S., Ramsay S. K., Takami H., eds, *Proc. SPIE Conf. Ser. Vol. 8446, Ground-based and Airborne Instrumentation for Astronomy IV*. SPIE, Bellingham, p. 84460T
 Dalton G. et al., 2012, in McLean I. S., Ramsay S. K., Takami H., eds, *Proc. SPIE Conf. Ser. Vol. 8446, Ground-based and Airborne Instrumentation for Astronomy IV*. SPIE, Bellingham, p. 84460P
 Dehnen W., Odenkirchen M., Grebel E. K., Rix H.-W., 2004, *AJ*, 127, 2753
 Diemand J., Kuhlen M., Madau P., Zemp M., Moore B., Potter D., Stadel J., 2008, *Nature*, 454, 735
 Dotter A., Sarajedini A., Anderson J., 2011, *ApJ*, 738, 74
 Erkal D., Belokurov V., 2015a, *MNRAS*, 450, 1136
 Erkal D., Belokurov V., 2015b, *MNRAS*, 454, 3542
 Erkal D., Sanders J. L., Belokurov V., 2016a, *MNRAS*, 461, 1590
 Erkal D., Belokurov V., Bovy J., Sanders J. L., 2016b, *MNRAS*, 463, 102
 Eyre A., 2010, *MNRAS*, 403, 1999
 Eyre A., Binney J., 2011, *MNRAS*, 413, 1852
 Fabricius M. H. et al., 2014, *ApJ*, 787, L26
 Fritz T. K., Kallivayalil N., 2015, *ApJ*, 811, 123
 Gelman A., Rubin D. B., 1992, *Stat. Sci.*, 7, 457
 Gelman A., Hwang J., Vehtari A., 2014, *Stat. Comput.*, 24, 997
 Gerhard O., 2011, *Mem. Soc. Astron. Ital. Suppl.*, 18, 185
 Gibbons S. L. J., Belokurov V., Evans N. W., 2014, *MNRAS*, 445, 3788
 Gibbons S. L. J., Belokurov V., Evans N. W., 2017, *MNRAS*, 464, 794
 Gnedin O. Y., Ostriker J. P., 1997, *ApJ*, 474, 223
 Górski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, 622, 759
 Grillmair C. J., 2009, *ApJ*, 693, 1118
 Grillmair C. J., Dionatos O., 2006, *ApJ*, 641, L37
 Gross E., Vitells O., 2010, *Eur. Phys. J. C*, 70, 525
 Hattori K., Erkal D., Sanders J. L., 2016, *MNRAS*, 460, 497
 Helmi A., 2004, *MNRAS*, 351, 643
 Helmi A., White S. D. M., 1999, *MNRAS*, 307, 495
 Ibata R., Lewis G. F., Irwin M., Totten E., Quinn T., 2001, *ApJ*, 551, 294

- Ibata R. A., Lewis G. F., Irwin M. J., Quinn T., 2002, *MNRAS*, 332, 915
- Ibata R. A., Lewis G. F., Martin N. F., 2016, *ApJ*, 819, 1
- Ishigaki M. N., Hwang N., Chiba M., Aoki W., 2016, *ApJ*, 823, 157
- Johnston K. V., 1998, *ApJ*, 495, 297
- Johnston K. V., Sigurdsson S., Hernquist L., 1999, *MNRAS*, 302, 771
- Johnston K. V., Spergel D. N., Haydn C., 2002, *ApJ*, 570, 656
- Johnston K. V., Law D. R., Majewski S. R., 2005, *ApJ*, 619, 800
- Just A., Berczik P., Petrov M. I., Ernst A., 2009, *MNRAS*, 392, 969
- Koposov S. E., Rix H.-W., Hogg D. W., 2010, *ApJ*, 712, 260
- Küpper A. H. W., MacLeod A., Heggie D. C., 2008, *MNRAS*, 387, 1248
- Küpper A. H. W., Kroupa P., Baumgardt H., Heggie D. C., 2010, *MNRAS*, 401, 105
- Küpper A. H. W., Balbinot E., Bonaca A., Johnston K. V., Hogg D. W., Kroupa P., Santiago B. X., 2015, *ApJ*, 803, 80
- Kuzma P. B., Da Costa G. S., Keller S. C., Maunder E., 2015, *MNRAS*, 446, 3297
- Lang D., Hogg D. W., Mykytyn D., 2016, *Astrophysics Source Code Library*, record ascl:1604.008
- Law D. R., Majewski S. R., 2010, *ApJ*, 714, 229
- Levi M. et al., 2013, preprint ([arXiv:1308.0847](https://arxiv.org/abs/1308.0847))
- Long K., Murali C., 1992, *ApJ*, 397, 44
- López-Corredoira M., Cabrera-Lavers A., Gerhard O. E., 2005, *A&A*, 439, 107
- Majewski S. R., Skrutskie M. F., Weinberg M. D., Ostheimer J. C., 2003, *ApJ*, 599, 1082
- Mastrobucchi-Battisti A., Di Matteo P., Montuori M., Haywood M., 2012, *A&A*, 546, L7
- Miyamoto M., Nagai R., 1975, *PASJ*, 27, 533
- Neal R., 2011, in Brooks S., Gelman A., Jones G. L., Meng X.-L., eds, *Handbook of Markov Chain Monte Carlo*. Chapman and Hall/CRC Press, Boca Raton, FL, p. 113
- Niederste-Ostholt M., Belokurov V., Evans N. W., 2012, *MNRAS*, 422, 207
- Odenkirchen M. et al., 2001, *ApJ*, 548, L165
- Odenkirchen M., Grebel E. K., Dehnen W., Rix H.-W., Cudworth K. M., 2002, *AJ*, 124, 1497
- Odenkirchen M. et al., 2003, *AJ*, 126, 2385
- Pearson S., Küpper A. H. W., Johnston K. V., Price-Whelan A. M., 2015, *ApJ*, 799, 28
- Peñarrubia J., Gómez F. A., Besla G., Erkal D., Ma Y.-Z., 2016, *MNRAS*, 456, L54
- Perryman M. A. C. et al., 2001, *A&A*, 369, 339
- Portail M., Wegg C., Gerhard O., Martínez-Valpuesta I., 2015, *MNRAS*, 448, 713
- Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P., 2007, *Numerical Recipes 3rd Edition: The Art of Scientific Computing*, 3 edn. Cambridge Univ. Press, Cambridge
- Price-Whelan A. M., Johnston K. V., Valluri M., Pearson S., Küpper A. H. W., Hogg D. W., 2016a, *MNRAS*, 455, 1079
- Price-Whelan A. M., Sesar B., Johnston K. V., Rix H.-W., 2016b, *ApJ*, 824, 104
- Rockosi C. M. et al., 2002, *AJ*, 124, 349
- Roman-Duval J., Jackson J. M., Heyer M., Rathborne J., Simon R., 2010, *ApJ*, 723, 492
- Sanders J. L., Bovy J., Erkal D., 2016, *MNRAS*, 457, 3817
- Schlegel D. J., Finkbeiner D. P., Davis M., 1998, *ApJ*, 500, 525
- Schönrich R., Binney J., Dehnen W., 2010, *MNRAS*, 403, 1829
- Springel V., 2005, *MNRAS*, 364, 1105
- Springel V. et al., 2008, *MNRAS*, 391, 1685
- Theano Development Team 2016, preprint ([arXiv:1605.02688](https://arxiv.org/abs/1605.02688))
- Thomas G. F., Ibata R., Famaey B., Martin N. F., Lewis G. F., 2016, *MNRAS*, 460, 2711
- Viel M., Becker G. D., Bolton J. S., Haehnelt M. G., 2013, *Phys. Rev. D*, 88, 043502
- Wilson A. G., 1955, *PASP*, 67, 27
- Yoon J. H., Johnston K. V., Hogg D. W., 2011, *ApJ*, 731, 58
- Zhu C., Byrd R. H., Lu P., Nocedal J., 1997, *ACM Trans. Math. Softw.*, 23, 550
- Zotos E. E., 2015, *MNRAS*, 446, 770

SUPPORTING INFORMATION

Supplementary data are available at [MNRAS](https://www.mnras.org/) online.

meas_mnras.dat

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the paper.

APPENDIX A: COORDINATE TRANSFORMATION MATRIX

In Section 2.2 we described the rotated coordinate system, (ϕ_1, ϕ_2) , which is approximately aligned with the stream. The transformation from (α, δ) to (ϕ_1, ϕ_2) is given by

$$\begin{bmatrix} \cos(\phi_1) \cos(\phi_2) \\ \sin(\phi_1) \cos(\phi_2) \\ \sin(\phi_2) \end{bmatrix} = \begin{bmatrix} -0.656057 & -0.754711 & 0.000636 \\ 0.609115 & -0.528995 & 0.590883 \\ -0.445608 & 0.388045 & 0.806751 \end{bmatrix} \times \begin{bmatrix} \cos(\alpha) \cos(\delta) \\ \sin(\alpha) \cos(\delta) \\ \sin(\delta) \end{bmatrix}$$

APPENDIX B: PAL 5 STREAM DENSITY IN THE DECALS DATA SET

To demonstrate the robustness of the recovered Pal 5 stream density profile, and to verify that the features observed are not caused by any issues related to the CFHT observations (such as data quality variation along the stream or an incomplete footprint), we have compared the measured stream density with a completely independent data set, namely the DECaLS (Blum et al. 2016). The DECaLS data set is comprised of *g*-, *r*- and *z*-band photometry covering the stream from $\phi_1 \sim -5^\circ$ to $\phi_1 \sim 20^\circ$ at a depth of 0.5–1 magnitude deeper than the SDSS. Here, we use the source catalogue from the second data release of the survey (DR2) and only include objects classified as point sources (`type = 'PSF'`; see Lang, Hogg & Mykytyn 2016). In this data release, the continuous coverage along the stream exists only in *r* and *z* bands. Therefore, we use the same matched filter analysis to select Pal 5 stream stars as described in Section 2 but in *r*, *z* bands (down to a limit of $z < 22$). Because the DECaLS data coverage is not restricted to a narrow region around the stream we can perform a proper background subtraction of the stream densities without doing model fitting, but instead using two background regions above and below the stream.

Fig. B1 shows the background-subtracted density of the CMD-selected Pal 5 stars within 0.25° of the stream track. Overplotted on top of the density recovered from the DECaLS data is the scaled best-fitting stream density model based on the CFHT data set (as measured in Section 3.5). We also note that a one-to-one match is not necessarily expected given that the DECaLS data is considerably shallower compared to the CFHT data as well as less homogeneous. None the less, most density fluctuations extracted from the CFHT data have clear counterparts in the DECaLS data set. In particular, the large overdensity/stream asymmetry at $\phi_1 \sim 3^\circ$, the epicyclic overdensity at $\phi_1 \sim -0.7^\circ$ and the underdensity at $\phi_1 \sim -3^\circ$ are all present in the DECaLS Pal 5 stream density profile. Accordingly, re-assured by this test, we conclude that the stream density model presented above gives a fair representation of the debris distribution along the Pal 5 tidal tails.

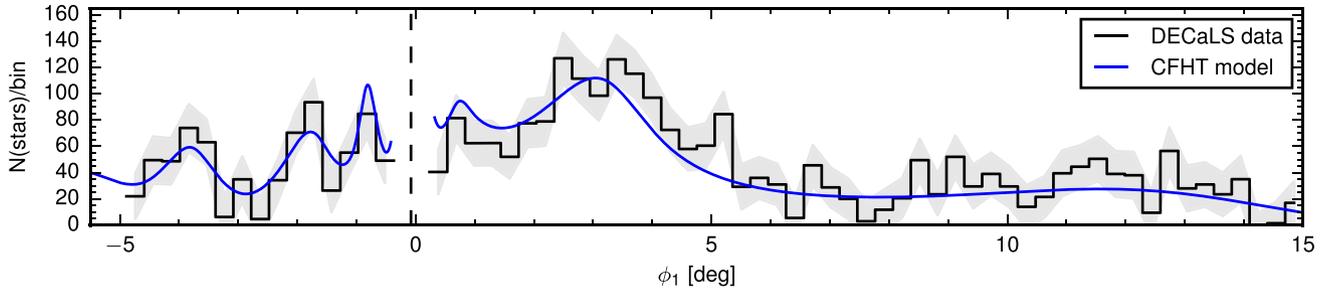


Figure B1. Comparison of the model Pal 5 stream density from CFHT data with data from the DECaLS survey. The histogram with the grey band shows the background-subtracted density of the candidate Pal 5 stream stars in DECaLS (selected using the optimal colour–magnitude mask in r , z bands) within $0^{\circ}.25$ of the stream track. The blue curve shows the best-fitting stream density model that was extracted from the CFHT data. The size of the bin is $0^{\circ}.3$. The grey bands represent the Poisson uncertainty of the stream star counts and incorporate the uncertainty from the background subtraction.

APPENDIX C: COMPARISON WITH SIMULATIONS

To test the performance of the non-parametric stream model described in Section 3, we have used it to extract density features from several Pal 5-like stellar streams generated using N -body simulations and perturbed with varying amounts of substructure. The simulations are broadly similar to those described in Section 4.3 and are run with a modified version of GADGET-3. We use a potential similar to MWPotential2014 from Bovy (2015) except the bulge is replaced with a $5 \times 10^9 M_{\odot}$ Hernquist profile with a scale radius of 500 pc. As above, the progenitor is taken to be a King profile with a mass of $2 \times 10^4 M_{\odot}$, $w = 2$, a core radius of 15 pc and is modelled with 10^5 equal mass particles and a softening length of 1 pc. The progenitor is given the best-fitting line-of-sight velocity and proper motions from Küpper et al. (2015), rewound for 3.4 Gyr

(the best-fitting age from Küpper et al. 2015), and then allowed to disrupt for the same amount of time. In the fiducial simulation, we do not include any substructure. To perturb the stream, we include the effect of substructure by taking the reported number density profile of subhaloes from Springel et al. (2008), scaling the profile down to a host mass of $10^{12} M_{\odot}$ as described in section 2.4 of Erkal et al. (2016b) and including the expected population of subhaloes with masses between 10^6 and $10^9 M_{\odot}$. Each subhalo is modelled as a single tracer particle which sources a Hernquist profile force given by the mass and scale radius of the subhalo. The relation between the mass and scale radius comes from fits to the v_{\max} – M_{tidal} relation in the public catalogues of Via Lactea II (Diemand et al. 2008) and is given by

$$r_s = 1.05 \text{ kpc} \left(\frac{M}{10^8 M_{\odot}} \right)^{1/2}. \quad (\text{C1})$$

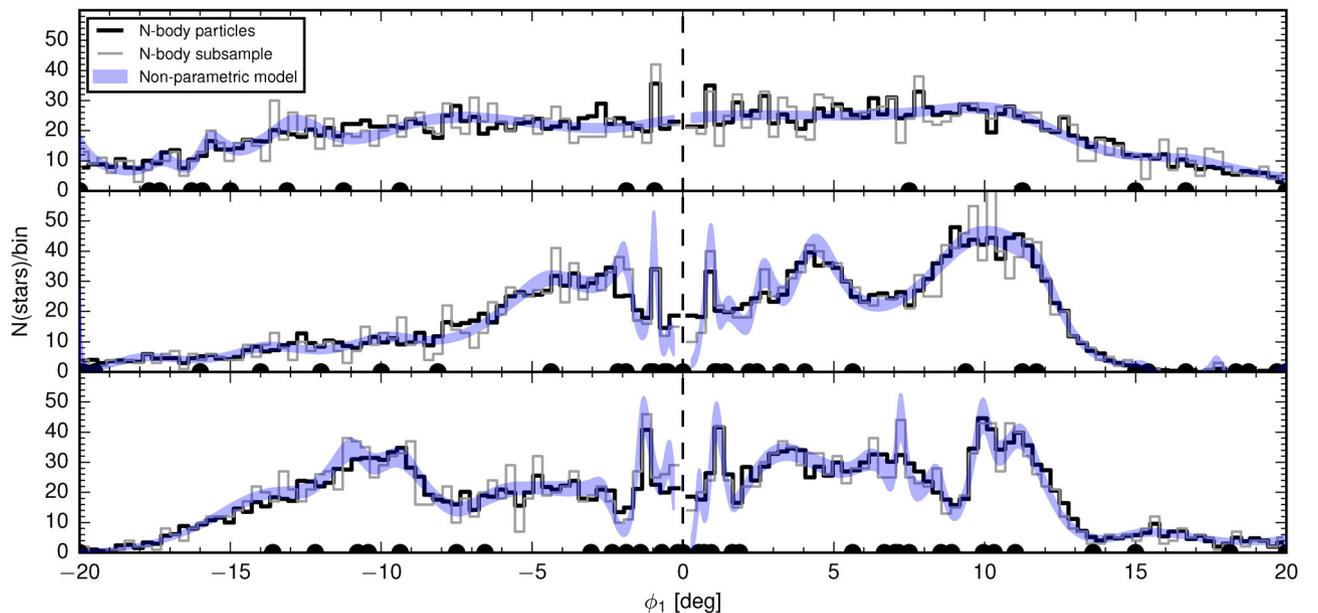


Figure C1. Performance of the non-parametric stream modelling method on simulated data. This shows the density of stars along simulated Pal 5-like streams that were perturbed by different amounts of substructure. *Top*: fiducial stream, not perturbed by any substructure. *Middle*: stream evolved with the expected amount of substructure between 10^6 and $10^9 M_{\odot}$ matching the Λ CDM predictions. *Bottom*: stream evolved with three times the expected amount of Λ CDM substructure between 10^6 and $10^9 M_{\odot}$. In each panel, black histogram gives the density of the stream particles in $0^{\circ}.3$ wide bins. The grey histogram shows the density of the 20 per cent subsample of all the particles used for stream density extraction. Blue bands show the non-parametric density measurement together with 16 per cent, 84 per cent credible intervals from the posterior samples. In all three cases the non-parametric method recovers well the majority of the structure in stream densities, including the small-scale epicyclic overdensities. Note that the number of subhaloes used does not account for the effect of the disc of the host which is expected to deplete the subhaloes by a factor of 3 (D’Onghia et al. 2010).

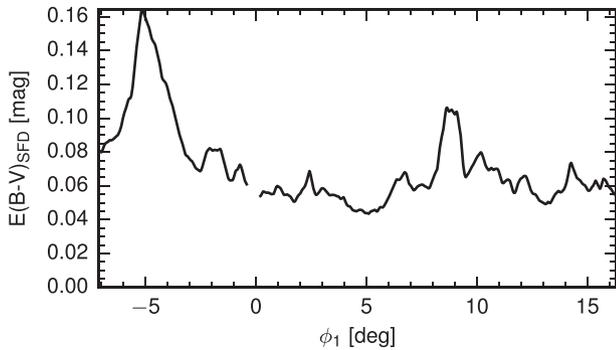


Figure D1. Average dust reddening within the measured width of the stream using the maps from Schlegel et al. (1998). The excess reddening at $-6^\circ < \phi_1 < -4^\circ$ may create a feature in the stream as we discuss in Section 3.5 but there are no dust features corresponding to the density variations of interest we find in Pal 5.

We then run a simulation in the expected Λ CDM background and a simulation in three times the expected background. Note that the number of subhaloes used does not account for the effect of the disc of the host which is known to deplete their number by a factor of 3 (e.g. D’Onghia et al. 2010).

For each stream in the simulations described above, we ran our measuring algorithm to extract a 1D stellar density profile, $I(\phi_1)$, as described in Section 3.2. Importantly, exactly the same procedure for dynamic spline node placement as described in Section 3.3 is used. Note that here only the 1D debris density distribution is determined, i.e. the track on the sky, stream width and background density are not modelled. None the less, we believe that even in this somewhat limited set-up, the simulated stream data provides a suitable test case for the algorithm. In order to make the simulated streams look closer to the actual Pal 5 stream as it appears in the CFHT data, we use only ~ 20 per cent of particles from the simulations. This subsampling produces a particle density similar to the median stellar density observed in the Pal 5 stream, namely two stars per arcminute (see Fig. 5). Fig. C1 shows the comparison of linear densities of the simulated streams as extracted by our non-parametric model (blue band shows 1σ credible intervals from the posterior samples) versus the simple histogram of the full set of particles in the simulation (thick black line) and the actual subset used in the fitting (thin grey line). We see that in all three cases, the density measurement matches very closely the true stream density, correctly extracting both large-scale fluctuations produced by the subhalo fly-bys and the small-scale features related to the epicyclic overdensities.

APPENDIX D: DUST EXTINCTION

As a check that dust extinction did not create any of the features of interest, we compute the average dust reddening within the width of the stream. Specifically, we take the stream track and width from Section 3.5 and compute the average $E(B - V)$ within one stream width using dust maps from Schlegel, Finkbeiner & Davis (1998). Fig. D1 shows this average with a clear excess of dust at -5° . While this excess reddening may create a feature in the stream (as discussed in Section 3.5), there are no dust features corresponding to the gap at $\phi_1 \sim -3^\circ$, the overdensity at $\phi_1 \sim 3^\circ$, or the broad underdensity around $\phi_1 \sim 8^\circ$. Thus, do not believe that extinction can explain the density variations of interest which we find in Pal 5.

APPENDIX E: SUPPLEMENTARY DATA

As supplementary material, we provide the summary of the Pal 5 stream properties shown in Fig. 5 in machine-readable form with the following columns:

- (i) `phi1`: angle along the stream ϕ_1 (deg);
- (ii) `phi2_16`, `phi2_50`, `phi2_84`: angle across the stream [deg] (16 per cent, 50 per cent, 84 per cent percentile);
- (iii) `sbstream_16`, `sbstream_50`, `sbstream_84`: surface brightness of the stream in stars per square arcminute (16 per cent, 50 per cent, 84 per cent percentile);
- (iv) `ldens_16`, `ldens_50`, `ldens_84`: linear density of the stream in stars per arcmin (16 per cent, 50 per cent, 84 per cent percentile);
- (v) `cumdens_16`, `cumdens_50`, `cumdens_84`: cumulative density of the stream from the progenitor (16 per cent, 50 per cent, 84 per cent percentile);
- (vi) `width_16`, `width_50`, `width_84`: Gaussian stream width (16 per cent, 50 per cent, 84 per cent percentile);
- (vii) `bgdens_16`, `bgdens_50`, `bgdens_84`: background stellar density (16 per cent, 50 per cent, 84 per cent percentile).

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.