Discovery of an equal-mass 'twin' binary population reaching 1000 + au separations

Kareem El-Badry, Hans-Walter Rix, Haijun Tian, Gaspard Duchêne, Maxwell Moe

Discovery of an equal-mass 'twin' binary population reaching 1000 + au separations

Kareem El-Badry [1,2]★ Hans-Walter Rix,[2] Haijun Tian,[2,3,4] Gaspard Duchêne [1,5] and Maxwell Moe[6]

[1]*Department of Astronomy and Theoretical Astrophysics Center, University of California Berkeley, Berkeley, CA 94720, USA*
[2]*Max Planck Institute for Astronomy, D-69117 Heidelberg, Germany*
[3]*Department of Astronomy, China Three Gorges University, Yichang 443002, China*
[4]*Center of Astronomy and Space Science Research, China Three Gorges University, Yichang 443002, China*
[5]*CNRS, Institut de Planétologie et d'Astrophysique de Grenoble (IPAG), Univ. Grenoble Alpes, F-38000 Grenoble, France*
[6]*Steward Observatory, University of Arizona, Tucson, AZ 85721, USA*

ABSTRACT

We use a homogeneous catalogue of 42 000 main-sequence wide binaries identified by *Gaia* to measure the mass ratio distribution, $p(q)$, of binaries with primary masses $0.1 < M_1/M_\odot < 2.5$, mass ratios $0.1 \lesssim q < 1$, and separations $50 < s/\mathrm{au} < 50\,000$. A well-understood selection function allows us to constrain $p(q)$ in 35 independent bins of primary mass and separation, with hundreds to thousands of binaries in each bin. Our investigation reveals a sharp excess of equal-mass 'twin' binaries that is statistically significant out to separations of 1000–10 000 au, depending on primary mass. The excess is narrow: a steep increase in $p(q)$ at $0.95 \lesssim q < 1$, with no significant excess at $q \lesssim 0.95$. A range of tests confirm the signal is real, not a data artefact or selection effect. Combining the *Gaia* constraints with those from close binaries, we show that the twin excess decreases with increasing separation, but its width ($q \gtrsim 0.95$) is constant over $0.01 < a/\mathrm{au} < 10\,000$. The wide twin population would be difficult to explain if the components of all wide binaries formed via core fragmentation, which is not expected to produce strongly correlated component masses. We conjecture that wide twins formed at closer separations ($a \lesssim 100$ au), likely via accretion from circumbinary discs, and were subsequently widened by dynamical interactions in their birth environments. The separation-dependence of the twin excess then constrains the efficiency of dynamical widening and disruption of binaries in young clusters. We also constrain $p(q)$ across $0.1 \lesssim q < 1$. Besides changes in the twin fraction, $p(q)$ is independent of separation at fixed primary mass over $100 \lesssim s/\mathrm{au} < 50\,000$. It is flatter than expected for random pairings from the initial mass function but more bottom-heavy for wide binaries than for binaries with $a \lesssim 100$ au.

Key words: binaries: general – binaries: visual – stars: formation – stars: statistics.

1 INTRODUCTION

Binary stars are ubiquitous: roughly half of all field stars have binary companions (Duquennoy & Mayor 1991; Duchêne & Kraus 2013; Moe & Di Stefano 2017), and the binary fraction is even higher in star-forming regions (e.g. Ghez, Neugebauer & Matthews 1993; Leinert et al. 1993; Mathieu 1994; Connelley, Reipurth & Tokunaga 2008; Sadavoy & Stahler 2017; Duchêne et al. 2018). Significant progress has been made in developing theoretical models to explain the population statistics of observed binaries (e.g. Kroupa 1995; Marks, Kroupa & Oh 2011; Bate 2012; Lomax et al. 2015),

but fundamental aspects of the binary formation process remain imperfectly understood.

The distribution of binary mass ratios has been a subject of interest for at least a century (e.g. van Biesbroeck 1916; Öpik 1924; Kuiper 1935). As a final outcome of the binary formation process, the mass ratio distribution provides useful constraints on theoretical models of star formation. Unlike the distributions of orbital separation and eccentricity, the mass ratio distribution has been suggested to be insensitive to dynamical evolution after formation (such that binaries of different mass ratios are disrupted at similar rates; e.g. Parker & Reggiani 2013). Mapping the mass ratio distribution over a range of binary masses and separations has thus been the focus of many studies (e.g. Trimble 1974, 1987, 1990; Eggleton, Fitchett & Tout 1989; Hogeveen 1992; Mazeh et al.

★ E-mail: kelbadry@berkeley.edu

1992; Shatsky & Tokovinin 2002; Burgasser et al. 2007; Söderhjelm 2007; Raghavan et al. 2010; Tokovinin 2014; Gullikson, Kraus & Dodson-Robinson 2016).

Observational studies of the mass ratio distribution are complicated by incompleteness. All binary detection methods are biased against low-mass ratio companions, which produce weaker radial velocity (RV) shifts at fixed separation (e.g. Shahaf & Mazeh 2019), contribute less light to the observed spectra of unresolved binaries (e.g. El-Badry et al. 2018a,b), cause weaker eclipses (e.g. Moe & Di Stefano 2013), and are less likely to be detected as visual companions (e.g. Tokovinin 2011; El-Badry & Rix 2018). The detection efficiency also varies with primary mass and orbital separation. This complicates measurement of the mass ratio distribution, because the distributions of separation, primary mass, and mass ratio are not independent (e.g. Moe & Di Stefano 2017). It is thus important for demographic studies of binaries that the selection function of observed samples is well understood. If possible, binaries of different masses and physical separations should be considered independently.

A puzzling feature of the mass ratio distribution identified by previous works is the so-called twin phenomenon, which refers to a purported statistical excess of nearly equal-mass binaries with mass ratios $0.95 \lesssim q < 1$. Most studies that find an excess of equal-mass twins have focused on spectroscopic binaries with close separations ($a \ll 1$ au, e.g. Lucy & Ricco 1979; Hogeveen 1992; Tokovinin 2000; Halbwachs et al. 2003; Lucy 2006; Pinsonneault & Stanek 2006; Simon & Obbie 2009; Kounkel et al. 2019). Indeed, several studies have reported a sharp drop-off in the fractional excess of twin binaries beyond periods of 40 d ($a \lesssim 0.2$ au; e.g. Lucy & Ricco 1979; Tokovinin 2000; Simon & Obbie 2009). Recent studies have not confirmed such a sharp drop-off in the twin excess for solar-type stars (Tokovinin 2014; Moe & Di Stefano 2017), but have still found it to decline monotonically with increasing separation. At least for $M_1 \gtrsim 1\,\mathrm{M}_\odot$ (the mass range on which most previous studies have focused), the twin excess has been found to decrease with increasing primary mass at fixed separation and to extend to wider separations for lower mass primaries (Moe & Di Stefano 2017).

Some works have also argued that the twin phenomenon may be a selection effect (Mazeh et al. 2003; Cantrell & Dougan 2014), since equal-mass binaries are brighter, can be detected at larger distances, and are preferentially selected by several binary detection methods (e.g. Branch 1976). Such biases are minimized for binary samples that are nearly volume-complete (e.g. Raghavan et al. 2010; Tokovinin 2014) and/or have well-understood completeness.

If it is a real effect, the physical origin of the twin phenomenon is not fully understood. Several mechanisms have been proposed that could lead to preferential formation of equal-mass binaries, including fragmentation during the late stages of protostellar collapse, mass transfer between pre-main-sequence (MS) stars, and competitive accretion (see Tokovinin 2000 for discussion of different formation mechanisms). Several simulations have predicted that accretion of high angular momentum gas, particularly from a circumbinary disc, tends to drive binary mass ratios towards unity (e.g. Bate 2000; Farris et al. 2014; Young & Clarke 2015). However, it is not obvious why, when averaged over a large population of binaries, accretion from a circumbinary disc would produce a sharp peak in the mass ratio distribution at $q \gtrsim 0.95$ as opposed to a gradual increase.

There have also been hints of an excess of twins among spatially resolved wide binaries with separations ranging from tens to thousands of au (Trimble 1987; Söderhjelm 2007). The selection functions of the wide binary samples studied in these works were poorly understood, causing investigators to remain agnostic of whether the excess of equal-brightness pairs in their catalogues was of astrophysical origin or rooted in selection biases. Recently, Moe & Di Stefano (2017) measured the twin excess at different separations in a small but relatively complete sample of solar-type binaries within 25 pc of the Sun. They found the twin excess to decline with separation, but found it inconsistent with 0 out to separations of 200 au. At even wider separations, they set an upper limit of $\sim$5 per cent on the excess twin fraction.

High-quality parallaxes and proper motions from the recent *Gaia* data releases (Gaia Collaboration 2016, 2018) have simplified the process of constructing samples of wide binaries with (a) little contamination from chance alignments and (b) a well-understood selection function. Using data from *Gaia* DR2, El-Badry & Rix (2018, hereafter ER18) constructed a high-purity catalogue of wide binaries within 200 pc of the Sun consisting mainly of AFGKM dwarfs. In this paper, we use a subset of that catalogue to constrain the mass ratio distribution over a wide range of primary masses ($0.1 \lesssim M/\mathrm{M}_\odot \lesssim 2.5$), mass ratios ($0.1 \lesssim q < 1$), and separations ($50 \lesssim s/\mathrm{au} < 50\,000$). The large size of the catalogue allows us to constrain $p(q)$ in narrow bins of primary mass and separation independently. This approach make it possible to measure variation in $p(q)$ with mass and separation, and it minimizes the sensitivity of our results to imperfectly known inputs such as the initial mass function (IMF) and separation distribution.

A striking result of our investigation is the unambiguous evidence that twins are not purely a close binary phenomenon: a significant excess of equal-mass ($q \gtrsim 0.95$) binaries persists out to separations as wide as 10 000 au. We derive constraints on the excess fraction of twins and the width of the twin excess as a function of mass and separation. We also provide constraints on the full mass ratio distribution over $0.1 \lesssim q \leq 1$ in all bins.

The remainder of this paper is organized as follows. Section 2 describes the binary catalogue and tests we have done to verify that the twin excess is real. In Section 3, we describe how we forward-model synthetic binary populations to fit for the intrinsic mass ratio distribution. Results of this fitting are presented in Section 4. In Section 5, we compare to previous work and discuss implications of our results for models of binary star formation and dynamical evolution. The appendices provide additional details about several aspects of the data and model. There we discuss sensitivity to the adopted parametric form of $p(q)$ (Appendix A), systematic uncertainties in our model (Appendix B), evidence for a twin excess in archival wide binary catalogues (Appendix C), empirical determination of the selection function (Appendix D), and validation of the adopted Galactic model and selection function (Appendix E). Constraints on fitting parameters are tabulated in Appendix G.

2 DATA

Our primary analysis uses the binary catalogue assembled in ER18. This catalogue contains $\sim$55 000 spatially resolved wide binaries with MS, giant, and white dwarf components, and projected physical separations of $50 \lesssim s/\mathrm{au} < 50\,000$. We refer to ER18 for a full description of the catalogue's contents. In brief, it was constructed by searching *Gaia* DR2 for nearby ($d < 200$ pc) pairs of stars whose positions, proper motions, and parallaxes are consistent with being gravitationally bound. Resolved higher order multiples and suspected members of bound and dissolving clusters were removed. The catalogue is designed to be pure but not complete: cuts on photometric and astrometric precision ensure that the contamination

rate from chance alignments is low ($\ll 1$ per cent), but they also reduce the number of faint binaries.

We do not use the full catalogue from ER18, but impose the following additional cuts:

(i) We only consider MS/MS binaries, removing binaries in which either component is suspected to be a white dwarf, subgiant, giant, or pre-MS star. We identify non-MS components from the colour–magnitude diagram (CMD; see below).

(ii) We require *both* components to have `parallax > 5`, `parallax_over_error > 20`, `phot_bp_mean_flux_over_error > 20`, and `phot_rp_mean_flux_over_error > 20`. ER18 used these same cuts for the primary, but used less stringent cuts for the secondary. Here, we apply the same cuts to both components in order to symmetrize the selection function.

(iii) We reject binaries in which the CMD-inferred mass of the primary falls outside $0.1 < M_1/\mathrm{M}_\odot < 2.5$. The ER18 catalogue contains fewer than 100 binaries with estimated primary masses above 2.5 $\mathrm{M}_\odot$, and none with estimated primary masses below $0.1\mathrm{M}_\odot$.

These additional cuts remove 23 per cent of the ER18 catalogue, leaving us with a sample of 42 338 MS/MS binaries.

Fig. 1 shows the CMD of all stars in the ER18 catalogue, with primaries and secondaries included on the same axes. Overplotted PARSEC isochrones show that the spread on the lower MS is primarily attributable to metallicity (upper right panel), while that on the upper MS and red giant branch is primarily due to age (lower left panel). A secondary sequence consisting mainly of unresolved binaries is visible above the MS. Wide 'binaries' with one component in this sequence are primarily hierarchical triples. We do not remove these from our sample but account for them in our model when fitting for the mass ratio distribution in Section 3.[1]

MS components that fall below an MS PARSEC isochrone with $[\mathrm{Fe/H}] = 0.5$ (black points in the lower right panel of Fig. 1) are treated as single stars. Those that fall above this isochrone but below an unresolved binary isochrone for two equal-mass stars with $[\mathrm{Fe/H}] = 0.5$ (yellow points) are treated as unresolved binaries. Finally, sources that fall above this binary isochrone (red points; likely a mix of unresolved triples, pre-MS stars, and giants/subgiants) are rejected, as are white dwarfs. Of the 42 382 wide binaries in our catalogue, 35 087 have two components consistent with having no unresolved sub-components (black points), and 7295 have at least one component suspected to be an unresolved binary (yellow points).

We estimate masses for both components of each binary based on their location in the CMD. The ratios of these masses are not used explicitly in fitting the mass ratio distribution, but they are used to assign the primary versus secondary components and to assign binaries to bins of primary mass. For MS stars that are suspected to be single, we estimate masses by interpolating from a grid of single-star PARSEC isochrones. This method is reasonably effective for single components, but it would yield biased results for the unresolved sub-components.

To estimate masses for components suspected to be unresolved binaries, we construct a population of synthetic unresolved binaries (see Section 3) and, for the subset of this population that falls in

the region of the CMD coloured in yellow in the bottom right panel of Fig. 1, we calculate the median *primary* mass as a function of M_G of the unresolved binary. For observed sources in that region of the CMD, we estimate the primary mass by interpolating from M_G on this median relation. That is, the mass assigned to unresolved components represents the mass of the primary of the unresolved component, not the total mass.

This method of assigning masses is not without drawbacks: masses assigned to unresolved binaries are imprecise, because the mass ratio is not known. In addition, some pre-MS stars may be mistaken for unresolved binaries, and some low-mass ratio unresolved binaries may be mistaken for higher metallicity single stars. However, we expect the typical accuracy to be $\lesssim 0.1\,\mathrm{M}_\odot$, which is good enough for our purpose of assigning binaries to different bins of primary mass prior to fitting. In modelling the mass ratio distribution, it is not critical that the mass ratio of any one binary be measured accurately, but rather that the *distribution* of magnitude difference be predicted self-consistently.

Fig. 2 shows 'postage-stamp' images[2] of some of the binaries in our catalogue from the Sloan Digital Sky Survey (SDSS; York et al. 2000). We note that SDSS photometry is not used in our analysis or in fitting the mass ratio distribution; we show it here because raw images from *Gaia* are not publicly available. Source contamination in the SDSS images is generally expected to be more severe than in *Gaia* photometry due to atmospheric seeing. To showcase the diversity of binaries in the catalogue, we choose a selection of binaries with roughly solar-mass primaries and a range of mass ratios (top panel), binaries with primary masses of $\approx 0.5\,\mathrm{M}_\odot$ and a range of mass ratios (middle panel), and some examples of equal-mass 'twin' binaries with component masses of $\approx 0.5\,\mathrm{M}_\odot$ (bottom panel). The twins are easily recognized because their magnitudes and colours are very similar.

An observable quantity that is closely related to the mass ratio is the difference in apparent magnitude of the two components, $\Delta G = |G_1 - G_2|$. Fig. 3 shows the distribution of ΔG as a function of apparent magnitude of the primary, angular separation, and projected physical separation. The sign of ΔG is randomized, such that the distribution about $\Delta G = 0$ is symmetric. This is helpful in making the equal-brightness population stand out, since it would otherwise be squished against the *x*-axis. An excess of equal-brightness binaries is evident over a wide range of apparent magnitudes; it is strongest at close separations. As we will show, the feature is quite narrow: the density of binaries is enhanced primarily at magnitude differences of $\Delta G < 0.25$ mag. This is much narrower than the range over which the selection function varies significantly: at $\theta > 2.5$ arcsec, the contrast sensitivity is basically constant over $0 < \Delta G < 2$ (see ER18; their appendix A).

The middle panel of Fig. 3 shows that there are no binaries with small separations and large ΔG in the catalogue. This is a consequence of photometric contamination at small angular separations. The dashed line shows the contrast sensitivity limit derived in Appendix D2; this is the value of ΔG at which the sensitivity is 50 per cent of its value at asymptotically large separations for a given θ. The contrast limit is derived from the correlation function of chance alignment sources subject to similar quality cuts as real binaries. The fall-off in sensitivity is quite steep beyond the contrast limit, leading to an envelope in $\Delta G(\theta)$ beyond which no binaries are found. The drop-off towards larger ΔG is less steep as a function of physical separation *s*, since

---

[1]We also experimented with removing binaries with suspected unresolved components from the sample; doing so does not qualitatively change any of our results.

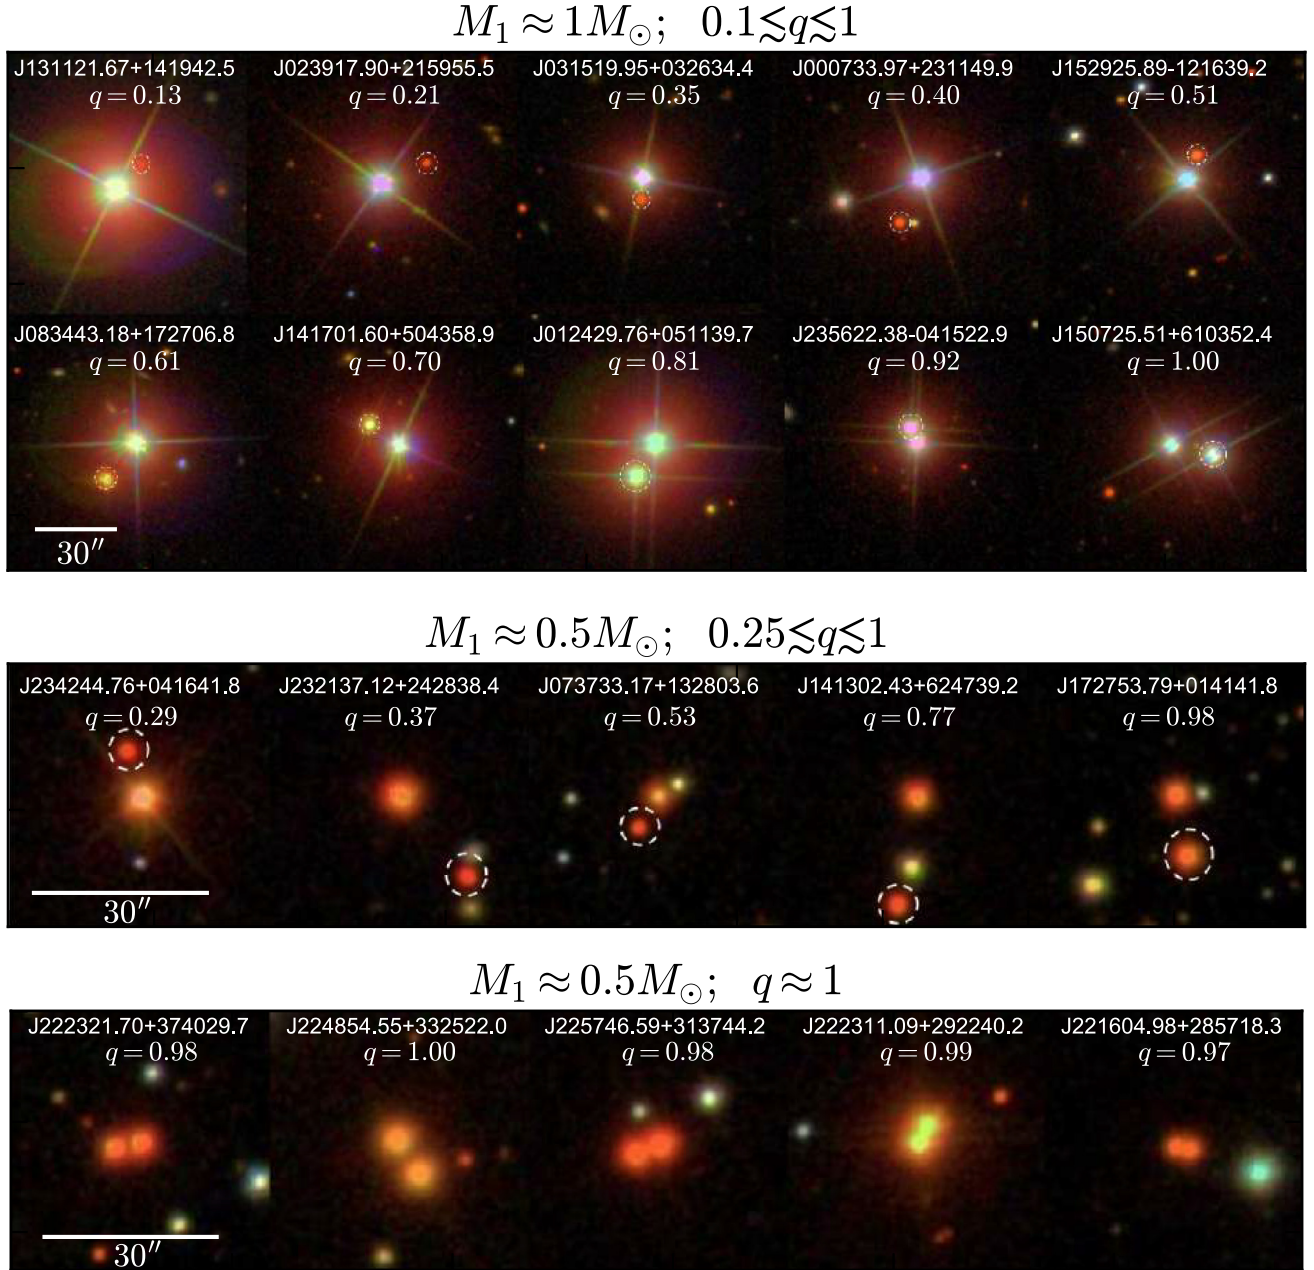[2]skyserver.sdss.org/dr14/en/tools/chart/listinfo.aspx
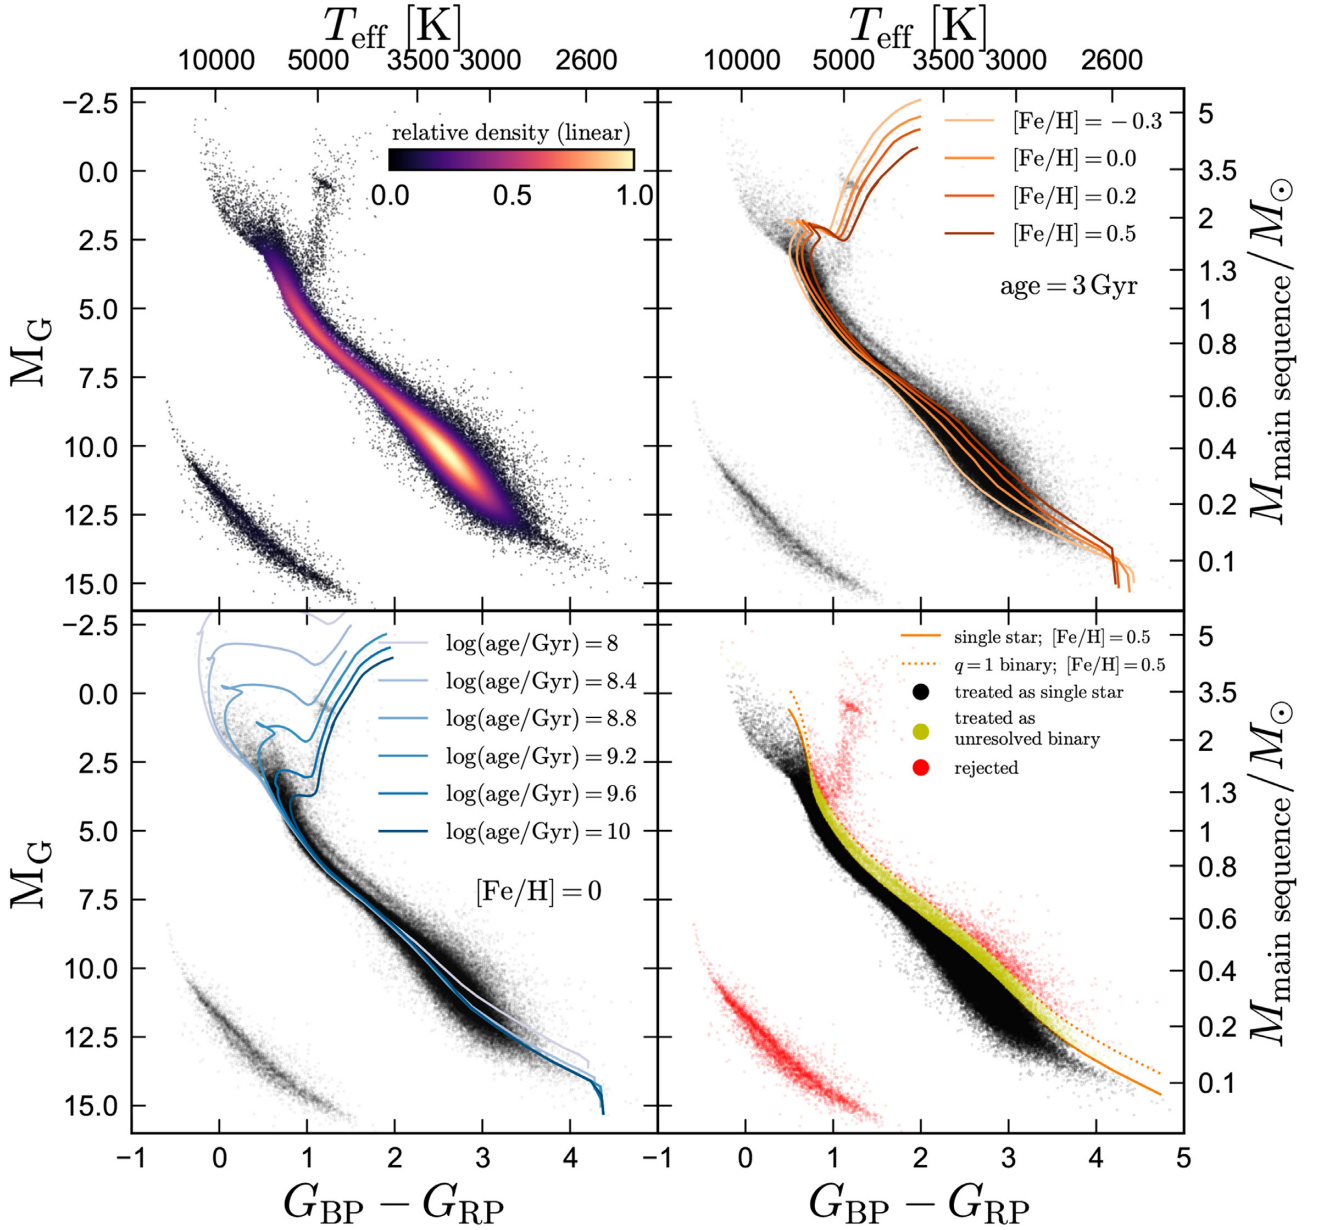
Figure 1. CMD of all stars (both primaries and secondaries) in the [ER18](#) catalogue of resolved wide binaries. Right *y*-axes and upper *x*-axes show the mass and effective temperature corresponding to a given magnitude and colour for main-sequence stars with [Fe/H] = 0. In the upper left panel, points are coloured by a Gaussian kernel density estimate of the local density. We compare PARSEC isochrones with a range of metallicity (upper right) and age (lower left) to the data. The bottom right panel divides the CMD into three regions. Black points (main-sequence components with no bright unresolved sub-components) comprise the bulk of our sample, and their masses are estimated using single-star isochrones. Yellow points (components of wide binaries with a bright unresolved companion) are kept, but their masses are estimated using unresolved binary isochrones. Red points (pre-main-sequence stars, unresolved triples, and evolved stars) are discarded.

binaries with similar *s* have different angular separations at different distances.

2.1 Is the twin excess real?

To test whether the narrow excess of binaries with $\Delta G \approx 0$ is a real astrophysical effect (as opposed to a data artefact), we produced a control sample of chance alignments with similar observable properties to the real binaries. This sample was produced by repeating the procedure used to produce the real binary catalogue (applying

the same quality cuts and limits on photometric and astrometric precision), but requiring that the two stars have parallaxes and proper motions that are *inconsistent* rather than consistent. This selects pairs of stars that are close together on the sky (and thus are affected by contamination and blending in the same way as real binaries) but are not physically associated. We applied the same procedure for removing resolved higher order multiples and potential members of bound and dissolving clusters that was used for real binaries. Because chance alignments are intrinsically rare at close angular separations, we searched out to 400 pc (rather than 200 pc for the fiducial

$$M_1 \approx 1 M_\odot; \quad 0.1 \lesssim q \lesssim 1$$

$$M_1 \approx 0.5 M_\odot; \quad 0.25 \lesssim q \lesssim 1$$

$$M_1 \approx 0.5 M_\odot; \quad q \approx 1$$

Figure 2. SDSS images of a selection of binaries from our catalogue. Top panel shows binaries in which the primary is a solar-type star ($0.9 < M_1/M_\odot < 1.1$) and the mass ratio varies from ≈ 0.1 to 1. The primary is at the centre of each image, and the secondary is circled. Middle panel shows binaries in which the primary is a late K dwarf ($0.4 < M_1/M_\odot < 0.6$), again for a range of mass ratios. Bottom panels shows examples of 'twin' binaries with near-identical components, each with $0.4 < M_1/M_\odot < 0.6$. Each image is 100 (top panel) or 45 (bottom panels) arcsec on a side. Our analysis uses photometry from *Gaia*, not SDSS. Because the SDSS photometry is ground based, blending and source contamination affect it more severely.

binary catalogue) to obtain better statistics. We verified that our conclusions are unchanged when only the sample within 200 pc is considered.

In Fig. 4, we compare the distributions of magnitude and colour difference for real binaries in our catalogue (left-hand panels) and chance alignments (right-hand panels). Chance alignments are more common at large angular separations. The broad distribution and outer envelope of ΔG and $\Delta(G_{BP} - G_{RP})$ at a given θ is similar

for binaries and chance alignments, reflecting *Gaia*'s contrast sensitivity. For the chance alignments, there is no sharp excess of pairs with nearly-equal magnitude and colour. Because chance alignments are subject to same cuts on astrometric and photometric quality and signal to noise as the real binaries, any aspects of the *Gaia* source detection algorithm that might be expected to produce a bias towards equal-brightness pairs should affect real binaries and chance alignments very similarly. We therefore interpret the lack of a
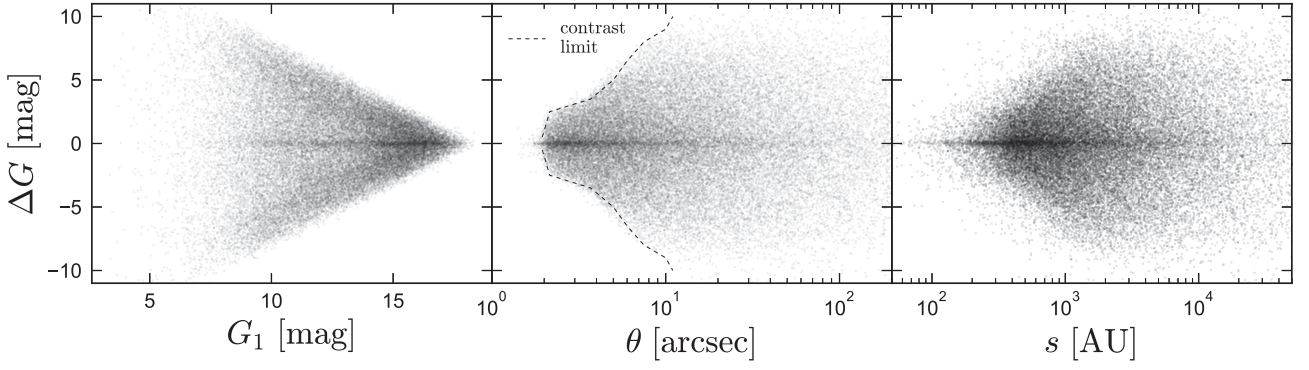
Figure 3. Apparent *G*-band magnitude difference between the two components of binaries in our catalogue as a function of the apparent magnitude of the primary (left), angular separation (middle), and projected physical separation (right). The sign of ΔG is randomized for easier visualization. There is a clear excess of equal-brightness binaries with $\Delta G \approx 0$. These 'twin' binaries are found over a large range of apparent magnitudes, preferentially at closer physical and angular separations. The middle panel shows the contrast sensitivity limit for our sample; at fixed angular separation, the probability of a companion passing our photometric quality cuts drops rapidly outside this limit due to source contamination (Appendix D2). The lack of binaries with large ΔG at small separations is a selection effect; the narrow excess at $\Delta G \approx 0$ is not.
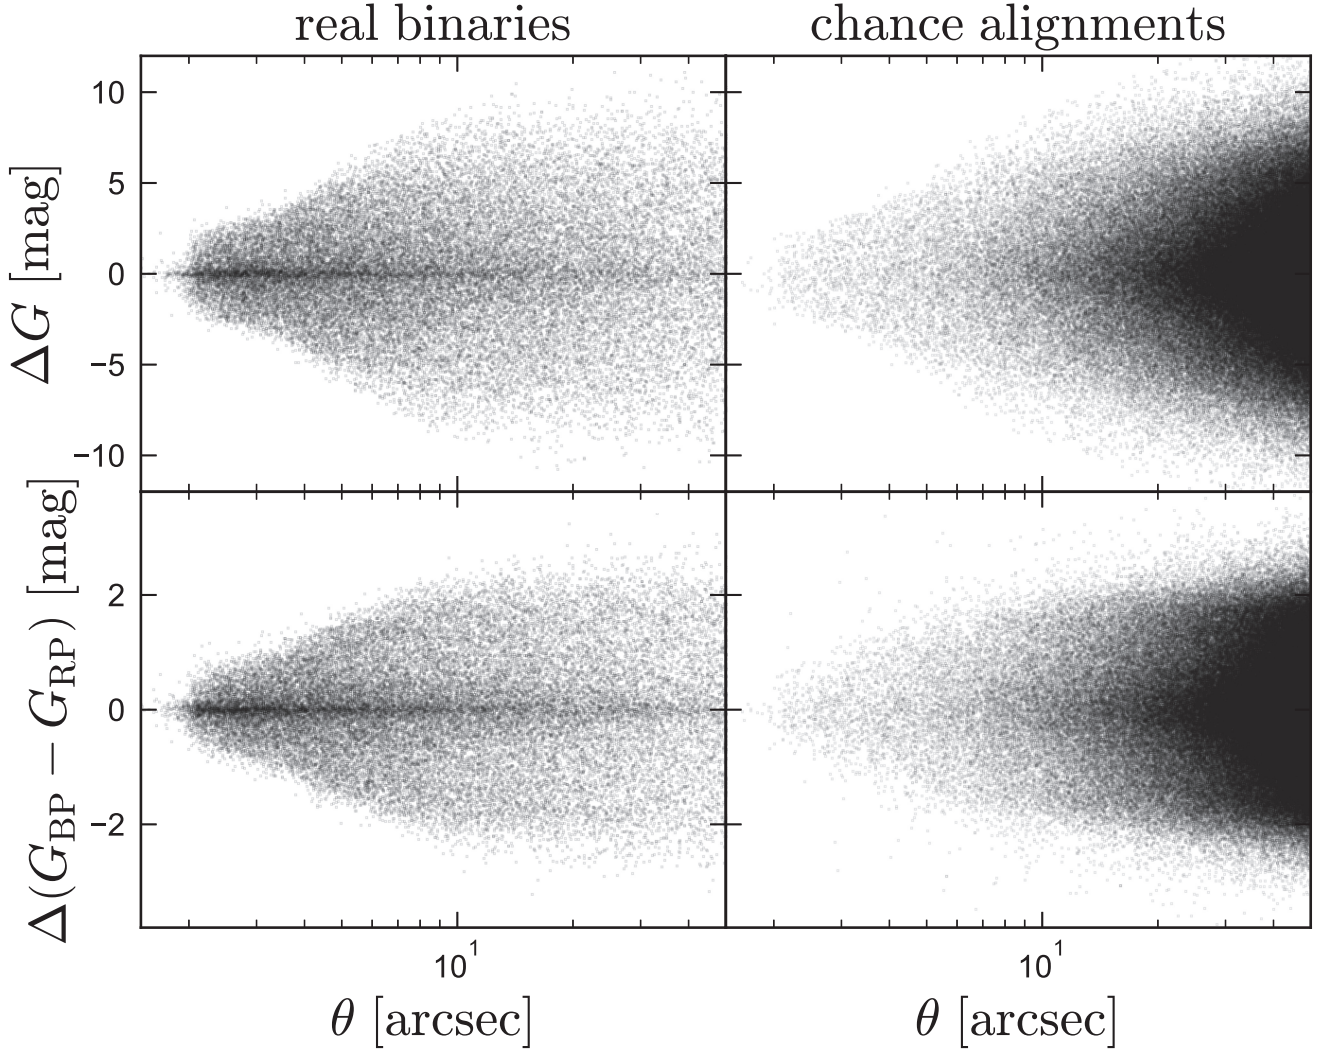
Figure 4. Difference in *G*-band magnitude (top) and colour (bottom) between the two components of pairs with a range of angular separations. We compare genuine binaries (left) to chance alignments (right). The chance alignments are required to pass the same quality cuts as the true binaries. A clear excess of pairs with $\Delta G \approx 0$ and $\Delta(G_{BP} - G_{RP}) \approx 0$ is evident for the true binaries, but not for the chance alignments. The absence of the excess among chance alignments bolsters our confidence that the feature found in the real binaries is due to a true excess of equal-mass twins and is not a selection effect or data artefact.
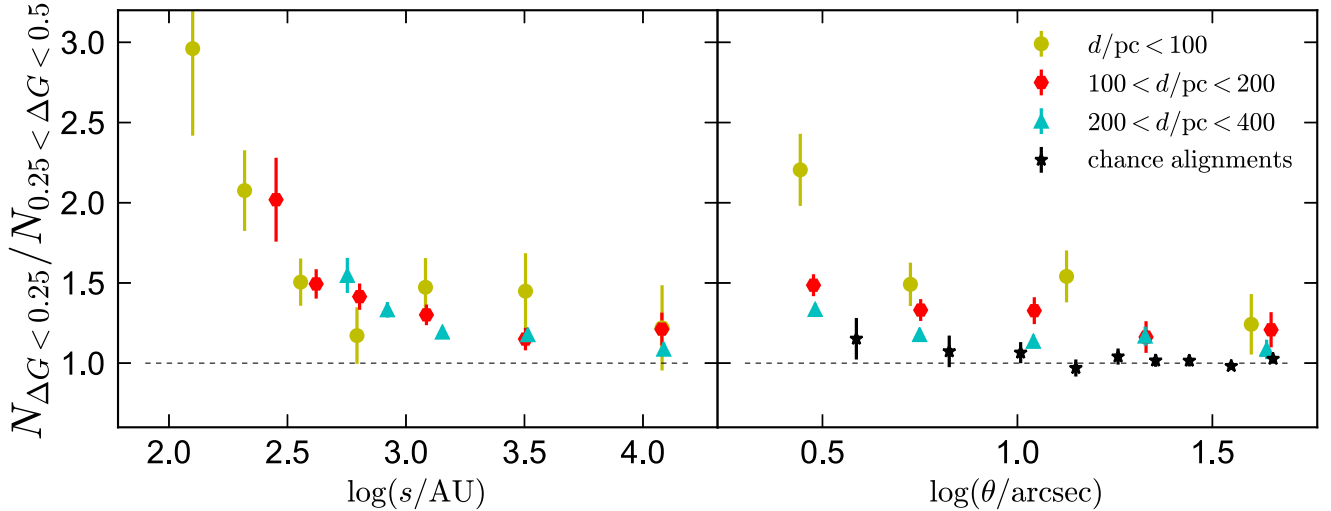
Figure 5. Ratio of the number of binaries with nearly equal magnitudes ($\Delta G < 0.25$) to the number with slightly different magnitude ($0.25 < \Delta G < 0.5$) as a function of physical separation (left) and angular separation (right). This ratio quantifies the excess of equal-brightness 'twins'. Error bars are 1σ. We compare binaries at different distances, as well as a control sample of chance alignments. At fixed *physical* separation, the twin excess is consistent with being independent of distance. At fixed *angular* separation, it decreases with distance. Together, these trends strongly suggest that the $\Delta G \approx 0$ excess (Figs 3 and 4) reflects a real excess of equal-mass binaries, not a selection effect or photometric issue at close angular separations.

thin excess at $\Delta G \approx 0$ and $\Delta(G_{BP} - G_{RP}) \approx 0$ for chance alignments as strong evidence that the twin excess among real binaries is astrophysical.

An initial worry was that apparent twins might be duplicate *Gaia* sources that were not properly removed: if a source were observed twice in two different scans without being identified as a duplicate, it could manifest in our catalogue as an apparent binary pair in which the two components had essentially identical astrometry and photometry. We verified that unrecognized duplicate sources are *not* the source of the twin signal using the SDSS images: we visually inspected the postage stamps of several hundred equal-brightness binaries in the SDSS footprint, and all the systems indeed contain two stars.

Another useful test to verify that the $\Delta G \approx 0$ excess is physical is to determine whether its strength depends primarily on physical or angular separation. This can be accomplished by comparing the ΔG distributions of binaries at different distances. If the excess of equal-brightness pairs were due to an observational bias against binaries that are not nearly equal mass (or if there were an issue with the *Gaia* photometry causing close pairs to erroneously have the same reported magnitude), one would expect the twin excess to depend on angular separation, manifesting itself at different *physical* separations for binaries at different distances. If the excess is due to an astrophysical preference for equal-mass binaries, then it should be primarily a function of s, manifesting itself at the same physical separation but different θ for binary samples at different distances.[3]

In Fig. 5, we compare the excess of equal-brightness pairs for binaries in three different distance bins as a function of physical separation (left) and angular separation (right). We measure the excess as the ratio of the number of binaries with $0 < \Delta G <$

0.25 (nearly-equal magnitudes) to the number with $0.25 < \Delta G < 0.5$ (slightly different, but still similar magnitudes). In order to include binaries at distance larger than 200 pc, we repeated the binary search from ER18, this time searching out to 400 pc. The additional binaries with $200 < d/\mathrm{pc} < 400$ are considered for this test only and are not included in the sample used to fit for the mass ratio distribution.

The left-hand panel of Fig. 5 shows that at fixed *physical* separation, the fractional excess of twins increases towards closer separations in a manner that is consistent across the three different distance samples. The right-hand panel shows that as expected, the twin excess at fixed *angular* separation varies with distance. At fixed θ, larger distances correspond to wider physical separation. Because the twin excess decreases with physical separation, it decreases with distance at fixed θ. The excess of equal-brightness chance alignments is also shown as a function of angular separation in the right-hand panel of Fig. 5. As expected, this is nearly consistent with 0 at all angular separations, meaning that there is no strong bias towards equal-brightness pairs compared to pairs with slightly different brightness. At the closest angular separations ($\theta \lesssim 5$ arcsec), there is a slight excess of equal-brightness chance alignments, suggesting that contrast sensitivity begins to play a role. The excess for the chance alignments is less than that found for the true binaries at all distances and is self-consistently accounted for in the selection function (Appendix D2).

Fig. 6 shows distributions of ΔG for binaries in five different bins of primary mass. Both the overall shape of the distribution and the strength of the excess at $\Delta G \approx 0$ vary substantially between mass bins. For the highest mass bin, $1.2 < M_1/\mathrm{M}_\odot < 2.5$, the excess at $\Delta G \approx 0$ is weak, and the observed distribution (without accounting for incompleteness) peaks at $\Delta G \approx 8$, corresponding to $q \approx 0.3$ (see Fig. 8). For $0.8 < M_1/\mathrm{M}_\odot < 1.2$ and $0.6 < M_1/\mathrm{M}_\odot < 0.8$, there is a clear excess of twins out to $s \approx 1000$ au, with the observed peaks in the ΔG distribution corresponding to $q \approx 0.45$ and $q \approx 0.6$, respectively. In the two lowest mass bins, the visible twin excess appears to extend to larger separations, and there is no secondary peak in the distributions of ΔG. For all mass

[3]An implicit assumption here is that the fractional excess of twins is independent of the intrinsic properties of a binary (e.g. mass), since binaries that pass our quality cuts at larger distances will be more massive on average than those that are nearby. We show in Section 4 that this assumption does not hold up in detail, which could lead to modest variation with distance in the observed twin excess at fixed physical separation.
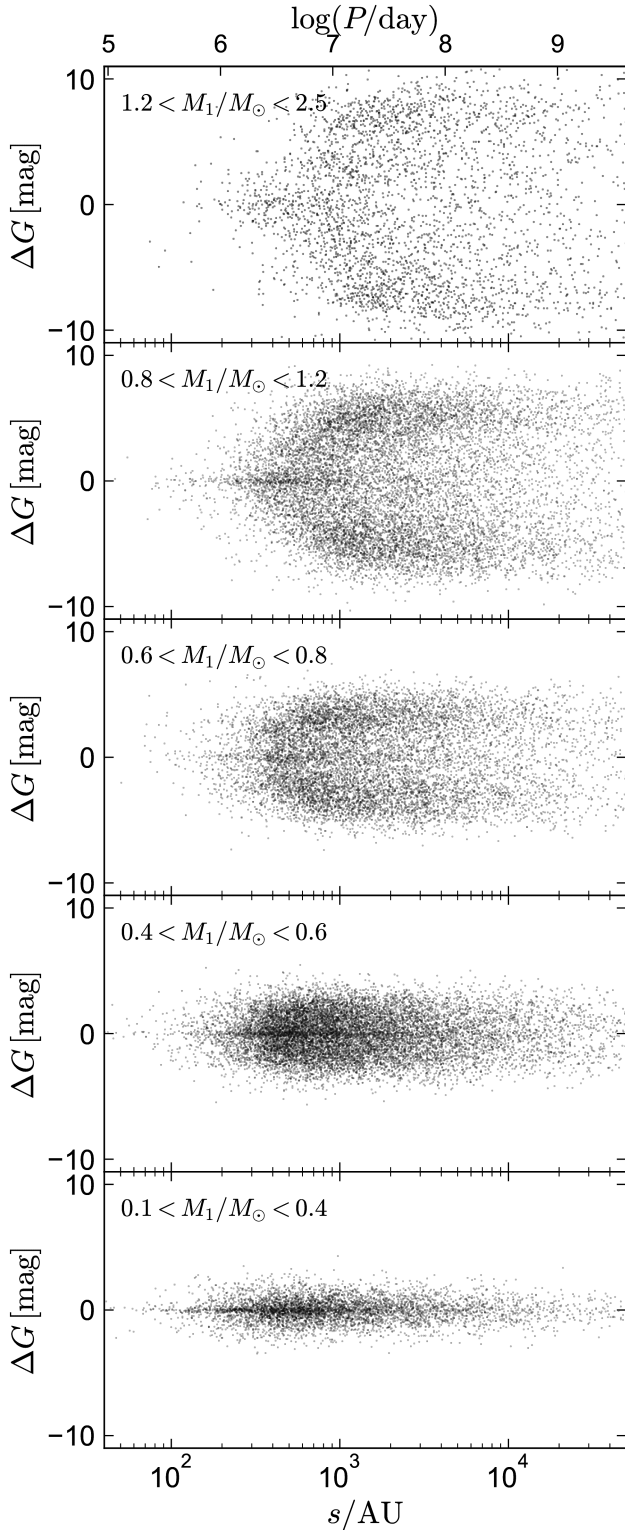
Figure 6. Magnitude difference as a function of projected physical separation for binaries with different primary masses. The sign of ΔG is randomized for easier visualization, and masses are computed from the CMD. The thin band of 'twin' binaries with $\Delta G \approx 0$ is clearly visible in all but the highest mass bin. In the top three bins, there is a clear preference for unequal-mass (large ΔG, low-q) binaries. Incompleteness to faint companions prevents the detection of large ΔG companions to low-mass secondaries.

bins, the binaries with the largest ΔG have secondaries with $M \approx 0.1 \, M_\odot$. The maximum ΔG in each panel is probably set largely by observational incompleteness, since at moderately old ages ($\gg 100 \, \mathrm{Myr}$), objects below the hydrogen burning limit are several magnitudes fainter than those just above it. Incompleteness at small separations due to the angular resolution and contrast sensitivity of our catalogue becomes significant at different projected physical separations for different mass bins, because the low-mass binaries that pass our quality cuts are on average at closer distances than the higher mass binaries.

3 MODELLING

We now turn to the intrinsic mass ratio distribution, $p(q)$. Because we expect that $p(q)$ may vary with primary mass and/or separation, we split the observed sample into 5 bins of primary mass and 7 bins of projected physical separation, fitting each of the resulting 35 subsamples independently. Our choice of bins is designed to balance the number of binaries in each bin while still covering a large dynamic range of mass and separation. We constrain $p(q)$ for each subsample by forward-modelling a simulated population of binaries with a given distribution of primary mass, age, metallicity, distance, physical separation, and mass ratio, passing them through the selection function, and comparing to the data. The 'data' we consider is the distribution of the observed binaries in the three-dimensional space of angular separation θ, magnitude difference ΔG, and parallax ϖ. The best-fitting $p(q)$ is then the one that best matches the observed data, and uncertainties in $p(q)$ are estimated via Markov chain Monte Carlo from the range of $p(q)$ that adequately reproduce the observed data. This approach requires a parametrized form of $p(q)$ (Section 3.1), knowledge of the selection function (Section 3.2), and a Galactic model from which simulated binaries are drawn (Section 3.3).

To generate a model prediction for a given set of fitting parameters, we generate a population of $N = 10^6$ synthetic binaries and forward-model their distribution into the space of observables. Masses, ages, metallicities, and distances are Monte Carlo sampled from their respective distribution functions. For mass ratios and separations, we use a regular 1000×1000 grid, weighting the synthetic binary at each gridpoint by the mass ratio and separation distributions. This approach is chosen because in the fiducial model, we fit for the mass ratio and separations distributions but leave the distributions of age, mass, distance, and metallicity fixed. The number of synthetic binaries generated must be large enough that Poisson noise is negligible. We verified that $N = 10^6$ binaries in each mass and separation bin is large enough that our constraints are converged and insensitive to the random seed.

Synthetic photometry is calculated for both components from isochrones (including the effects of unresolved companions in hierarchies; see Section 3.3), and the observable properties of each binary are passed through the selection function. In constructing the distribution of mock-observables to be compared with the data, each synthetic binary is weighted by the selection function evaluated for its observables. Finally, the observed and simulated distributions are binned on a regular 3D grid. We default to using 100 bins in ΔG, 100 bins in $\log\theta$, and 5 bins in ϖ (because typical errors in ϖ are larger than those in ΔG or θ). The maximum value of the ΔG grid for a given primary mass bin is chosen to include the largest ΔG value for the data in that bin, and the 5 bins in ϖ are chosen so that roughly the same number of observed binaries fall in each bin. Our constraints are not sensitive to the choice of bins,

since they are small compared to the scale on which the data exhibit substructure.

We re-scale the binned model prediction such that the total counts match the observed data. We calculate the likelihood for a particular set of model parameters by summing over all cells in the 3D distribution, assuming that the distribution of counts in each cell is set by a Poisson process. The log-likelihood function is

$$\ln L = \sum_{m_i \neq 0} \left[d_i \ln m_i - m_i - \ln(d_i!) \right], \qquad (1)$$

where m_i and d_i are the counts in the ith cell of the binned model and data ('!' denotes a factorial). We sample from the posterior using EMCEE (Foreman-Mackey et al. 2013), using priors described in Section 3.1. We use 200 walkers and draw 20 000 samples for each bin of mass and separation after a burn-in period of 200 steps per walker. Inspecting the chains, we find this to be sufficient for convergence in all cases. We carried out tests with mock data that was drawn from a known mass ratio distribution assigned with realistic observational uncertainties in order to verify that our approach yields unbiased constraints on $p(q)$.

This fitting procedure is qualitatively very similar to the method commonly used to constrain population properties such as the IMF, star formation history, unresolved binary fraction, or initial–final mass relation from CMDs (e.g. Dolphin 2002; Bonatto, Lima & Bica 2012; Geha et al. 2013; El-Badry, Rix & Weisz 2018c). The difference between our approach and these studies is that we are forward-modelling the distribution of binaries in the space of angular separation, magnitude difference, and parallax rather than colour and magnitude. The approach can in principle be generalized to include other observables, such as colour difference or apparent magnitude of the primary, but the computational expense increases rapidly with the dimension of the space in which the likelihood function is calculated.

We note that the data uncertainties do not enter equation (1). The implicit assumption (which does hold for our problem setup) is that the uncertainties are small compared to the scale of the bins in all quantities. We also note that for a fine grid, a majority of grid cells will contain 0 or 1 real binaries. This is not a problem; equation (1) does not make any assumptions about the magnitude of d_i.

3.1 Parametrization

We fit the mass ratio distribution by assuming a parametrized form of $p(q)$ and then obtaining constraints on the parameters. Our fiducial parametrization is motivated by the one used by Moe & Di Stefano (2017) and is shown in Fig. 7. $p(q)$ is parametrized as a broken power law with logarithmic slope γ_{smallq} at $q < q_{\mathrm{break}}$ and γ_{largeq} at $q > q_{\mathrm{break}}$. A possible excess (or deficit) of twins is added on top of the power-law component at $q > q_{\mathrm{twin}}$. This excess is modelled as a step function, with the magnitude such that the integrated excess of twins divided by the total number of binaries with $q > 0.3$ is F_{twin}. The reason for this choice (as opposed to normalizing relative to *all* binaries) is that, compared to the constraints at $q > 0.3$, the constraints on $p(q)$ at small q are often weak due to incompleteness. $p(q)$ is set to 0 at $q < 0.05$ to prevent divergence when $\gamma_{\mathrm{smallq}} < -1$. Because we are not sensitive to companions with $q < 0.05$, this has little effect on our results.

There is no a priori motivation for this particular parametrization: because the physics that set $p(q)$ are imperfectly understood, we simply require a functional form that is sufficiently flexible to reproduce the observed distributions of ΔG. We have experimented with other forms of $p(q)$, including adding an additional break point to the
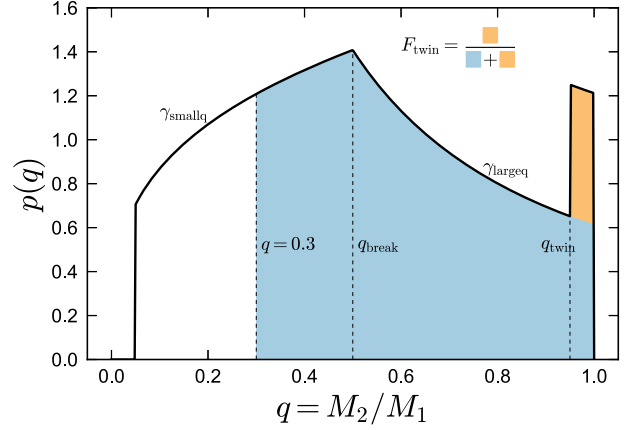
Figure 7. Parametrized mass ratio distribution. The broad part of the distribution is modelled with a broken power law of logarithmic slope γ_{smallq} at $q < q_{\mathrm{break}}$ and γ_{largeq} at $q > q_{\mathrm{break}}$. F_{twin} is the *excess* fraction of nearly equal-mass binaries with $q > q_{\mathrm{twin}}$, relative to the underlying power-law distribution for $q > 0.3$. For this particular example, $F_{\mathrm{twin}} = 0.04$, $q_{\mathrm{twin}} = 0.95$, $\gamma_{\mathrm{smallq}} = 0.3$, $\gamma_{\mathrm{largeq}} = -1.3$, and $q_{\mathrm{break}} = 0.5$.

power law and modelling the twin excess as ramping up linearly instead of increasing stepwise. However, we find the form shown in Fig. 7 to provide a good fit in all mass and separation bins, with more complicated models providing only marginal improvements. We thus use this functional form for our main analysis. In Appendix A, we show results of using alternative parametrizations, including one that smoothly transitions between the two power laws instead of including a sharp break and one that leaves the shape of the twin excess flexible.

There are two differences between our parametrization of $p(q)$ and the one used by Moe & Di Stefano (2017), who focused primarily on binaries with $M_1 \gtrsim \mathrm{M}_\odot$. First, they fixed $q_{\mathrm{break}} = 0.3$, a value that was appropriate for their analysis because several studies have found the mass ratio distribution to peak at $q \approx 0.2 - 0.3$ for binaries with $M_1 \gtrsim \mathrm{M}_\odot$ (Duquennoy & Mayor 1991; Gullikson et al. 2016; Murphy et al. 2018). We also use $q_{\mathrm{break}} = 0.3$ for our highest mass bin, but we find $q_{\mathrm{break}} = 0.5$ to provide a better fit at lower masses (see Appendix A2). Secondly, they fixed $q_{\mathrm{twin}} = 0.95$, roughly the value found for spectroscopic binaries (e.g. Tokovinin 2000). In order to identify or rule out trends with mass and separation, we leave q_{twin} as a free parameter.

Along with the parameters of the mass ratio distribution, we also fit for γ_s, the local logarithmic slope of the separation distribution in each bin of primary mass and separation. Opik's law (a uniform distribution of $\log(s)$) corresponds to $\gamma_s = 0$.

Our adopted priors are listed in Table 1. For the lowest mass bin, the smooth component of $p(q)$ is modelled as a single power law, because no useful constraints can be obtained at $q \lesssim 0.5$. Our priors on γ_{smallq} and γ_{largeq} are loosely motivated by constraints from the literature (particularly Tokovinin 2014 and Moe & Di Stefano 2017) but are fairly weak. We use a uniform flat prior on F_{twin}. We require $q_{\mathrm{twin}} > 0.93$ in order to prevent cases where a large F_{twin} is combined with $q_{\mathrm{twin}} \ll 1$ such that a broad 'twin excess' simply modifies the overall shape of $p(q)$ but does not actually correspond to a sharp increase near $q = 1$. We show in Appendix A1 that, where a twin excess is significant, it is always narrow ($q_{\mathrm{twin}} \gtrsim 0.94$).

Table 1. Summary of priors adopted in each primary mass bin. We use the same priors for all separation bins. $\mathcal{U}(a, b)$ represents a uniform distribution over $[a, b]$, and $\mathcal{N}(\mu, \sigma)$ represents a normal distribution with mean μ and dispersion σ. γ_s is the logarithmic slope of the local separation distribution; i.e. $p(s) \propto s^{\gamma_s}$. Other parameters are described in Fig. 7.

	$0.1 < M_1/\mathrm{M}_\odot < 0.4$	$0.4 < M_1/\mathrm{M}_\odot < 0.6$	$0.6 < M_1/\mathrm{M}_\odot < 0.8$	$0.8 < M_1/\mathrm{M}_\odot < 1.2$	$1.2 < M_1/\mathrm{M}_\odot < 2.5$
F_{twin}	$\mathcal{U}(-1, 1)$	$\mathcal{U}(-1, 1)$	$\mathcal{U}(-1, 1)$	$\mathcal{U}(-1, 1)$	$\mathcal{U}(-1, 1)$
q_{twin}	$\mathcal{U}(0.93, 1)$	$\mathcal{U}(0.93, 1)$	$\mathcal{U}(0.93, 1)$	$\mathcal{U}(0.93, 1)$	$\mathcal{U}(0.93, 1)$
γ_{largeq}	$\mathcal{N}(0.5, 1)$	$\mathcal{N}(-1, 1)$	$\mathcal{N}(-1, 1)$	$\mathcal{N}(-1, 1)$	$\mathcal{N}(-1, 1)$
γ_{smallq}	–	$\mathcal{N}(0.5, 0.5)$	$\mathcal{N}(0, 0.5)$	$\mathcal{N}(0, 0.5)$	$\mathcal{N}(0, 0.5)$
γ_{s}	$\mathcal{N}(-1.5, 1)$	$\mathcal{N}(-1.5, 1)$	$\mathcal{N}(-1.5, 1)$	$\mathcal{N}(-1.5, 1)$	$\mathcal{N}(-1.5, 1)$
q_{break}	–	0.5	0.5	0.5	0.3

3.2 Selection function

Because both components of a binary must pass astrometric and photometric quality cuts, the binaries in our catalogue are relatively bright. The median apparent magnitude of all stars in our fitting sample (considering primaries and secondaries together) is $\langle G \rangle =$ 14.5, and 90 per cent (99 per cent) of stars fall in the range $9.1 < G < 17.4$ ($6.6 < G < 18.1$). For stars in this magnitude range, *Gaia* DR2 is nearly complete outside of crowded fields (Arenou et al. 2018; Sollima 2019). The completeness is not quite 100 per cent due to a variety of issues, but it is primarily a function of position on the sky, not colour or magnitude. The selection function for our sample is thus determined by the cuts imposed on astrometric and photometric precision.

In order for a binary to appear in the catalogue, (a) both components must be bright enough that they individually pass the cuts we impose on `parallax_over_error` and photometric signal to noise and (b) they must not be so close on the sky that the photometry of either component is significantly contaminated. The selection function for binaries can thus be expressed as a product of the two components' single-star detection probabilities and a contrast sensitivity cross-term:

$$s_{\mathrm{binary}} = s_1 \times s_2 \times s_{\Delta G}(\theta). \tag{2}$$

Here, s_1 and s_2 represent the independent probabilities of detecting an isolated star with the observable properties of star 1 and star 2; they depend primarily on apparent magnitude and colour. $s_{\Delta G}(\theta)$ quantifies the reduction in the probability of detecting star 1 and star 2 together, *relative to the probability of detecting them at asymptotically large separation*. It is primarily a function of the angular separation of the two stars and their flux ratio. For example, $s_{\Delta G}(\theta) \approx 0$ at $\theta < 2$ arcsec and 1 at $\theta > 10$ arcsec; at intermediate separations, it depends strongly on ΔG. We calculate the single-star selection function given our quality cuts in Appendix D1 and $s_{\Delta G}(\theta)$ in Appendix D2. The derived selection functions are then validated in Appendix E.

3.3 Model inputs

We draw primary masses and system ages assuming a Kroupa (2001) IMF and a constant star formation history over the last 10 Gyr. Because suspected members of bound and dissolving clusters are removed from our binary catalogue, we remove synthetic binaries with age < 100 Myr. We assume that the wide binary fraction, f_{wb}, scales with mass as $f_{\mathrm{wb}} \propto M_1^{\alpha_{\mathrm{wb}}}$, where $\alpha_{\mathrm{wb}} = 0.4$ is a constant (i.e. higher mass primaries are more likely to have wide binary companions). The effect of this assumption is that primary masses are drawn from a distribution with logarithmic slope $\alpha_{\mathrm{IMF}} + \alpha_{\mathrm{wb}}$, where α_{IMF} is the local logarithmic slope of the IMF ($\alpha_{\mathrm{IMF}} = -2.3$ for $M_1/\mathrm{M}_\odot > 0.5$ and $\alpha_{\mathrm{IMF}} = -1.3$ for $M_1/\mathrm{M}_\odot < 0.5$). We

find that with this choice of α_{wb}, our model predicts a distribution of primary magnitudes in reasonably good agreement with that of the binary catalogue when the selection function is taken into account.

We model the intrinsic spatial distribution of all stars as a plane-parallel exponential with the Sun at the mid-plane. The exponential scale height increases with stellar age (e.g. Nordström et al. 2004; Seabroke & Gilmore 2007), because older stars (a) were born from kinematically hotter gas and (b) have been dynamically heated more since their formation (e.g. Ting & Rix 2018). We use a fit to the empirical age-scale height relation recently measured by Sollima (2019) using *Gaia* star counts:

$$\log(h_z/\mathrm{pc}) = 0.53 \log(\mathrm{age}/\mathrm{yr}) - 2.65. \tag{3}$$

That is, the scale height increases from 40 pc for stars of age 100 Myr, to 130 pc at age 1 Gyr, to 450 pc at age 10 Gyr. We show in Appendix E1 that this leads to a predicted distance distribution in good agreement the data.

We use the tabulated empirical metallicity distribution function (MDF) for our binary catalogue that was measured in El-Badry & Rix (2019) by considering a subset of ~ 8000 binaries in the catalogue for which [Fe/H] was measured spectroscopically for at least one component. Most of the binaries in our catalogue are disc stars, with a median metallicity of $\langle[\mathrm{Fe/H}]\rangle \approx -0.1$ and tails extending to $[\mathrm{Fe/H}] \approx -1.0$ and $[\mathrm{Fe/H}] \approx +0.4$. We do not include any variation in the MDF with age or distance.

We generate synthetic photometry in the *Gaia* DR2 bands from Evans et al. (2018) using PARSEC[4] isochrones (Bressan et al. 2012; Chen et al. 2014). Just as for the real data, we remove synthetic binaries in which either component has evolved off the MS. For companions with $M_2 < 0.1\,\mathrm{M}_\odot$, we supplement the PARSEC models with `BT-Settl` models for very low-mass stars and brown dwarfs (Allard, Homeier & Freytag 2012; Allard 2014).[5] We include companions with masses as low as $0.01\,\mathrm{M}_\odot$ in our model for completeness. However, these have very little effect on our results, because substellar companions are too faint and red to pass our quality cuts except at very young ages.

Fig. 8 shows the relation between q and ΔG predicted for PARSEC isochrones. We show relations for five different primary masses, corresponding approximately to the median masses of primaries in each of our five mass bins, and use a minimum companion mass of $0.1\mathrm{M}_\odot$ in all cases. Comparing to Fig. 6, it is evident that the lowest mass secondaries have $M_2 \approx 0.1\mathrm{M}_\odot$ in all bins of primary mass. As for the excess of equal-brightness binaries,

[4]http://stev.oapd.inaf.it/cgi-bin/cmd
[5]Synthetic photometry in *Gaia* DR2 bands is computed for the `BT-Settl` models using the Pheonix web simulator, available at phoenix.ens-lyon.fr/simulator-jsf22-26/index.faces.
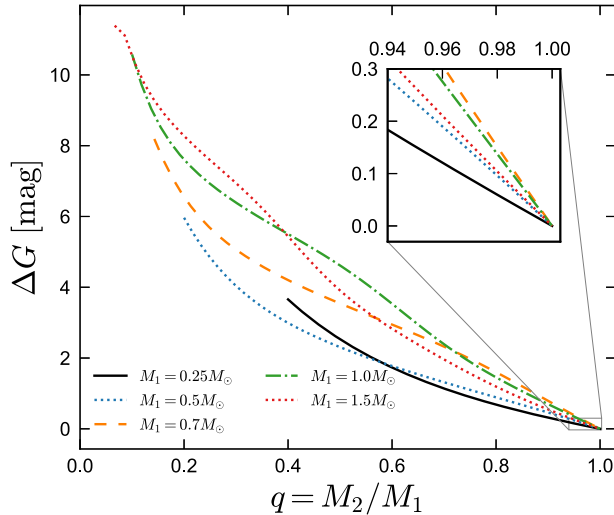
Figure 8. Magnitude difference versus mass ratio for binaries of different primary mass, computed from PARSEC isochrones. We assume both components have [Fe/H] = 0, adopting an age of 1 Gyr for the highest primary mass and 5 Gyr for the others. Because the mass–luminosity relation varies with mass, the relation between ΔG and q does as well. A magnitude difference of 0.25 mag (roughly where the twin excess becomes significant in our catalogue) corresponds to a mass ratio of 0.93–0.97, depending on primary mass.

it is primarily manifest over $0 < \Delta G < 0.25$ (though there is some variation with mass; see Fig. 10), which corresponds to $0.93 \lesssim q < 1$.

Some components of wide binaries have their own spatially unresolved close companions. We incorporate these in our models following observed binary statistics. The probability that a star has a close companion is taken to be a function of its mass, increasing from 20 per cent at $M < 0.2\,M_\odot$, to 30 per cent at $0.2 < M/M_\odot < 0.5$, to 35 per cent at $0.5 < M/M_\odot < 0.8$, to 45 per cent at $M/M_\odot > 0.8$ (Duchêne & Kraus 2013). For components that are assigned a binary companion, we draw the unresolved companion mass from a mass ratio distribution that is uniform between $q_{\min} = 0.1\,M_\odot/M_1$ and 1. We assume that the two components' probabilities of having a close companion are independent and neglect dynamical stability constraints. This prescription reproduces the morphology of the observed CMD reasonably well. In particular, the fraction of stars within 200 pc that fall in the 'suspected unresolved binary' region of the CMD (yellow points in the bottom right panel of Fig. 1) at a given M_G is reproduced within $\sim$25 per cent over $0.1 < M/M_\odot < 1.0$.

We do not include extinction or reddening due to dust in our model and do not attempt to correct the data for it. Because the stars in our catalogue are nearby ($d < 200$ pc), the effects of extinction are expected to be modest. The morphology of the CMD (e.g. the compactness of the red clump and MS; see Fig. 1) validates this assumption. Moreover, because the two components of a binary have similar position on the sky (within an arcminute in most cases) and similar distance, the extinction towards both components is expected to be similar.

We discuss the sensitivity of our constraints on the mass ratio distribution to various model ingredients in Appendix B. The largest systematic uncertainties come from the choice of stellar models. All the systematics we consider primarily affect constraints at low q. This is also true for uncertainties in the completeness function: because the two components of binaries with $q \approx 1$ have similar magnitudes and colours, incompleteness affects them similarly. The

translation between the distribution of ΔG and $p(q)$ is thus more straightforward at $q \approx 1$ than at $q \ll 1$.

4 RESULTS

Our fitting produces samples from the posterior distribution of free model parameters for each bin of mass and separation. These samples translate to marginalized constraints on each parameter of the mass ratio distribution and the covariances between them. An example for a single mass and separation bin is shown in Fig. B1. Marginalized constraints for all bins are listed in Appendix G.

The 35 panels of Fig. 9 show median and 2σ constraints on the mass ratio distribution for 35 bins of primary mass and separation. We do not show the *y*-axis ticks to avoid cluttering the figure, as the limits are different in each panel. The uncertainties (shaded regions) are derived by drawing 1000 samples from the posterior, calculating a normalized $p(q)$ for each sample, and then shading the middle 95.4 per cent range of these samples at each q. Solid black lines show the median at each q. Because $p(q)$ is normalized, uncertainty in the mass ratio distribution at *any* q translates to uncertainty in the normalization of $p(q)$ at *all* q. This reflects the fact that it is impossible to know the total fraction of all binaries that fall in some mass ratio range if one does not know how many binaries with low-mass ratios are missed. However, γ_{largeq}, the slope of the mass ratio distribution at $q > q_{\text{break}}$, is usually well constrained. Dashed vertical lines in each panel show the lowest mass ratio observed binary included in that bin. This roughly corresponds to the observational completeness limit and marks the mass ratio below which meaningful constraints cannot be obtained.

The fact that $p(q)$ is modelled as a double power law with a sharp break leads to some unphysical artefacts in Fig. 9, including a sharp change in slope at $q = q_{\text{break}}$ and artificially low uncertainty just below q_{break}, which is caused by $p(q)$ 'pivoting' about this point as γ_{smallq} varies and $p(q)$ is renormalized. We show in Appendix A3 that these features are not present when we fit a more flexible 'smoothly broken' power-law model. However, doing so introduces parameter covariances that are not present for the fiducial form of $p(q)$. We use the simpler sharply broken power law as our fiducial model to facilitate easier comparison between different mass and separation bins and comparison with the literature.

Fig. 10 compares the observed 1D distributions of ΔG to the predictions of the best-fitting model; this is useful for assessing the quality of the fits. In each panel, we note the number of observed binaries and the marginalized 1σ constraints on F_{twin}. In generating the model predictions, we use the median of the 1D marginalized posterior distribution for each free parameter. We then generate a Monte Carlo realization of the binary population, drawing primary masses, mass ratios, separations, ages, metallicities, and distances from the appropriate joint distributions, calculating synthetic photometry, and weighting each binary by the selection function evaluated for its observables.

Overall, the model predictions are in good agreement with the observed distributions of ΔG. This indicates that our parametrization of $p(q)$ is suitable and sufficiently flexible. In panels where F_{twin} is inconsistent with 0, there is a clear excess of equal-brightness binaries. The distributions of ΔG near $\Delta G = 0$ are also adequately reproduced by the model, indicating that the simple 'step function' model for the twin excess is consistent with the data. Although it is not shown here, we also find the best-fitting models to predict distributions of other observables (angular separation, parallax, and apparent magnitude of the primary) in good agreement with the data.
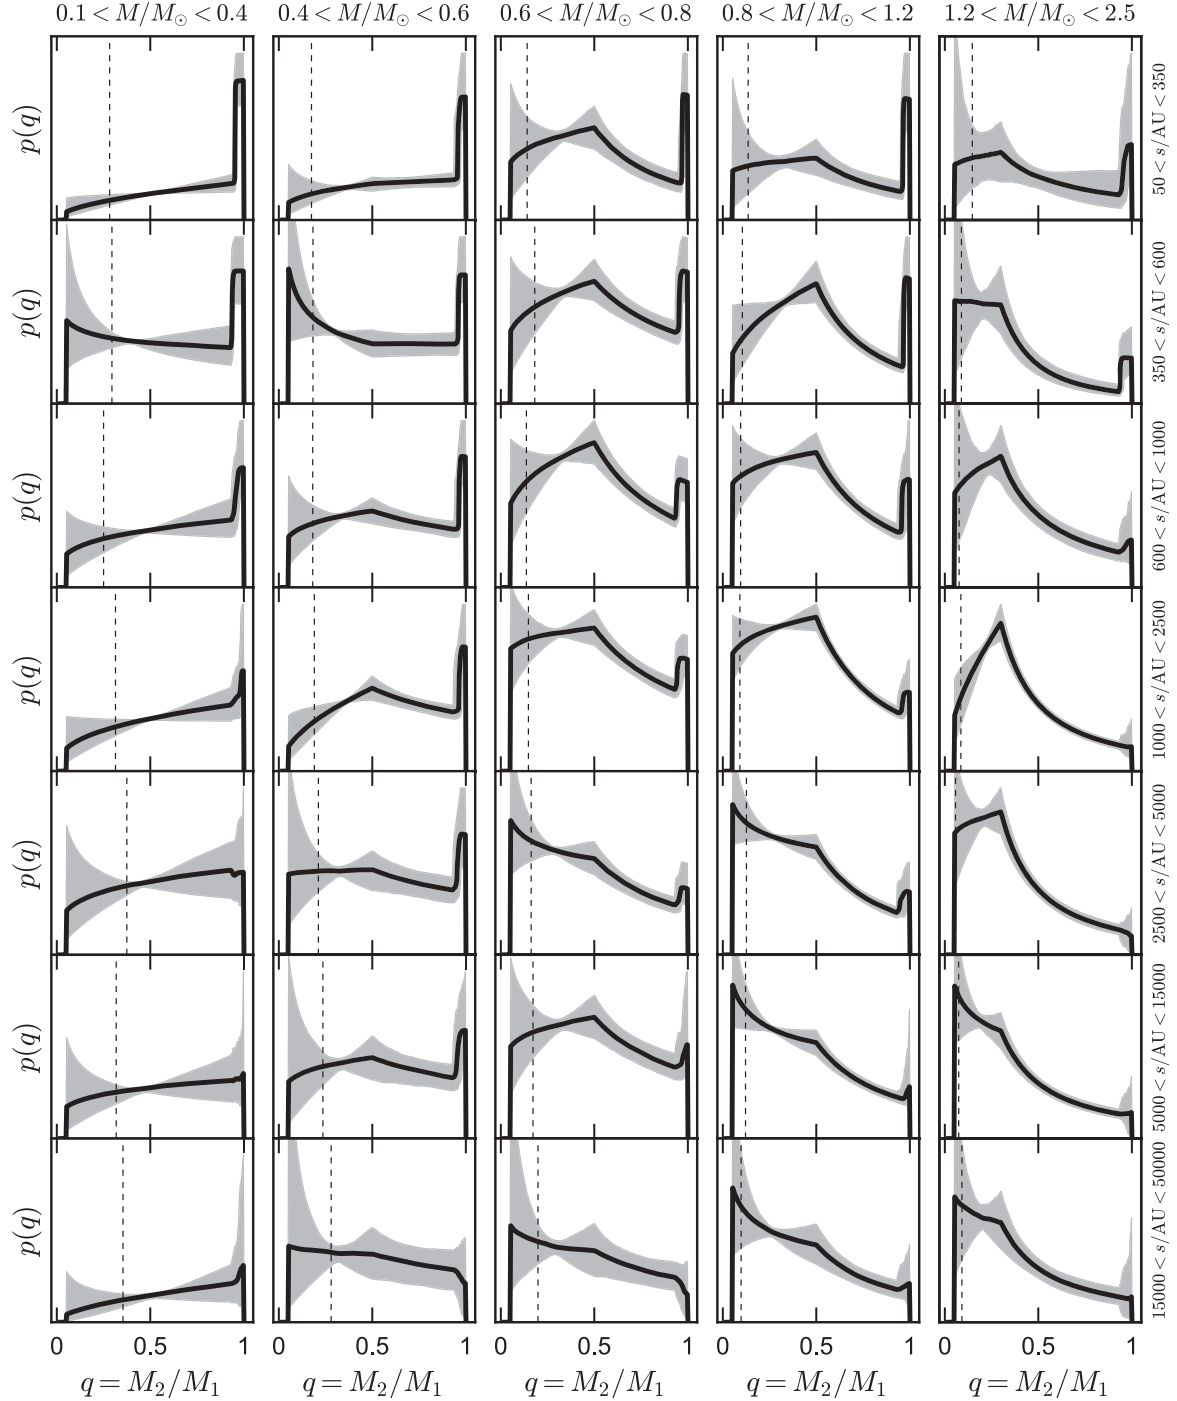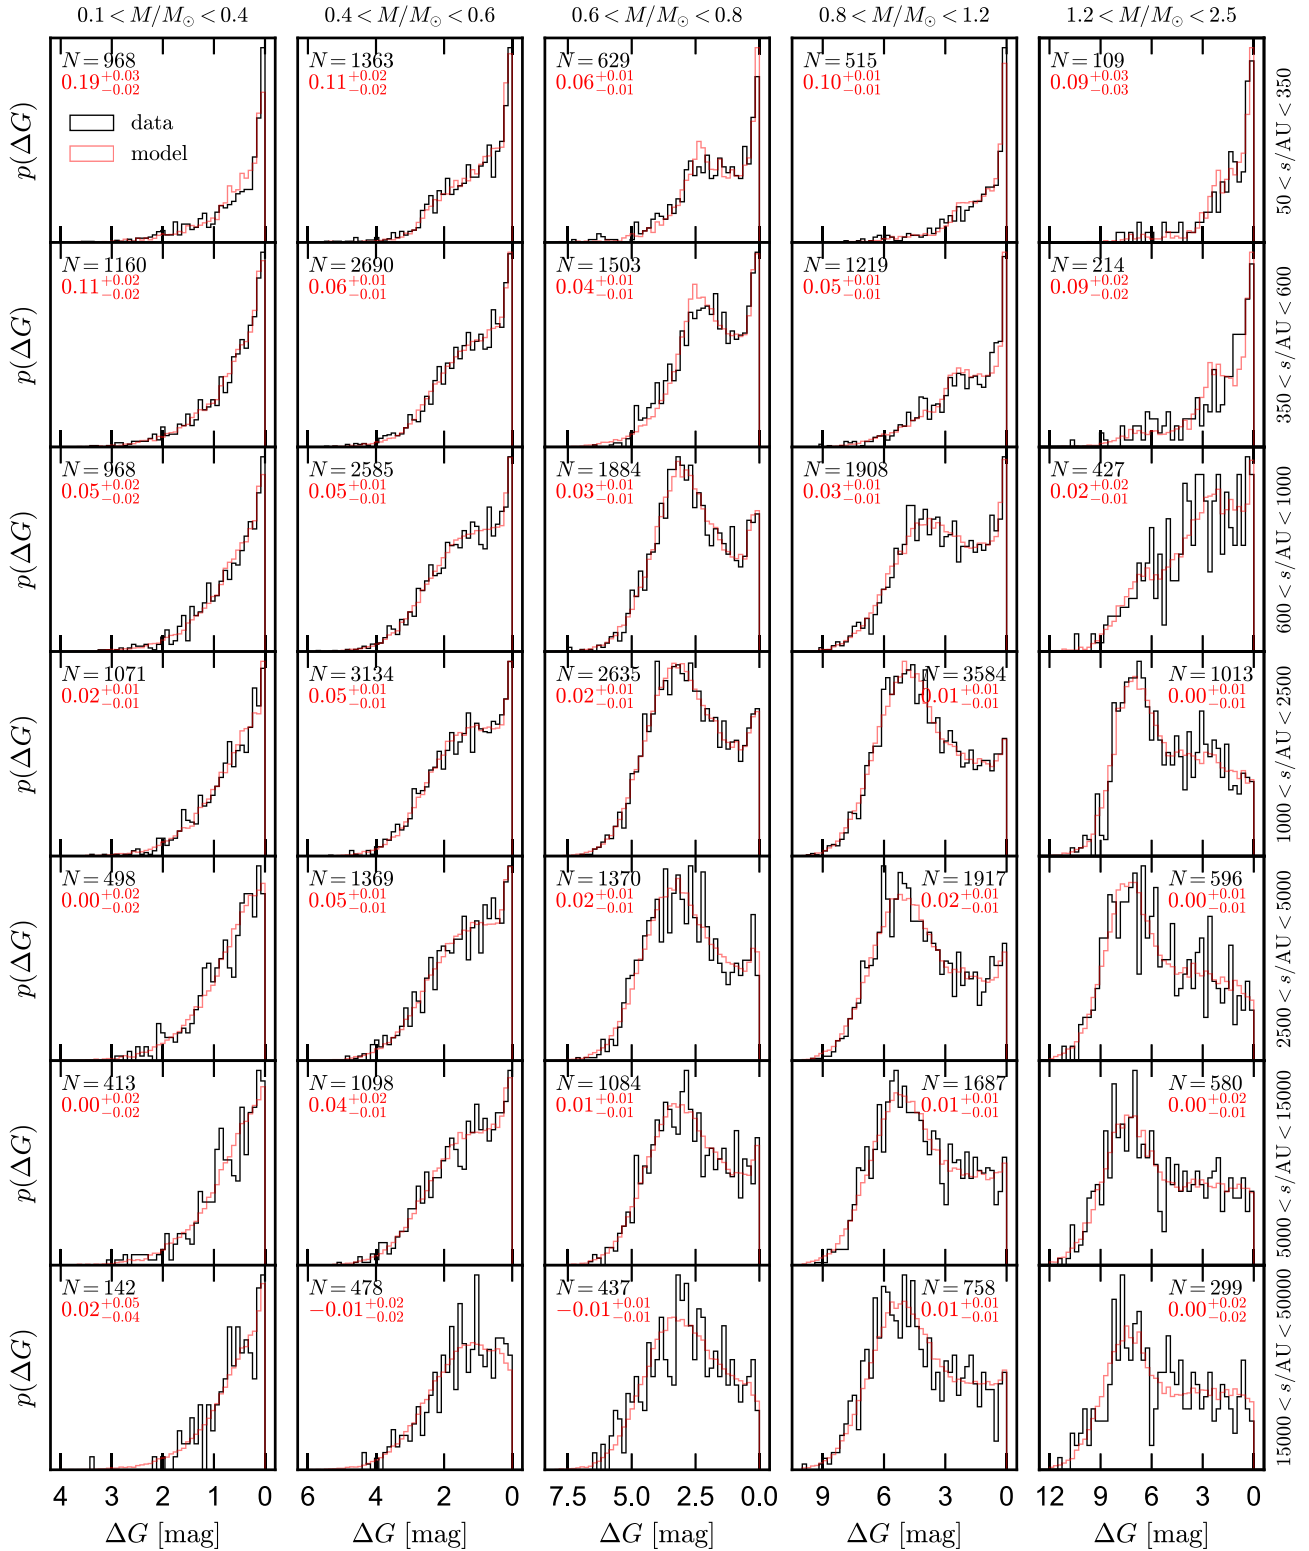
Figure 9. Mass ratio distribution constraints. Each panel corresponds to a single bin of primary mass (increasing left to right) and physical separation (increasing top to bottom). Grey shaded regions show 95.4 per cent probability. Normalization is arbitrary; the scale of the *y*-axis is linear and begins at 0. Vertical dashed lines show the lowest mass ratio of observed binaries in each bin.

Fig. 11 shows the constraints on $F_{\rm twin}$ as a function of mass and projected physical separation. Light and dark error bars show middle 95.4 per cent and 68.2 per cent of the marginalized posterior distributions. Note that the *y*-axis scale is different in the top panel. In all bins of primary mass, the excess twin fraction falls with increasing separation and is negligible at the largest separations we consider ($s > 15\,000$ au). At fixed separation, the magnitude of the excess varies with primary mass. For close separations ($50 < s/$au

< 350), $F_{\rm twin}$ is largest in the lowest bin of primary mass ($0.1 < M_1/{\rm M}_\odot < 0.4$). This may be in part because low-mass primaries in our sample are at closer distances on average, such that the median separation within the ($50 < s/$au < 350) bin is smaller than for higher mass primaries.

The maximum separation out to which there is a significant twin excess also varies with primary mass. For the highest mass bin, $F_{\rm twin}$ is consistent with 0 at $s > 600$ au. For solar-type primaries,

Figure 10. Magnitude difference between the two components of a binary (linear *y*-axis; normalization is arbitrary). Black histograms show real binaries, split into bins of primary mass (increasing left to right) and physical separation (increasing top to bottom). The twin excess can be seen in many panels as an excess of binaries with $\Delta G \approx 0$. Red histograms show Monte Carlo populations generated from the best-fitting model. The number of observed binaries (black) and the marginalized 1σ constraints on F_{twin} (red) are listed in each panel. Overall, the model is quite successful in matching the observed distributions.
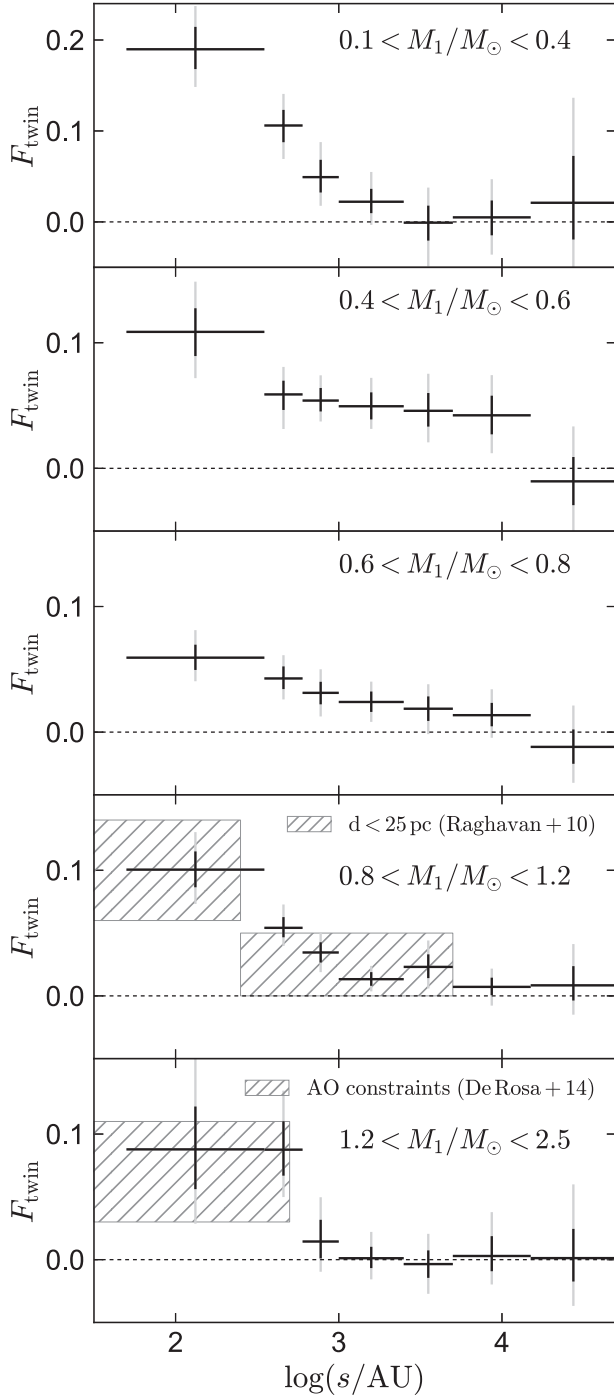
Figure 11. Marginalized constraints on F_{twin}, the fractional excess of binaries with nearly equal mass, versus physical separation. Each panel shows a separate bin of primary mass. Dark and light error bars show 1σ and 2σ constraints. In all mass bins, F_{twin} declines with increasing separation and is consistent with 0 at the largest separations. The twin excess reaches the widest separations for $0.4 < M_1/M_\odot < 0.6$, where F_{twin} is ~ 5 per cent out to 15 000 au. We compare to constraints from the nearby solar-type binary sample of Raghavan et al. (2010, panel 4) and the adaptive optics imaging survey from De Rosa et al. (2014, panel 5); these were derived by Moe & Di Stefano (2017).

F_{twin} is only consistent with 0 at $s > 5000$ au, but it is < 3 per cent for $s > 1000$ au. The fall-off in F_{twin} with increasing separation is shallowest for $0.4 < M_1/M_\odot < 0.6$ and $0.6 < M_1/M_\odot < 0.8$, with a larger normalization for $0.4 < M_1/M_\odot < 0.6$. Here, F_{twin} is almost independent of separation over $350 < s/\mathrm{au} < 15\,000$. Finally, the fall-off steepens again in the lowest primary mass bin, where $F_{twin} \approx 0$ beyond 2500 au. We discuss possible interpretations of these trends in Section 5.2.1.
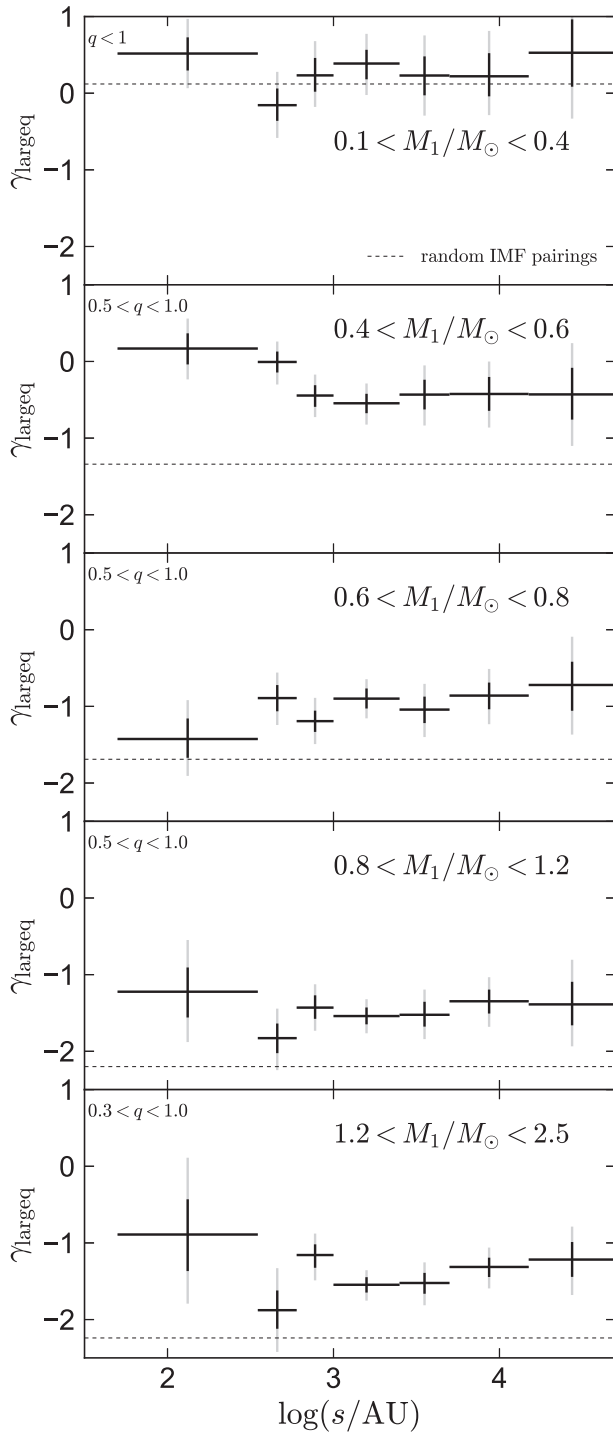
In the bottom two panels of Fig. 11, we compare our constraints on F_{twin} to the 1σ constraints obtained by Moe & Di Stefano (2017, their tables 8 and 11) for binaries in similar mass and separation ranges. The constraints for solar-type stars ($0.8 < M_1/M_\odot < 1.2$) were obtained from the solar neighbourhood sample of Raghavan et al. (2010). Those shown in the bottom panel were obtained from the AO-assisted survey of visual binaries with A-type primaries ($1.7 < M_1/M_\odot < 2.3$) described in De Rosa et al. (2014). Reassuringly, constraints for both mass bins are consistent with those obtained from our catalogue. Because the binary sample we analyse is larger than the Raghavan et al. (2010) and De Rosa et al. (2014) samples, we can tighten the uncertainties on F_{twin} at large separations, showing, for instance, that F_{twin} for solar-type primaries is inconsistent with 0 out to $s \approx 5000$ au.

Fig. 12 shows constraints on γ_{largeq} as a function of separation, again separating binaries by primary mass. The dashed line in each panel shows the slope that would be expected for random pairings from the IMF. This is obtained by fitting a power-law mass ratio distribution over the same range of primary masses and mass ratios to a simulated population of binaries in which the masses of both components are drawn from a Kroupa (2001) IMF and paired randomly.

Consistent with previous work (e.g. Lépine & Bongiorno 2007; Reggiani & Meyer 2011; Duchêne & Kraus 2013; Moe & Di Stefano 2017), we find that the mass ratio distribution is *not* consistent with random pairings from the IMF, but is weighted towards higher mass ratios than would be expected in such a scenario. Whether this is an imprint of the binary formation process or in part reflects the fact that binaries with higher mass ratio have higher binding energy and are thus more difficult to disrupt is an open question. Any trends in γ_{largeq} with separation are weak over the separation range we probe: at the 2σ level, our constraints are consistent with a separation-independent γ_{largeq} over $100 \lesssim s/\mathrm{au} < 50\,000$ in all mass bins. However, they are not consistent with being independent of primary mass: $p(q)$ becomes increasingly bottom-heavy (lower γ_{largeq}) with increasing M_1.

The fact that γ_{largeq} does not vary much with separation can serve as a strong constraint on formation models for wide binaries. It has frequently been argued that while binaries with separations of $100 \lesssim s/\mathrm{au} \lesssim 5000$ formed primarily by core fragmentation, those with $s \gtrsim 5000$ au (the size of typical cloud cores) formed by another process. Candidate processes include cluster dissolution (Kouwenhoven et al. 2010; Moeckel & Bate 2010), unfolding of hierarchical triples (Reipurth & Mikkola 2012), or pairing of adjacent cores (Tokovinin 2017a). One might naively expect a change in the mass ratio distribution at $s \sim 5000$ au if the binary formation mechanism changes there, but none is observed.

The mass ratio distribution for wide solar-type binaries is not uniform, but is weighted towards low-mass ratios. For example, companions with $q \approx 0.5$ are roughly twice as common as those with $q \approx 0.9$. The dominance of low-mass ratio companions can be seen clearly in the data at wide separations (Fig. 10) and cannot be due to selection effects, which all work *against* low-mass ratio binaries. $p(q)$ is thus more bottom-heavy at wide separations than

at close separations, where it is basically uniform (Mazeh et al. 1992; Tokovinin 2014). Analysing the 25-pc Raghavan et al. (2010) sample of solar-type binaries, Moe & Di Stefano (2017) found the mass ratio distribution to transition from $\gamma_{largeq} = -0.4 \pm 0.3$ (close to uniform) across $10 < s/\text{au} < 200$ to $\gamma_{largeq} = -1.1 \pm 0.3$ across $200 < s/\text{au} < 5000$.[6] Combined with our constraints at wide separations, this implies that the transition between a uniform mass ratio distribution at close separations and a bottom-heavy distribution at wide separations occurs relatively abruptly at $s \sim 100$ au. Several other binary population properties are observed to change at $s \sim 100$ au (see El-Badry & Rix 2019 and references therein), perhaps due to a transition in the dominant binary formation mechanism at this separation. This sharp transition, and the fact that γ_{largeq} is nearly constant over $300 \lesssim s/\text{au} \lesssim 50\,000$, provides a useful constraint for star formation models. Because the effects of dynamical processing after formation on the mass ratio distribution are imperfectly understood, similar constraints obtained in star-forming environments will prove useful for disentangling the primordial mass ratio distribution from the effects of dynamical processing.

We show constraints on other parameters obtained from our fitting in Fig. 13. The left column shows q_{twin}. This parameter is only meaningful if F_{twin} is non-zero, so we do not show constraints for bins of mass and separation where F_{twin} is consistent with 0 at the 2σ level. There are no strong trends in q_{twin} with mass or separation; the typical best-fitting value is $q_{twin} \approx 0.95$. In a few bins (e.g. $1.2 < M_1/M_\odot < 2.5$ and $350 < s/\text{au} < 600$), the constraint runs up against the prior, implying that a broader excess may be preferred. However, the width of the observed $\Delta G \approx 0$ excess is still reproduced reasonably well in these cases (Fig. 10).

Constraints on γ_{smallq} are shown in the middle column. Most constraints are consistent with $\gamma_{smallq} = 0$ (a flat distribution at small q), and trends with separation are weak. For the lowest mass bin ($M_1 < 0.4\,M_\odot$), we fit a single power law, so $\gamma_{smallq} = \gamma_{largeq}$. Unlike for γ_{largeq}, there are no strong trends with primary mass. We note that our data cannot strongly constrain γ_{smallq} for low-mass primaries, so the constraints are influenced somewhat by the priors (Section 3.1).

Finally, the right column of Fig. 13 shows constraints on the local slope of the separation distribution γ_s, which is defined such that $p(s) \propto s^{\gamma_s}$. The dashed line shows $\gamma_s = -1.6$, which is approximately the value that has been found for wide binaries when marginalizing over a larger range of primary masses and separations (Andrews, Chanamé & Agüeros 2017; ER18). The constraints we find here are similar to this value on average but show some evidence for a steepening in $p(s)$ with increasing separation. Any trends with primary mass at fixed separation are weak. We emphasize that these constraints come from the gradient in binary counts as a function of *s measured within a narrow bin of s*, which is necessarily noisy: trends in γ_s with separation represent the second derivative of $p(s)$.

5 DISCUSSION

5.1 Comparison to previous work

Although most work on twin binaries to date has focused on spectroscopic binaries, hints of a twin excess at wide separation

Figure 12. Marginalized constraints on γ_{largeq}, the logarithmic slope of the power-law component of the mass ratio distribution at large q, versus projected separation. Dark and light error bars show 1σ and 2σ constraints. The range of q over which γ_{largeq} is fit in each mass bin is noted in the upper left of each panel. Dashed lines show the slope expected if binary component masses were drawn from the IMF and paired randomly. Except in the lowest mass bin, $p(q)$ is more bottom-heavy than uniform, but it is always more top-heavy than expected for random IMF pairings.

---

[6]Because we use $q_{break} = 0.5$ for solar-type primaries and Moe & Di Stefano (2017) used $q_{break} = 0.3$, our measurements should not be directly compared. $p(q)$ flattens at $q < q_{break}$, so at fixed γ_{largeq}, a lower q_{break} corresponds to a more bottom-heavy mass ratio distribution.
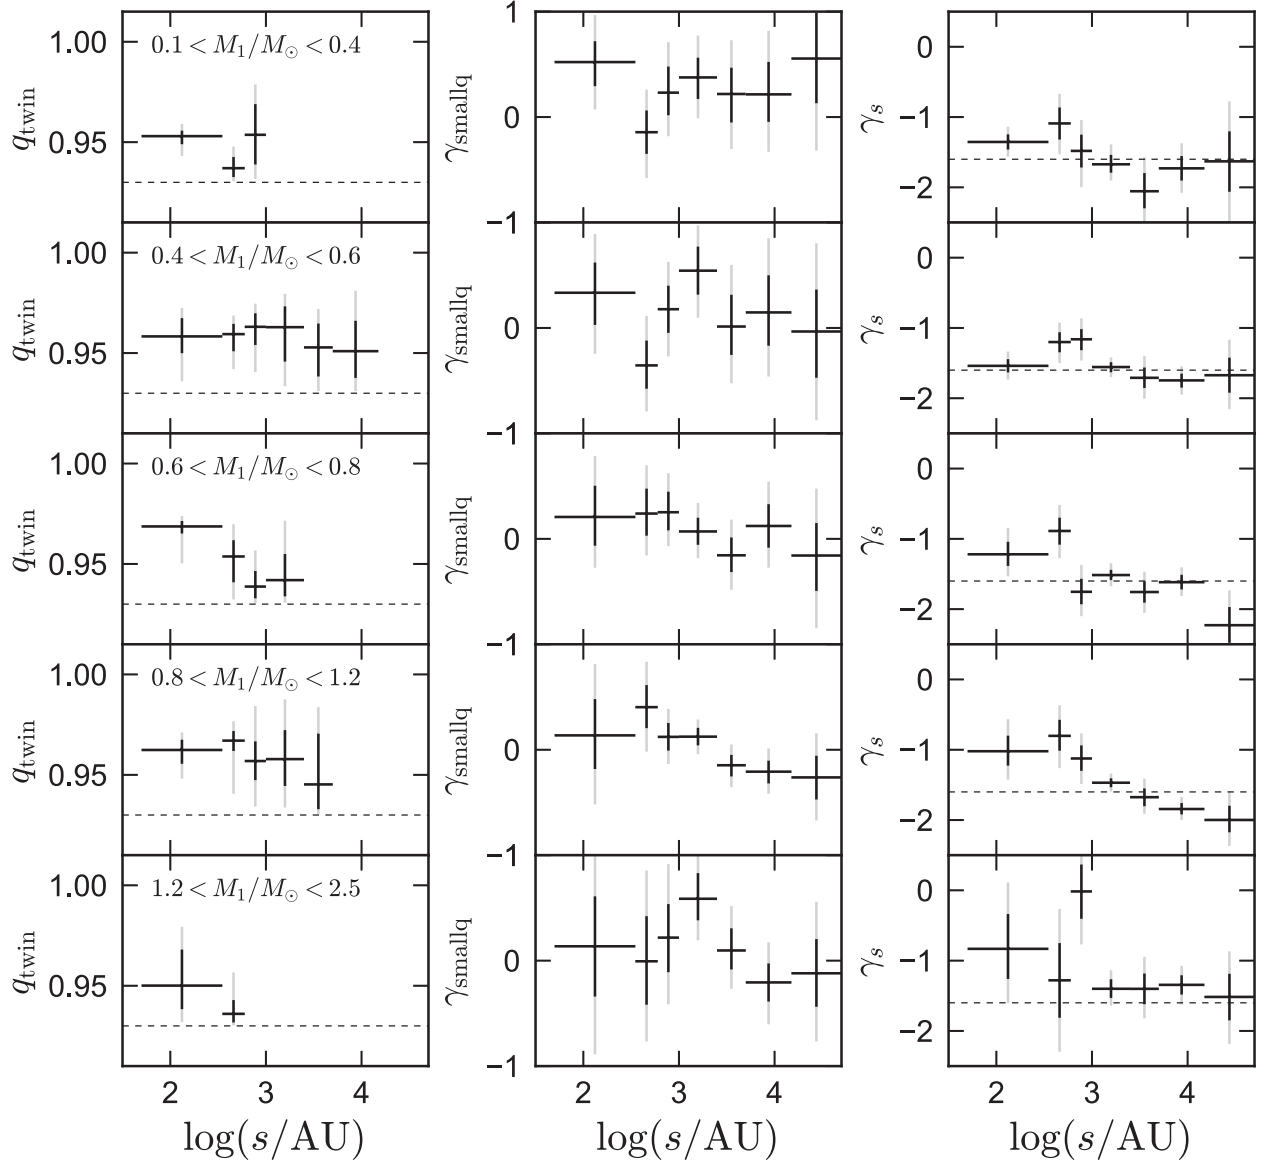
Figure 13. Constraints on q_{twin} (left; the mass ratio above which the twin excess begins), γ_{smallq} (middle; the logarithmic slope of the mass ratio distribution at small q), and γ_s (right; the logarithmic slope of the local separation distribution). Dark and light error bars show 1σ and 2σ constraints. Primary mass increases from top to bottom. Constraints on q_{twin} are only shown for mass and separation bins where F_{twin} is inconsistent with 0 at the 2σ level. Dashed line in the left-hand panels shows $q = 0.93$, the lower limit set on q_{twin} by the prior. Dashed lines in the right-hand panels show $\gamma_s = -1.6$, the value found by ER18 when averaging over all masses and separations.

have been reported in several previous works. Trimble (1987) found a significant excess of equal-brightness pairs among a sample of bright, nearby visual binaries. She suggested that this excess might point towards a formation mechanism that favours equal-mass systems but found significantly different distributions of magnitude difference when comparing different samples of wide binaries and thus did not rule out the possibility that the $q \sim 1$ peak was the result of selection effects. Similar conclusions were reached by Giannuzzi (1987). Halbwachs (1988) argued that the mass ratio distribution of wide, common proper motion binaries was likely consistent with random pairings from the IMF once selection effects were corrected for.

Larger and more homogeneous samples of bright wide binaries were identified using astrometry from the *Hipparcos* satellite for one or both components (e.g. Söderhjelm 2000, 2007; Eggenberger et al. 2004; Lépine & Bongiorno 2007; Shaya & Olling 2011). The mass ratio distribution of *Hipparcos* binaries with A and F star primaries was modelled in detail by Söderhjelm (2007), who found evidence for a $q \approx 1$ peak at $100 \lesssim s/\mathrm{au} \lesssim 1000$. The strength of the peak decreased with primary mass. He argued that the twin excess was not the result of any known selection effect, although he did not reject the possibility that an unknown bias in the *Hipparcos* input catalogue could explain it. The twin feature identified by Söderhjelm (2007) is likely the same feature apparent in our catalogue. We note, however, that most of the binaries observed by *Hipparcos* fall in our highest primary mass bin (where the twin excess is weaker than at lower masses), as *Hipparcos* only observed bright stars ($G \lesssim 13$).
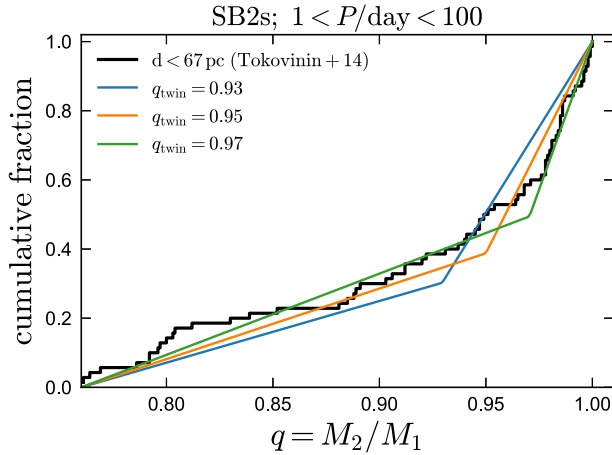
Figure 14. Cumulative distribution function of short-period double-lined spectroscopic binaries with $q > 0.76$ in the volume-limited 67-pc sample of FG dwarfs from Tokovinin (2014). We compare predictions for models with three values of q_{twin} (see Section 5.1.1). For this sample, we find $q_{twin} = 0.964 \pm 0.013$, consistent with the values we find for wide binaries with periods $P = 10^{5-9}$ d.

A preference for equal-brightness pairs was also noticed among wide common proper motion disc and halo binaries identified by Chanamé & Gould (2004) and Dhital et al. (2010). The authors could not rule out the possibility that it was the result of selection effects, which were not well understood for their samples. The twin excess in our catalogue is visually quite striking (e.g. Fig. 4), so a natural question is why it was not as clear in earlier binary catalogues. In Appendix C, we show that a twin excess *is* apparent in the large, low-mass wide binary catalogue produced by Dhital et al. (2015) using SDSS photometry, but it only becomes obvious once objects with blended photometry are removed. It is also clearer in the *Gaia* photometry than in the ground-based SDSS photometry.

More broadly, the twin excess we identify is visually striking because it is narrow, but it represents only a few per cent of the total twin population. This means that it will only become obvious when (a) the photometry is sufficiently precise and uncontaminated that the difference in the components' apparent magnitude can be measured with precision that is good compared to the intrinsic width of the twin excess and (b) the sample considered is large, containing (at least) hundreds of objects. A twin excess among solar-type visual binaries is also observed in the Raghavan et al. (2010) 25-pc and Tokovinin (2014) 67-pc samples, extending out to $s \approx 200$ au at a statistically significant level Moe & Di Stefano (2017). This excess is consistent with our constraints (e.g. Fig. 11); at wider separations, these samples did not contain enough objects to detect or rule out a few-per-cent excess with high significance.

5.1.1 Width of the twin excess

Among spectroscopic binaries, there has been some disagreement in the literature over whether the twin excess is limited to a narrow peak in the mass ratio distribution at $q \gtrsim 0.95$ (e.g. Tokovinin 2000) or is a broader feature, corresponding simply to a positive slope in $p(q)$ at $q \gtrsim 0.8$ (e.g. Halbwachs et al. 2003; see Lucy 2006 for further discussion). In our sample, the twin feature is unambiguously narrow, only becoming significant above $q_{twin} \approx 0.95$ (Fig. 13). In Appendix A1, we show that $p(q)$ is always consistent with a flat or negative power law at $q < 0.94$; for the majority of the mass and

separation bins, the twin excess only becomes strong at $q > 0.97$ (see Fig. A1). To allow better comparison between close and wide binaries, we now re-examine the twin excess among spectroscopic binaries.

In the volume-limited 67-pc sample of FG dwarfs, Tokovinin (2014) identified 98 double-lined spectroscopic binaries (SB2s) with periods $P = 1-100$ d, 70 of which have dynamical mass ratios $0.76 < q < 1$. Nearly all binaries with $P < 100$ d and $q > 0.76$ will appear as SB2s, so the 70 observed systems in this parameter space represent a relatively complete subsample. In Fig. 14, we plot the cumulative mass ratio distribution of the 70 short-period SB2s with $q > 0.76$, about half of which have $q > 0.95$. We model a uniform mass ratio distribution across $0.76 < q < 1$ with an excess twin fraction F_{twin} above $q > q_{twin}$. We use the maximum-likelihood method described in Moe & Di Stefano (2017) to fit the two free parameters F_{twin} and q_{twin} and draw 1000 bootstrap samples to estimate their uncertainties. We show in Fig. 14 the best-fitting models obtained when fixing $q_{twin} = 0.93, 0.95,$ and 0.97. The value $q_{twin} = 0.93$ is inconsistent with the data ($p = 0.014$), while $q_{twin} = 0.95$ and 0.97 both provide reasonable fits. We formally measure $q_{twin} = 0.964 \pm 0.013$ (1σ uncertainties).[7] By using a larger and more complete sample of short-period SB2s, we thus confirm the conclusions of Tokovinin (2000) and Moe & Di Stefano (2017) that close solar-type binaries with $a < 0.5$ au exhibit a large excess twin fraction and that the twins are narrowly distributed above $q_{twin} \gtrsim 0.95$.

In Fig. 15, we show constraints on F_{twin} and q_{twin} for solar-type binaries across a wide range of periods and separations. At wide separations, the constraints from *Gaia* wide binaries are reproduced from Figs 11 and 13. At closer separations, we show the constraints on F_{twin} obtained by Moe & Di Stefano (2017) from the Raghavan et al. (2010) sample, as well as the constraint on q_{twin} at $P = 1-100$ d derived above for the Tokovinin (2014) sample (Fig. 14). At intermediate periods ($P = 10^{2-6}$ d), we show $q_{twin} \approx 0.95 \pm 0.02$ based on the Moe & Di Stefano (2017) analysis of the Raghavan et al. (2010) sample. This constraint is not the result of formal fitting but provided a good fit to the data (see fig. 30 of Moe & Di Stefano 2017). Fig. 15 shows that while F_{twin} decreases with separation, $q_{twin} \approx 0.95$ is consistent with being constant over all separations in this mass range. Similar values of q_{twin} are also found for massive binaries at the short periods where there is a significant twin excess (Moe & Di Stefano 2013).

5.2 Origin of twin binaries

It is typically assumed that the components of binaries wider than a few hundred au formed nearly independently of one another (e.g. White & Ghez 2001; Moe & Di Stefano 2017; Tokovinin 2017a; Moe, Kratter & Badenes 2019; El-Badry & Rix 2019) during turbulent core fragmentation (for binaries with separation less than a few thousand au; e.g. Offner et al. 2010) or by becoming bound at slightly later times (for those with the widest separations; e.g. Parker et al. 2009; Kouwenhoven et al. 2010; Moeckel & Bate 2010; Tokovinin 2017b).

The existence of a narrow twin excess at $q \gtrsim 0.95$ suggests that the components of a fraction of binaries with $s \gg 100$ au formed at closer separations in a highly correlated way. We do

---

[7]We do not present constraints on F_{twin} for this sample, because F_{twin} depends on $p(q)$ at $0.3 < q < 1$, and many lower mass ratio binaries will not be double-lined.
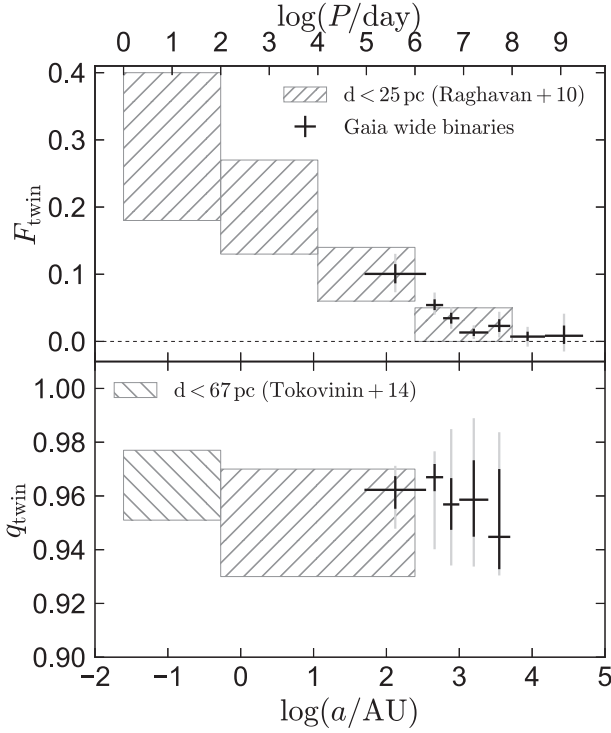
Figure 15. F_{twin} (top) and q_{twin} (bottom) for solar-type binaries ($0.8 \lesssim M_1/M_\odot \lesssim 1.2$). We compare results from *Gaia* wide binaries (dark and light error bars show 1σ and 2σ constraints) to 1σ constraints at closer separations from the samples of Raghavan et al. (2010) and Tokovinin (2014, see Fig. 14). F_{twin} decreases with separation and is consistent with 0 at $a \gtrsim 5000$ au. However, $q_{twin} \approx 0.95$ is basically constant at all separations.

expect that the dynamical process of becoming and staying bound may lead to a preference for roughly equal-mass binaries (say, $q > 0.5$), because these have higher binding energy. This could quite reasonably explain, at least in part, why the power-law component of the mass ratio distribution is shallower than expected for random pairings from the IMF (Fig. 12). But we do not expect dynamical processes to produce a *sharp* twin feature like what is observed: at fixed primary mass, the binding energy at $q = 0.9$ is not much less than that at $q = 1$. And indeed, simulations of binary formation during cluster dissolution find larger typical mass ratios at wide separations than predicted for random pairings, but they do not predict a narrow excess of twins (Kroupa 1998; Kouwenhoven et al. 2010; Moeckel & Bate 2010).

The excess twin fraction uniformly decreases with separation and eventually goes to 0 at $s > 15\,000$ au in all mass bins. The shape of the twin excess (i.e. q_{twin} and the slope of $p(q)$ at $q > q_{twin}$) does not vary much between 50 and 15 000 au in our catalogue. Moreover, it is effectively the same for spectroscopic binaries (with separations as close as 0.01 au) and wide binaries (Fig. 15). Invoking Occam's razor, it seems more likely that the wide binary twin phenomenon is an extension of the phenomenon that has previously been observed at $s \lesssim 100$ au than that it is produced by a qualitatively different process.

Even for close binaries, there is not a clear consensus in the literature about the physical origin of the twin phenomenon. Some models for the formation of twins can only apply to very close binaries. In the first paper to highlight the twin phenomenon for spectroscopic binaries, Lucy & Ricco (1979) suggested that twins were formed by fragmentation of rapidly rotating pre-MS stars

during the late stages of dynamical collapse, at scales of $a \ll 1$ au. Alternatively, Krumholz & Thompson (2007) proposed that twins could be produced by mass transfer between stars of initially different masses during pre-MS evolution.[8] It seems implausible that such mechanisms can explain the twin phenomenon among wide binaries, because there is no known mechanism to widen the orbits of twins from the separations at which they operate – a few, or at most a few tens of, solar radii – to the separations at which they are observed today. Such widening would require a very strong velocity kick, the magnitude of which would have to be fine-tuned in order to not unbind the binaries completely.

A more plausible formation mechanism for equal-mass twins at wider separations is through competitive accretion from a circumbinary disc. Many studies have found that the accretion rate from a circumbinary disc is usually higher for the secondary than the primary (e.g. Bate & Bonnell 1997; Bate 2000; Farris et al. 2014; Young & Clarke 2015; Nelson & Marzari 2016; Matsumoto, Saigo & Takakuwa 2019). Because the secondary's orbit is larger than that of the primary, it sweeps out a larger radius in the disc and can accrete more rapidly than the primary, unless the material being accreted has low angular momentum.[9] Preferential accretion on to the secondary will necessarily drive the mass ratio towards unity. If such accretion continues for long enough, binaries within circumbinary discs should thus end up with $q \approx 1$. An appeal of this formation mechanism is that it can operate at scales comparable to the size of circumbinary discs, $s \lesssim 100$ au.

It is plausible but not obvious that preferential accretion on to the secondary will give rise to a sharp twin feature like the one found observationally. In order to end up at $q \approx 1$, a binary that initially had an intermediate-mass ratio must accrete a large fraction of its mass from a circumbinary disc, such that there is enough time to drive the mass ratio to 1 even while the primary continues to accrete. If twins are formed by accretion from circumbinary discs, then the width of the twin feature (i.e. q_{twin}) can tell us about the fraction of the total mass accreted from the disc, as well as the mass ratio above which the accretion rate on to the two components becomes nearly equal. A sharp twin feature could be expected if a fraction of binaries accrete most of their mass from a circumbinary disc (becoming twins) and the rest either do not develop circumbinary discs or only accrete a subdominant fraction of their total mass from them.

In accretion-driven explanations of the twin excess, the increase in F_{twin} towards close separations has been interpreted as evidence that gravitational torques within circumbinary discs (e.g. Artymowicz et al. 1991; Shi et al. 2012) shrink the orbits of

---

[8]These authors sought to explain the observed twin excess in massive stars. The specific mechanism they proposed, which relies on deuterium shell burning causing protostars to expand and overflow their Roche lobes, cannot operate in solar-type or lower mass stars. Moreover, the twin excess for massive stars appears to be limited to close separations (Moe & Di Stefano 2017), so mass transfer may adequately explain it. Here, we simply suppose, for the sake of argument, that there is some mechanism through which stable mass transfer in lower mass stars could drive the mass ratio to unity.

[9]It is worth noting that there has not been a full consensus in the literature whether this mechanism works: some simulations of accretion from a circumbinary disc have actually predicted the opposite trend, with accretion favouring the primary (Ochi, Sugimoto & Hanawa 2005; Hanawa, Ochi & Ando 2010; de Val-Borro et al. 2011). These simulations assumed a higher gas temperature than those which have found accretion to favour the secondary; Young & Clarke (2015) showed that accretion only favours the secondary when accreted gas is cold, as is appropriate for stellar binaries.

twins (Young & Clarke 2015). Indeed, the observed twin excess is largest at separations $a < 0.2$ au (Fig. 15), too close for binaries to have formed at their current separations. This implies that some combination of gravitational torques, viscous dissipation, and dynamical interactions (e.g. Bate 2012; Moe et al. 2019) must have shrunk the orbits of twins at very close separations.

However, several recent simulations of circumbinary discs have found that, contrary to previous results in the literature, accretion can also *widen* binaries within circumbinary discs, when the advective torque dominates over the gravitational torque (Miranda, Muñoz & Lai 2017; Moody, Shi & Stone 2019; Muñoz, Miranda & Lai 2019). Whether gravitational or advective torques dominate depends on details such as the sink prescription used for accretion (Tang, MacFadyen & Haiman 2017). Further work is needed to clarify the effects of circumbinary discs on orbital evolution. However, the fact that the twin excess extends to very wide separations suggests that orbit shrinkage is not a ubiquitous outcome of accretion from circumbinary discs.

High-resolution studies of the dynamics of individual circumbinary discs are generally too idealized, run for too short a time, and are focused on too narrow a range of initial conditions to make ab initio predictions of the full mass ratio distribution. However, they do typically find that gas is preferentially accreted on to the secondary for realistic accretion geometries once a steady state is reached. On the other hand, global simulations of the fragmentation of molecular clouds (e.g. Bate 2009, 2014, 2019) are reaching the point where they can make realistic predictions of the mass ratio distribution with minimal fine-tuning. Such studies have lower resolution than idealized simulations of individual binaries, so it is not guaranteed that the dynamics within discs are well resolved, but they are able to predict the accretion rate and angular momentum distribution of accreted material, and the mass and size distribution of discs (e.g. Bate 2018). Such global calculations predict an excess of equal-mass binaries out to separations of order 100 au (see e.g. Bate 2014, fig. 18). Because they typically only form a few dozen binaries in a cloud, such calculations do not yet have the statistical power to predict or rule out a few-per cent effect at wider separations.

Accretion can plausibly explain an excess of twins out to significantly larger separations than mass transfer or late-stage fragmentation, but it alone cannot explain a signal reaching out to thousands of au. Observed circumstellar and circumbinary gas discs have typical radii of order 100 au (Ansdell et al. 2018; Eisner et al. 2018). The *largest* observed circumbinary discs have radii of order 500 au and host binaries with separations of 50–200 au (e.g. Hioki et al. 2007; Brinch et al. 2016; Tobin et al. 2016; Takakuwa et al. 2017; Comerón et al. 2018; Czekala et al. 2019); these preferentially host relatively massive binaries. There are no observed discs with radii exceeding 1000 au, and simulations also predict the largest circumbinary discs to have radii of several hundred au (Bate 2018). It thus seems exceedingly unlikely that twin binaries with $s \gtrsim 200$ au formed at their present-day separation by accretion from a circumbinary disc. This implies that either (a) twin binaries formed at closer separations and their orbits were subsequently widened or (b) some other process is responsible for producing twins at very wide separations. Lacking a good candidate mechanism for (b), we here consider the plausibility of orbit widening.

5.2.1 Dynamical orbit widening in young clusters

In the Galactic field, dynamical interactions have a negligible effect on most binaries with $s \lesssim 10\,000$ au (e.g. Weinberg, Shapiro &

Wasserman 1987). However, dynamical interactions are expected to be more efficient in binaries' birth environments, where the typical stellar density is higher. The dynamical evolution of binaries within their birth clusters has been the subject of considerable study. As a general rule, interactions within birth clusters are expected to widen the orbits of binaries with orbital velocities lower than the cluster velocity dispersion, and to tighten the orbits of those with orbital velocities greater than it (Heggie 1975; Hills 1975). There are several complicating factors in real clusters. For example, the mass distribution within clusters at early times is not smooth but clumpy (Dorval et al. 2017), both stars and gas can be dynamically important, and cluster density and velocity dispersion change as clusters age, in part because energy is redistributed among binaries (see e.g. Kroupa & Burkert 2001; Parker et al. 2009; Goodwin 2010). Observed binary populations provide constraints on models for binaries' dynamical evolution, but a unified model to explain the diversity of binary populations found in different environments does not exist. We summarize some relevant observational constraints on disruption below.

(i) Over a wide range of separations ($10 \lesssim s/\mathrm{au} < 3000$), the binary fraction in low-density star-forming regions (e.g. Taurus) is higher than in the field by roughly a factor of 2 (e.g. Leinert et al. 1993; Duchêne 1999). This is true especially at wider separations, where the separation distribution of young binaries is roughly log-uniform, but that of field stars declines more steeply (e.g. Connelley et al. 2008; Kraus et al. 2011).

(ii) In dense young clusters (e.g. the ONC), the binary fraction at separations of $s \gtrsim 100$ au is similar to the field (lower than in low-density star-forming regions) and declines steeply at wide separations (Reipurth et al. 2007). At closer separations ($10 \lesssim s/\mathrm{au} \lesssim 60$), the binary fraction in dense clusters is comparable to that in low-density star-forming regions, and higher than that found in the field (Duchêne et al. 2018).

(iii) At very close separations ($s \lesssim 5$ au), the binary fraction in star-forming regions (over a range of densities) is consistent with that in the field (e.g. Kounkel et al. 2019).

Some models (e.g. Kroupa 1995; Marks et al. 2011; Marks & Kroupa 2012) postulate that the initial binary fraction and separation distribution are insensitive to local properties, such that observed variation in binary populations must be due to disruption. These models interpret the higher wide binary fraction in low-density clusters as the primordial population, which is transformed into the field population by dynamical widening and disruption. Because disruption of wide systems is more rapid in dense clusters, these models also predict the binary fraction at $s \gtrsim 100$ au to decrease with cluster density, in agreement with observations. However, if such models are correct, it is not clear what happens to the excess of relatively tight binaries ($10 \lesssim s/\mathrm{au} \lesssim 100$), which are overrepresented relative to the field in both high- and low-density star-forming regions. Such systems have high enough binding energies that they can only be disrupted in very dense clusters. Thus, the lower binary fraction in the field at $10 \lesssim s/\mathrm{au} \lesssim 100$ would imply that a large fraction of field stars formed in dense environments.

In any case, the fact that the binary fraction in the field is lower than that in star-forming regions down to fairly close separations implies that a significant fraction of young binaries undergo quite energetic interactions, the cumulative effect of which is sufficient to disrupt binaries with initial separations as close as 10–100 au. In many cases, binaries will be disrupted by such interactions, but in some cases, they will only be widened (e.g. Kroupa &

Burkert [2001]). It is in this latter context that the twin fraction at wide separations is informative about the fraction of stars at a given present-day separation that formed at significantly closer separations.

The fact that the dependence of F_{twin} on separation varies with binary mass in a non-monotonic way (Fig. 11) suggests that, if the twin excess at $s \gg 100$ au is due to dynamical widening of orbits, then the primordial twin statistics (e.g. the twin fraction and the range of separations over which twins are form) must also vary with mass. This may not be unreasonable, since the physical properties of discs do vary with mass (Bate [2018]; Eisner et al. [2018]), but it means that disentangling the effects of dynamical widening and the primordial separation distribution of twins is non-trivial.

We use a simple toy model to explore how the separation-dependence of the twin excess may depend on the initial twin population and the strength of dynamical interactions. We suppose that twin formation is efficient out to a separation of $a_{twin, max} \approx 100$ au and does not operate at wider initial separations. We consider primordial populations of equal-mass binaries with components of mass M and a uniform distribution of $\log a$ between 10 and 10^4 au. At $a < a_{twin, max}$, we assign a random subset of 20 per cent of the binaries the label of 'twin'. We then assume that dynamical interactions add an energy E_{int} to each orbit (for both twins and non-twins), such that the new orbital energy is $E_{new} = E_{initial} + E_{int}$, where $E_{initial} = -GM^2/a_{initial}$. Values of E_{int} for each binary are drawn from a uniform distribution between 0 and $E_{max, interactions}$. If E_{new} is positive, then the binary is considered unbound and is discarded. The new separation after interactions is $a_{new} = -GM^2/E_{new}$. Finally, we measure what fraction of the surviving binaries bear the twin label as a function of final separation. This fraction is proportional to the fraction of binaries at a given present-day separation that formed with $a < a_{twin, max}$. We do not model dynamical hardening of close binaries, and thus implicitly assume that all binaries we consider are soft; i.e. that they reside in clusters with velocity dispersion larger than the highest orbital velocities of binaries being widened. For an initial separation of 100 au and binaries with typical component mass of $0.5 \, M_{\odot}$, this corresponds to a dispersion of $\sigma \approx 1.5 \, \mathrm{km \, s^{-1}}$; for an initial separation of 10 au, to $\sigma \approx 5 \, \mathrm{km \, s^{-1}}$.

We plot the results of this experiment in Fig. 16, varying $a_{twin, max}$ (top), $E_{max, interactions}$ (middle), and M (bottom). The final separation-dependence of the twin fraction depends significantly on all of these parameters. The twin excess extends to larger separations, and falls off less steeply with separation, when (a) the initial separation out to which twins form is larger, (b) dynamical interactions are more energetic, or (c) the binding energy of twins is lower. In order to obtain a non-zero twin fraction at very wide separations, it is necessary that a fraction of binaries undergo dynamical interactions energetic enough to unbind binaries at separation at which twins form. Once this is satisfied, twins can contribute significantly to the binary population at wider separations, because the same interactions that widen close binaries will unbind a large fraction of initially wider non-twin binaries. Wide twins are produced most efficiently in clusters where the velocity dispersion is comparable to the orbital velocity at a separation of $a_{twin, max}$: significantly denser clusters produce few wide binaries, since most binaries that are not close are disrupted completely.

The toy model generically predicts that F_{twin} decreases with separation and that low-mass twins can be more efficiently widened than high-mass twins due to their lower binding energy. It thus predicts that the excess twin fraction will fall off less steeply for lower mass binaries. This trend *is* found in the observed binaries
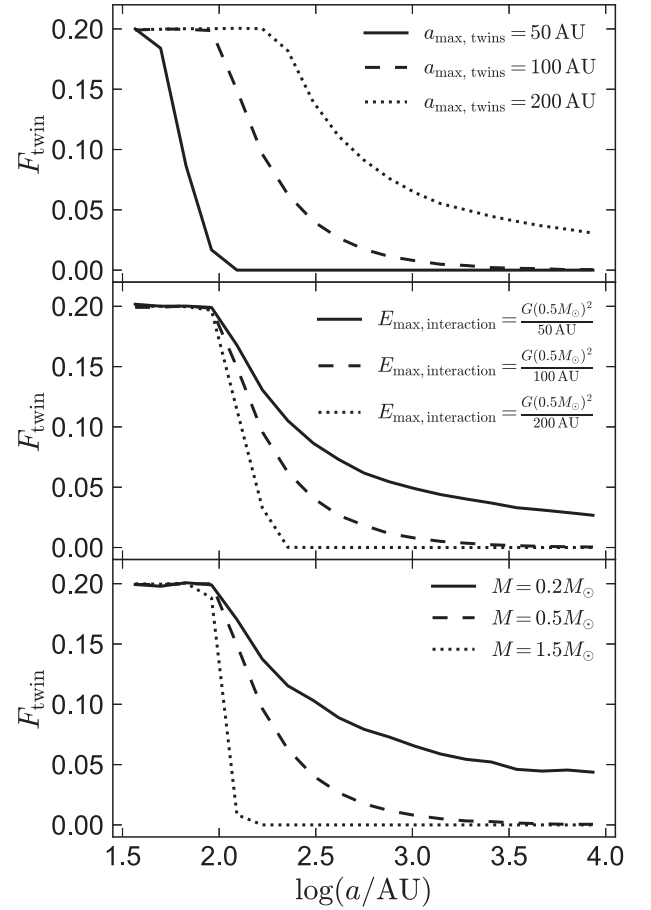

Figure 16. Predictions of the toy model for orbit widening described in Section 5.2.1. We assume that the initial twin fraction is 20 per cent for $a < a_{twin, max}$ (varied in the top panel; 100 au by default) and 0 at larger separations. We consider an initial separation distribution that is flat in log space. We then widen the orbits of all binaries by adding energy, representing the cumulative effect of gravitational interactions within a birth cluster. The total added energy is drawn from a uniform distribution between 0 and $E_{max, interaction}$. Binaries that are disrupted are discarded. We plot the final twin fraction as a function of separation, varying $E_{max, interaction}$ (middle panel; default is $G(0.5 M_{\odot})^2 / (100 \, au)$) and mass (bottom panel; default is $0.5 \, M_{\odot}$).

over $0.4 < M_1/M_{\odot} < 2.5$ (Fig. 11), but it is reversed for the lowest mass subsample: F_{twin} falls off significantly more steeply for $0.1 < M_1/M_{\odot} < 0.4$ than for $0.4 < M_1/M_{\odot} < 0.6$. This is not easily explained in the context of the toy model.

It is thought that a fraction of wide binaries (likely at $a \gg 1000$ au) form at later times (perhaps during cluster dissolution) than wide binaries at closer separations. Dilution of the twin excess due to these binaries is not accounted for in the toy model. If the fraction of binaries at fixed separation that formed during cluster dissolution were higher at lower masses, this could explain the observed steeper decline in F_{twin} with separation at low masses.

An addition complication is that the observed population of field binaries is an average over a wide range of formation environments, from low-density regions to dense clusters. The trends in Fig. 16 rely on a number of crude approximations and should not be used to directly interpret the observed trends in F_{twin} (and we have not attempted to tune the model to match observed trends). Here, we simply emphasize that the separation distribution of twins is quite
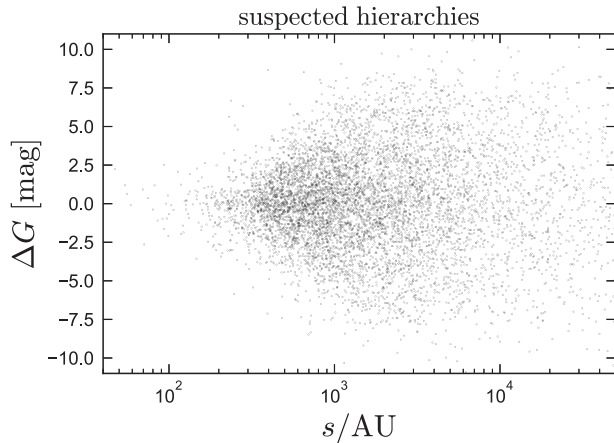
Figure 17. Magnitude difference between the two resolved components of all wide binaries in which one or both components is suspected to have a bright, unresolved companion ($N = 7,295$; yellow points in Fig. 1). There is no obvious excess at $\Delta G \approx 0$ (compare to the right-hand panel of Fig. 3).

sensitive to dynamical processing. We conclude that dynamical orbit widening provides a plausible explanation for the existence of the wide twin excess, but more realistic theoretical modelling is needed to determine whether the observed trends in F_{twin} with primary mass and separation can be reproduced when averaging over a realistic population of star-forming environments.

5.2.2 *Widening of unstable triple systems?*

Another possible mechanism for dynamically widening twins is through unfolding of unstable triple systems, in which initially close companions can be scattered to much wider orbits (e.g. Reipurth & Mikkola 2012). Triples are not uncommon: more than a third of all wide binaries contain subsystems (e.g. Tokovinin & Smekhov 2002; Tokovinin 2014), and a large fraction of binaries are thought to have formed as higher order multiples that subsequently decayed (e.g. Sterzik & Durisen 1998).

Widening of the outer orbit in triples generally comes at the expense of shrinking of the inner orbit. Since the outer orbit is the more fragile one, a triple-origin of the wide twin excess would imply that one of the component stars in wide twins had, and in most cases still has, a close companion, which in most cases would be spatially unresolved. However, the twin excess at $\Delta G \lesssim 0.25$ mag cannot be due to systems in which either resolved component has an unresolved companion any less than ≈ 2.5 magnitudes fainter than it; otherwise the light from the unresolved companion would make it more than 0.25 mag brighter than the other resolved component. Fig. 17 shows that there is indeed no obvious excess at $\Delta G \approx 0$ for wide binaries in which either component has a bright unresolved companion. There is thus a photometric upper limit on the mass ratio of any unresolved companions to twins with $\Delta G \approx 0$; it ranges from $q < 0.45$ to $q < 0.7$, depending on primary mass (see Fig. 8).

Unresolved companions with lower mass ratios cannot be ruled out based on photometry, but at close separations, they can be detected from RV variability. In Appendix F, we show that the *Gaia* visit-to-visit RV scatter implies a *lower* close binary fraction for the components of wide twins than for components of non-twins with similar separations and masses. This speaks against a triple-origin

of wide twins. It is, however, consistent with our expectations if twins are formed through accretion in a circumbinary disc, since the presence of a third close companion in the disc would complicate the mechanism through which accretion drives the mass ratio of two stars in a disc towards unity.

6 SUMMARY AND CONCLUSIONS

We have analysed a pure and homogeneous sample of $\sim$42 000 MS wide binaries selected from *Gaia* DR2 to constrain the mass ratio distribution of binaries with projected separations $50 \lesssim s/\mathrm{au} < 50\,000$ and primary masses $0.1 < M_1/\mathrm{M}_\odot < 2.5$. High-precision photometry allows us to measure mass ratios with unprecedented accuracy, and a well-understood selection function makes it possible to account for biases arising from the magnitude and contrast sensitivity limits of *Gaia* DR2 given our quality cuts.

A striking feature of the catalogue is a high-significance excess of 'twin' binaries with nearly equal brightness (Fig. 3). The excess is present over a narrow range of magnitude differences, $0 < \Delta G \lesssim 0.25$, corresponding to mass ratios $0.95 \lesssim q \lesssim 1$, and extends over a wide range of masses and separations. The twin excess is reminiscent of the excess of equal-mass binaries historically reported at very close separations ($a < 0.2$ au) and recently found to extend to $a \sim 100$ au, but it extends to separations of several thousand au, where binary formation models do not predict strongly correlated component masses.

We have done a variety of tests to confirm that the twin feature is caused by an astrophysical excess of equal-mass binaries, not selection effects or data artefacts. We first repeated the search that produced the binary catalogue, but required that the two components of a pair have *inconsistent* rather than consistent astrometry. This yields a catalogue of physically unassociated 'chance alignments' that are subject to essentially the same selection function as the real binaries and have similar distributions of most observable properties. A narrow excess of equal-brightness pairs is not found among chance alignments (Fig. 4).

By considering binaries at different distances, we verified that the strength of the twin feature varies primarily with physical, not angular, separation (Fig. 5). This speaks to the physicality of the twin feature, since most observational biases are expected to depend on angular separation, which is the observable quantity. Finally, we verified that the twin feature is not an artefact of the *Gaia* photometry; it is visible in photometry from other surveys. We also identified the same twin excess in archival data from another binary catalogue after removing objects with contaminated photometry (Appendix C).

In order to measure the intrinsic mass ratio distribution, we forward-model the joint distribution of magnitude difference, angular separation, and parallax given a Galactic model, a selection function, and a parametrized mass ratio distribution. We measure the selection function for our catalogue empirically, using chance alignments to quantify the contrast sensitivity as a function of angular separation. Our primary results are as follows:

(i) *Twin fraction*: We quantify the twin excess as F_{twin}, the fractional excess of binaries with $q > q_{twin}$ relative to the full population with $q > 0.3$ (Fig. 7; $q_{twin} \approx 0.95$ quantifies the width of the excess). Typical values of F_{twin} are 10 per cent at 100 au and 3 per cent at 1000 au, with some dependence on mass (Fig. 11). These values are lower than the value found for spectroscopic

binaries (Fig. 15), which is $F_{twin} \approx 25$ per cent, but they are clearly inconsistent with 0.

(ii) *Width of the twin excess*: At all masses and separations where it is statistically significant, the twin excess we find is quite narrow, only becoming significant at $q \gtrsim 0.95$ (Fig. 13). We experimented with different functional forms for the enhancement near $q \approx 1$, but we find a step function increase at $q \gtrsim 0.95$ to perform as well as more complicated models (Fig. A1). We also re-measured q_{twin} at close separations using a volume-limited sample of spectroscopic binaries (Fig. 14). The width of the twin excess, q_{twin}, is basically constant over six orders of magnitude in separation, from the closest spectroscopic binaries to wide binaries at $s \sim 10\,000$ au.

(iii) *Mass and separation dependence*: The fractional excess of twins declines with separation and is consistent with 0 at $s > 15\,000$ au for all mass bins (Fig. 11). The twin excess at $s > 1000$ au is strongest for primary masses $0.4 < M_1/M_\odot < 0.6$. At closer separations ($50 < s/$au < 1000), it is strongest for low-mass primaries ($M_1 < 0.4\,M_\odot$) and roughly consistent for other mass bins. The slope of the fall-off in F_{twin} with separation varies non-monotonically with mass; it is shallowest at $0.4 < M_1/M_\odot < 0.6$.

(iv) *Mass ratio distribution at lower q:* We provide broken power-law fits to the full mass ratio distribution for all mass and separation bins (Fig. 9; Appendix G). These fits reproduce the observed data well (Fig. 10). For solar-type stars, the power-law slope is $\gamma_{largeq} \approx -1.3$ for $q > 0.5$ (i.e. weighted towards lower mass ratios than the uniform distribution found at closer separations) and $\gamma_{smallq} \approx 0$ for $q < 0.5$. $p(q)$ becomes more bottom-heavy with increasing primary mass but is always flatter than expected for random pairings from the IMF. Besides variation in the excess twin fraction, $p(q)$ does not vary much with separation over $100 \lesssim s/$AU $< 50\,000$ (Figs 12 and 13).

(v) *Origin of the twin excess*: No theoretical models that have been proposed to explain twin binaries at close separations predict them to form at $s \gtrsim 100$ au. Dynamical processes may lead to a formation bias against low-q binaries, but they are not expected to produce a sharp peak at $q \approx 1$. Given the monotonic fall-off in F_{twin} with separation, the similar width of the twin feature between close and wide binaries, and the lack of a plausible mechanism for forming twins at very wide separations, we conjecture that the excess twins must have formed at closer separations ($s \lesssim 100$ au; likely through accretion from a circumbinary disc) and subsequently been widened by dynamical interactions.

In this scenario, the separation-dependence of the twin fraction is an imprint of dynamical orbit widening in binaries' birth environments (see Section 5). The plausibility of this explanation is not straightforward to assess because (a) present-day field binaries formed in a wide range of environments and (b) existing models for widening and disruption of binaries in star-forming regions do not fully explain the diversity of observed binary populations in young clusters and in the field. A simple toy model suggests that a separation-dependence in F_{twin} comparable to that which is observed is predicted if typical dynamical interactions are strong enough to disrupt binaries at the separation inside which twins are expected to form ($a \lesssim 100$ au). However, the mass-dependence of F_{twin} at wide separations is not fully explained in such models. Further theoretical work is required to (a) predict the primordial separation distribution of twins at different masses and (b) constrain the efficiency of dynamical orbit widening for a realistic ensemble of star-forming environments. In future work, we will search for a wide twin excess in observed star-forming regions to shed more light on the primordial separation distribution of twins and the density-dependence of orbit widening.

REFERENCES

Allard F., 2014in Booth M., Matthews B. C., Graham J. R., eds, IAU Symp. Vol. 299, Exploring the Formation and Evolution of Planetary Systems. p. 271

Allard F., Homeier D., Freytag B., 2012, Phil. Trans. R. Soc. A, 370, 2765

Andrews J. J., Chanamé J., Agüeros M. A., 2017, MNRAS, 472, 675

Ansdell M. et al., 2018, ApJ, 859, 21

Arenou F. et al., 2018, A&A, 616, A17

Artymowicz P., Clarke C. J., Lubow S. H., Pringle J. E., 1991, ApJ, 370, L35

Bate M. R., 2000, MNRAS, 314, 33

Bate M. R., 2009, MNRAS, 392, 590

Bate M. R., 2012, MNRAS, 419, 3115

Bate M. R., 2014, MNRAS, 442, 285

Bate M. R., 2018, MNRAS, 475, 5618

Bate M. R., 2019, MNRAS, 484, 2341

Bate M. R., Bonnell I. A., 1997, MNRAS, 285, 33

Bonatto C., Lima E. F., Bica E., 2012, A&A, 540, A137

Bovy J., Rix H.-W., Green G. M., Schlafly E. F., Finkbeiner D. P., 2016, ApJ, 818, 130

Branch D., 1976, ApJ, 210, 392

Brandeker A., Cataldi G., 2019, A&A, 621, A86

Bressan A., Marigo P., Girardi L., Salasnich B., Dal Cero C., Rubele S., Nanni A., 2012, MNRAS, 427, 127

Brinch C., Jørgensen J. K., Hogerheijde M. R., Nelson R. P., Gressel O., 2016, ApJ, 830, L16

Burgasser A. J., Reid I. N., Siegler N., Close L., Allen P., Lowrance P., Gizis J., 2007in Reipurth B., Jewitt D., Keil K., eds, Protostars and Planets V. Univ. Arizona Press, Tucson, AZ, p. 427

Cantrell A. G., Dougan T. J., 2014, MNRAS, 445, 2028

Chambers K. C. et al., 2016, preprint (arXiv:1612.05560)

Chanamé J., Gould A., 2004, ApJ, 601, 289

Chen Y., Girardi L., Bressan A., Marigo P., Barbieri M., Kong X., 2014, MNRAS, 444, 2525

Choi J., Dotter A., Conroy C., Cantiello M., Paxton B., Johnson B. D., 2016, ApJ, 823, 102

Choi J., Conroy C., Ting Y.-S., Cargile P. A., Dotter A., Johnson B. D., 2018, ApJ, 863, 65

Comerón F., Reipurth B., Yen H.-W., Connelley M. S., 2018, A&A, 612, A73

Connelley M. S., Reipurth B., Tokunaga A. T., 2008, AJ, 135, 2526

Czekala I., Chiang E., Andrews S. M., Jensen E. L. N., Torres G., Wilner D. J., Stassun K. G., Macintosh B., 2019, preprint (arXiv:1906.03269)

De Rosa R. J. et al., 2014, MNRAS, 437, 1216

de Val-Borro M., Gahm G. F., Stempels H. C., Pepliński A., 2011, MNRAS, 413, 2679

Dhital S., West A. A., Stassun K. G., Bochanski J. J., 2010, AJ, 139, 2566

Dhital S., West A. A., Stassun K. G., Schluns K. J., Massey A. P., 2015, AJ, 150, 57

Dolphin A. E., 2002, MNRAS, 332, 91

Dorval J., Boily C. M., Moraux E., Roos O., 2017, MNRAS, 465, 2198

Duchêne G., 1999, A&A, 341, 547

Duchêne G., Kraus A., 2013, ARA&A, 51, 269

Duchêne G., Lacour S., Moraux E., Goodwin S., Bouvier J., 2018, MNRAS, 478, 1825

Duquennoy A., Mayor M., 1991, A&A, 500, 337

Eggenberger A., Halbwachs J.-L., Udry S., Mayor M., 2004in Allen C., Scarfe C., eds, Rev. Mex. Astron. Astrofis. Conf. Ser. Vol. 27. p. 28

Eggleton P. P., Fitchett M. J., Tout C. A., 1989, ApJ, 347, 998

Eisner J. A. et al., 2018, ApJ, 860, 77

El-Badry K., Rix H.-W., 2018, MNRAS, 480, 4884

El-Badry K., Rix H.-W., 2019, MNRAS, 482, L139

El-Badry K., Weisz D. R., Quataert E., 2017, MNRAS, 468, 319

El-Badry K., Rix H.-W., Ting Y.-S., Weisz D. R., Bergemann M., Cargile P., Conroy C., Eilers A.-C., 2018a, MNRAS, 473, 5043

El-Badry K. et al., 2018b, MNRAS, 476, 528

El-Badry K., Rix H.-W., Weisz D. R., 2018c, ApJ, 860, L17

Evans D. W. et al., 2018, A&A, 616, A4

Farris B. D., Duffell P., MacFadyen A. I., Haiman Z., 2014, ApJ, 783, 134

Flewelling H. A. et al., 2016, preprint (arXiv:1612.05243)

Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, PASP, 125, 306

Gaia Collaboration, 2016, A&A, 595, A1

Gaia Collaboration, 2018, A&A, 616, A1

Geha M. et al., 2013, ApJ, 771, 29

Ghez A. M., Neugebauer G., Matthews K., 1993, AJ, 106, 2005

Giannuzzi M. A., 1987, Ap&SS, 135, 245

Goodwin S. P., 2010, Phil. Trans. R. Soc. A, 368, 851

Gullikson K., Kraus A., Dodson-Robinson S., 2016, AJ, 152, 40

Halbwachs J. L., 1988, Ap&SS, 142, 139

Halbwachs J. L., Mayor M., Udry S., Arenou F., 2003, A&A, 397, 159

Hanawa T., Ochi Y., Ando K., 2010, ApJ, 708, 485

Heggie D. C., 1975, MNRAS, 173, 729

Hills J. G., 1975, AJ, 80, 809

Hioki T. et al., 2007, AJ, 134, 880

Hogeveen S. J., 1992, Ap&SS, 196, 299

Jurić M. et al., 2008, ApJ, 673, 864

Katz D. et al., 2019, A&A, 622, A205

Kounkel M. et al., 2019, AJ, 157, 196

Kouwenhoven M. B. N., Goodwin S. P., Parker R. J., Davies M. B., Malmberg D., Kroupa P., 2010, MNRAS, 404, 1835

Kraus A. L., Ireland M. J., Martinache F., Hillenbrand L. A., 2011, ApJ, 731, 8

Kroupa P., 1995, MNRAS, 277, 1507

Kroupa P., 1998, MNRAS, 298, 231

Kroupa P., 2001, MNRAS, 322, 231

Kroupa P., Burkert A., 2001, ApJ, 555, 945

Krumholz M. R., Thompson T. A., 2007, ApJ, 661, 1034

Kuiper G. P., 1935, PASP, 47, 15

Leinert C., Zinnecker H., Weitzel N., Christou J., Ridgway S. T., Jameson R., Haas M., Lenzen R., 1993, A&A, 278, 129

Lépine S., Bongiorno B., 2007, AJ, 133, 889

Lindegren L. et al., 2018, A&A, 616, A2

Lomax O., Whitworth A. P., Hubber D. A., Stamatellos D., Walch S., 2015, MNRAS, 447, 1550

Lucy L. B., 2006, A&A, 457, 629

Lucy L. B., Ricco E., 1979, AJ, 84, 401

Marks M., Kroupa P., 2012, A&A, 543, A8

Marks M., Kroupa P., Oh S., 2011, MNRAS, 417, 1684

Mathieu R. D., 1994, ARA&A, 32, 465

Matsumoto T., Saigo K., Takakuwa S., 2019, ApJ, 871, 36

Mazeh T., Goldberg D., Duquennoy A., Mayor M., 1992, ApJ, 401, 265

Mazeh T., Simon M., Prato L., Markus B., Zucker S., 2003, ApJ, 599, 1344

Miranda R., Muñoz D. J., Lai D., 2017, MNRAS, 466, 1170

Moe M., Di Stefano R., 2013, ApJ, 778, 95

Moe M., Di Stefano R., 2017, ApJS, 230, 15

Moe M., Kratter K. M., Badenes C., 2019, ApJ, 875, 61

Moeckel N., Bate M. R., 2010, MNRAS, 404, 721

Moody M. S. L., Shi J.-M., Stone J. M., 2019, ApJ, 875, 66

Muñoz D. J., Miranda R., Lai D., 2019, ApJ, 871, 84

Murphy S. J., Moe M., Kurtz D. W., Bedding T. R., Shibahashi H., Boffin H. M. J., 2018, MNRAS, 474, 4322

Nelson A. F., Marzari F., 2016, ApJ, 827, 93

Nordström B. et al., 2004, A&A, 418, 989

Ochi Y., Sugimoto K., Hanawa T., 2005, ApJ, 623, 922

Offner S. S. R., Kratter K. M., Matzner C. D., Krumholz M. R., Klein R. I., 2010, ApJ, 725, 1485

Öpik E., 1924, Publ. Tartu Astrofiz. Obs., 25, 1

Parker R. J., Reggiani M. M., 2013, MNRAS, 432, 2378

Parker R. J., Goodwin S. P., Kroupa P., Kouwenhoven M. B. N., 2009, MNRAS, 397, 1577

Pinsonneault M. H., Stanek K. Z., 2006, ApJ, 639, L67

Raghavan D. et al., 2010, ApJS, 190, 1

Reggiani M. M., Meyer M. R., 2011, ApJ, 738, 60

Reipurth B., Mikkola S., 2012, Nature, 492, 221

Reipurth B., Guimarães M. M., Connelley M. S., Bally J., 2007, AJ, 134, 2272

Sadavoy S. I., Stahler S. W., 2017, MNRAS, 469, 3881

Seabroke G. M., Gilmore G., 2007, MNRAS, 380, 1348

Shahaf S., Mazeh T., 2019, MNRAS, 487, 3356

Shatsky N., Tokovinin A., 2002, A&A, 382, 92

Shaya E. J., Olling R. P., 2011, ApJS, 192, 2

Shi J.-M., Krolik J. H., Lubow S. H., Hawley J. F., 2012, ApJ, 749, 118

Simon M., Obbie R. C., 2009, AJ, 137, 3442

Söderhjelm S., 2000, Astron. Nachr., 321, 165

Söderhjelm S., 2007, A&A, 463, 683

Sollima A., 2019, preprint (arXiv:1904.05646)

Sterzik M. F., Durisen R. H., 1998, A&A, 339, 95

Takakuwa S., Saigo K., Matsumoto T., Saito M., Lim J., Hanawa T., Yen H.-W., Ho P. T. P., 2017, ApJ, 837, 86

Tang Y., MacFadyen A., Haiman Z., 2017, MNRAS, 469, 4258

Ting Y.-S., Rix H.-W., 2018, preprint (arXiv:e-print)

Tobin J. J. et al., 2016, Nature, 538, 483

Tokovinin A. A., 2000, A&A, 360, 997

Tokovinin A., 2011, AJ, 141, 52

Tokovinin A., 2014, AJ, 147, 87

Tokovinin A., 2017a, MNRAS, 468, 3461

Tokovinin A., 2017b, ApJ, 844, 103

Tokovinin A. A., Smekhov M. G., 2002, A&A, 382, 118

Trimble V., 1974, AJ, 79, 967

Trimble V., 1987, Astron. Nachr., 308, 343

Trimble V., 1990, MNRAS, 242, 79

van Biesbroeck G., 1916, AJ, 29, 173

Weinberg M. D., Shapiro S. L., Wasserman I., 1987, ApJ, 312, 367

White R. J., Ghez A. M., 2001, ApJ, 556, 265

York D. G. et al., 2000, AJ, 120, 1579

Young M. D., Clarke C. J., 2015, MNRAS, 452, 3085

Ziegler C. et al., 2018, AJ, 156, 259

# APPENDIX A: FUNCTIONAL FORM OF $P(Q)$

A1 How sharp is the twin feature?

In our default model, the twin feature is characterized by two number: F_{twin}, the fractional excess of binaries with near-equal
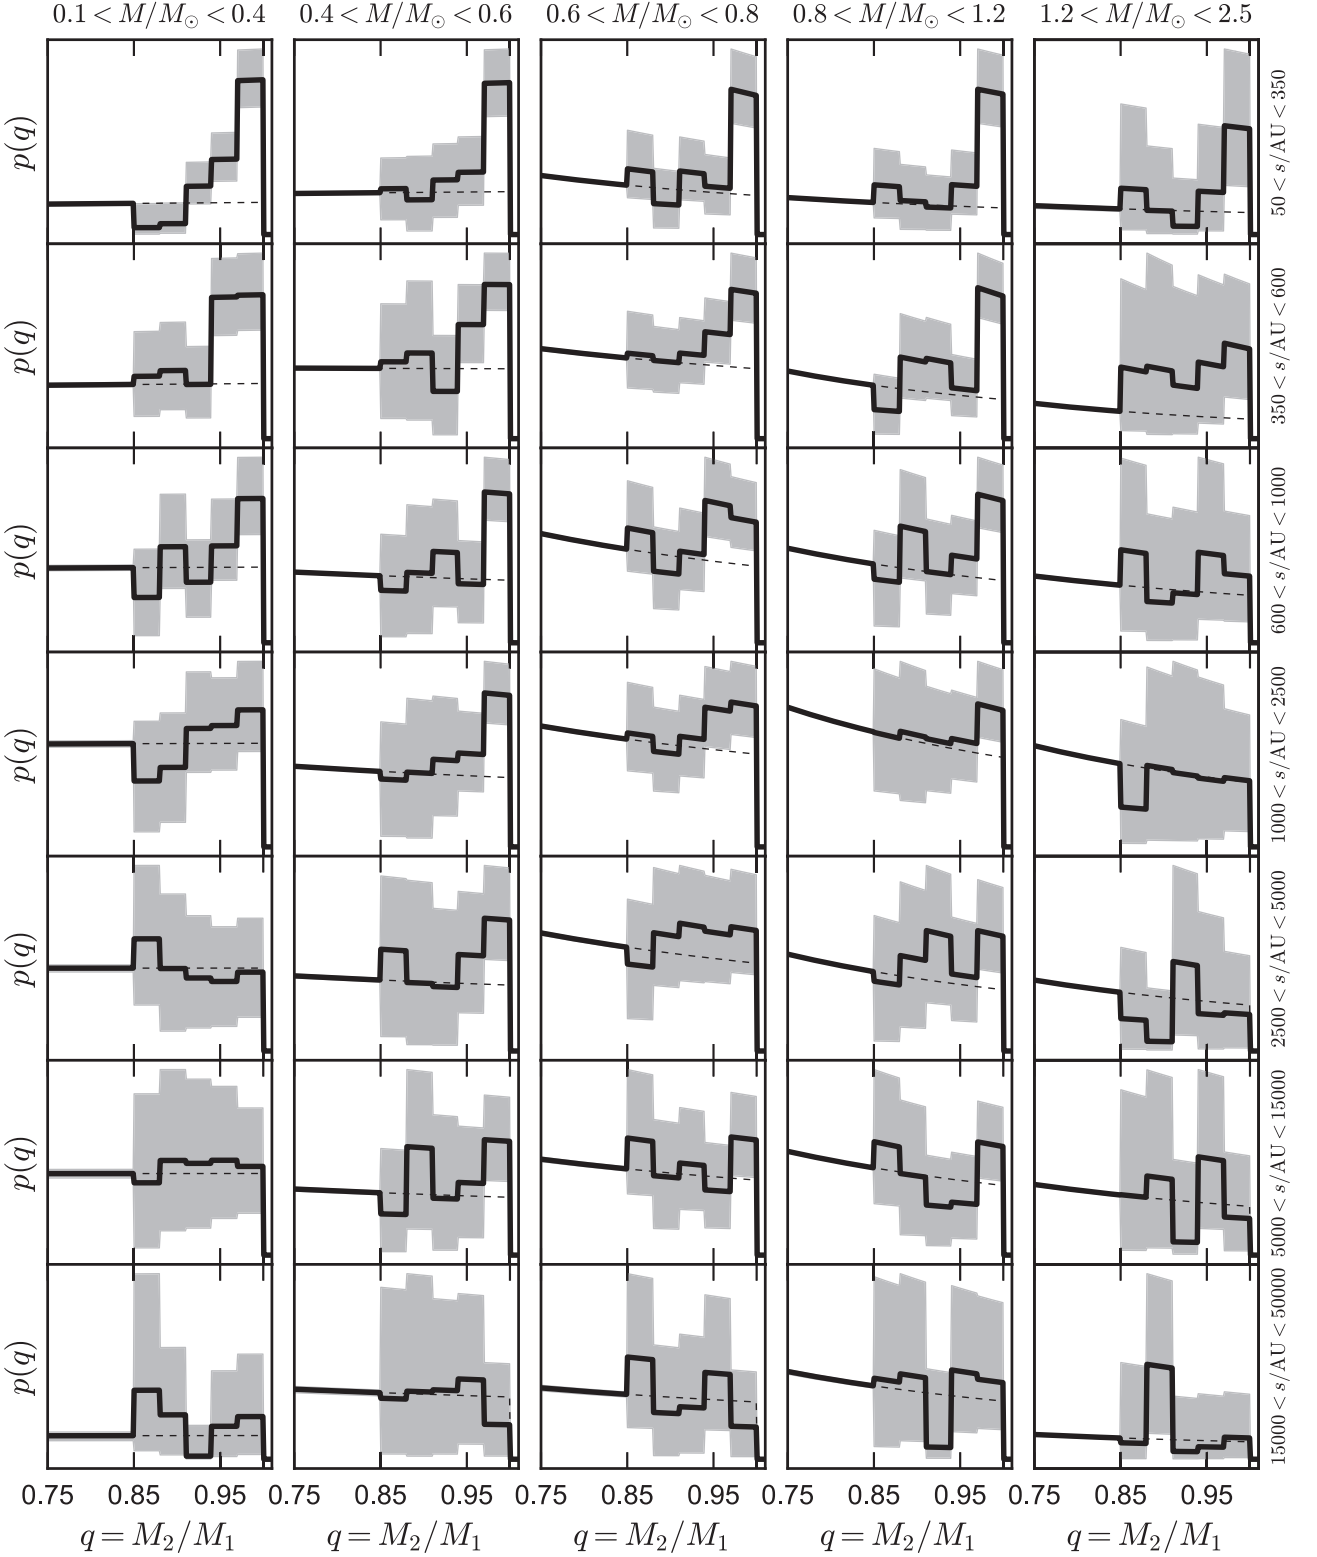
Figure A1. Constraints on $p(q)$ near $q = 1$. Solid lines and grey shaded regions represent median and 2σ (middle 95.4 per cent) constraints; y-axis scale is linear. We model the mass ratio distribution over $0.85 < q < 1$ as a histogram with 5 bins of width 0.03, fitting for the height of each bin as a free parameter. The twin excess generally becomes statistically significant only at $0.97 < q < 1$, in some cases with a modest enhancement at $0.94 < q < 0.97$. At lower q, the mass ratio distribution is generally consistent with a smooth power law (dashed line).
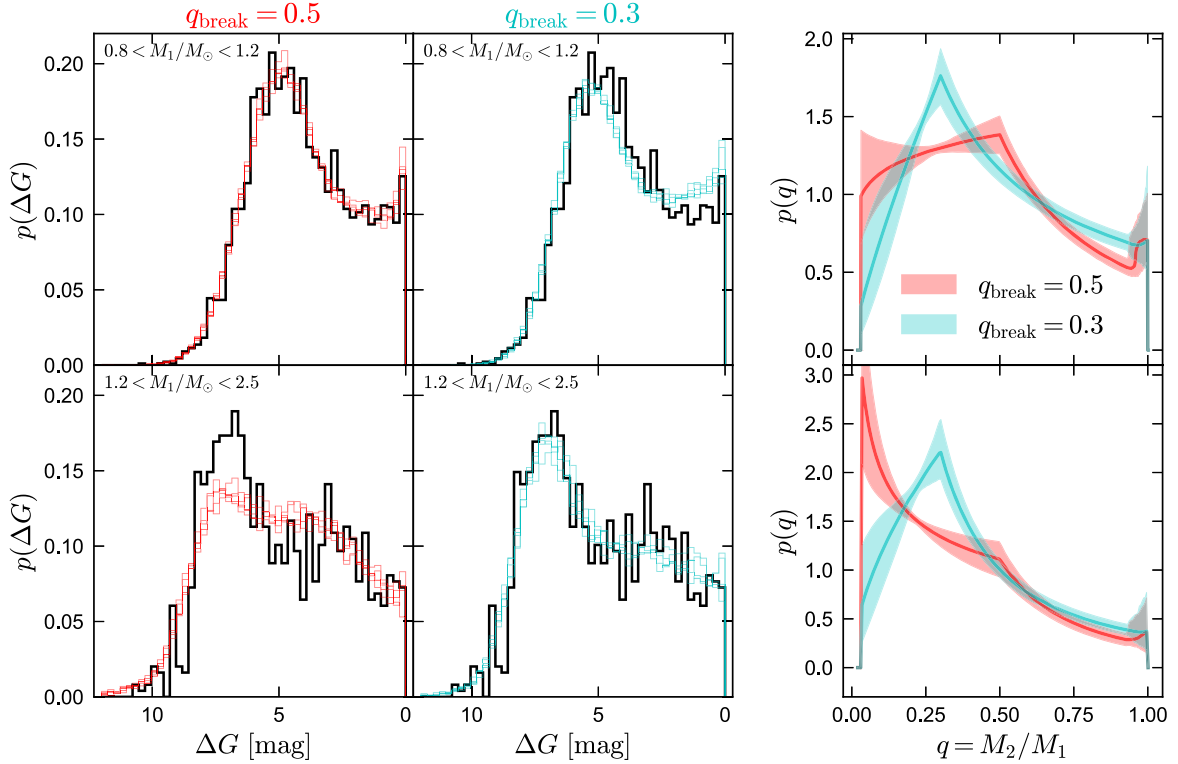
Figure A2. Mass ratio distribution constraints (right) and predicted ΔG distributions (left) for fits with $q_{break} = 0.5$ (red) and $q_{break} = 0.3$ (cyan). We show a single separation bin ($1000 < s < 2500$). Top panels show solar-type stars ($0.8 < M_1/M_\odot < 1.2$); bottom panels show A and F stars ($1.2 < M_1/M_\odot < 2.5$). For each choice of q_{break}, we show the best-fitting constraints when q_{break} is fixed at that value and other parameters are left free. For solar-type stars, a significantly better fit can be obtained with $q_{break} = 0.5$ than with $q_{break} = 0.3$. The opposite is true for A and F stars, where the data strongly favour $q_{break} = 0.3$.

masses, and q_{twin}, the mass ratio above which the excess manifests itself. The implicit assumption in this model is that (a) the increase in the mass ratio distribution near $q = 1$ occurs abruptly, not gradually, and (b) the amount of twin excess is constant between $q = q_{twin}$ and $q = 1$. In this section, we assess the validity of this assumption by using a more flexible model for the twin excess.

We model the twin excess between $q = 0.85$ and $q = 1$ as a histogram with five equally spaced bins of width $\Delta q = 0.03$. As in the fiducial model, we assume the smooth underlying distribution in this region is $p(q) \sim q^{\eta_{largeq}}$. We then introduce a set of uncorrelated weights, w_i with $i = 0 \ldots 5$, and multiply $p(q)$ in the ith bin by w_i. With this parametrization, $w_i = 1$ entails no deviation from a smooth power law in the ith bin, $w_i = 2$ corresponds to a factor-of-two enhancement relative to the underlying power law, and $w_i = 0$ means that there are no binaries at all with mass ratios in the ith bin. With a sufficiently large number of bins, this parametrization is flexible enough to represent any arbitrary shape of $p(q)$. Five bins of width 0.03 is a pragmatic choice given the size of our binary catalogue, since the shot noise uncertainty increases as the bin size is reduced.

Fig. A1 shows results of fitting a histogram model. Consistent with Fig. 9, a statistically significant excess of twins is found out to several thousand au for $0.1 < M_1/M_\odot < 1.2$ and $s \lesssim 2500$ au, with the excess reaching $s = 15\,000$ au for $0.4 < M_1/M_\odot < 0.6$. In most mass and separation bins, the data is consistent with a smooth power-law distribution (with no twin excess) all the way up to $q = 0.97$. In a few mass bins, there is also a significant excess relative to the underlying power law at $0.94 < q < 0.97$, but never at $q < 0.94$. That is, in all bins where it is significant, the twin feature is 'thin': there is not a broad peak at $q \gtrsim 0.8$, but a narrow excess only

at $q \gtrsim 0.95$. It is because of the narrowness of the twin peak that the excess of equal-mass pairs is so clearly apparent in the observed data (e.g. Fig. 3), even though twins only make up a small fraction (<5 per cent; see Fig. 11) of the total binary population at wide separations.

A2 CHOICE OF q_{break}

As discussed in Section 3, we model the mass ratio distribution in all but the lowest mass bin as a broken power law whose slope changes at $q = q_{break}$ (Fig. 7). In order to make comparison between different separation bins more straightforward, we do not leave q_{break} free, but instead fix it to a single value in a given mass bin: $q_{break} = 0.3$ for $1.2 < M_1/M_\odot < 2.5$, and $q_{break} = 0.5$ in all other mass bins. These values were chosen by trial and error, but it *is* necessary to use different values of q_{break} in different mass bins, as we now show.

Fig. A2 shows constraints on the mass ratio distribution (right) and corresponding predicted ΔG distributions (left) for two choices of q_{break} and two bins of primary mass. A single separation bin is shown for illustrative purposes: $1000 < s < 2500$, where the twin excess is weak in both mass bins. In the left-hand panels, we compare the observed distributions of ΔG to Monte Carlo binary populations produced for five draws from the posterior.

The top panels show that for solar-mass primaries ($0.8 < M_1/M_\odot < 1.2$), the best-fitting model obtained while assuming $q_{break} = 0.5$ provides a significantly better fit to the observed ΔG distribution than the best-fitting model with $q_{break} = 0.3$: in the latter case, the observed distribution is poorly reproduced both at $0 < \Delta G \lesssim 2$ and at $\Delta G \approx 5$. On the other hand, the bottom panels show that $q_{break} = 0.3$ provides a much better fit for A and F star primaries.
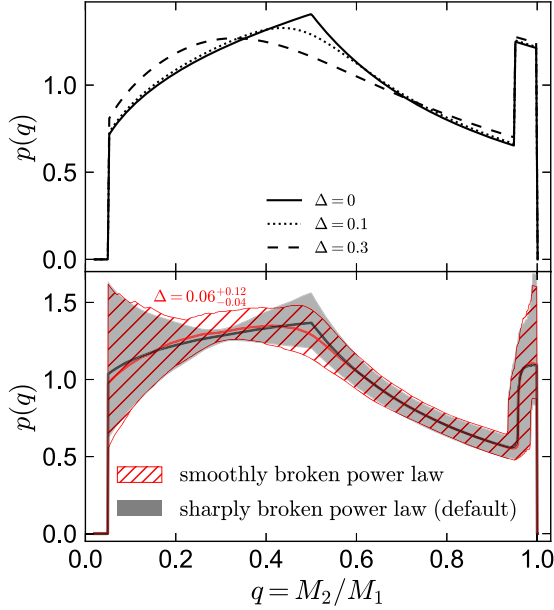
Figure A3. Top panel shows example smoothly broken power laws for different values of the smoothing parameter, Δ. For all cases, $\gamma_{\mathrm{smallq}} = 0.3$, $\gamma_{\mathrm{largeq}} = -1.2$, $F_{\mathrm{twin}} = 0.04$, and $q_{\mathrm{twin}} = 0.95$. Increasing Δ softens the transition between $p \sim q^{\gamma_{\mathrm{smallq}}}$ and $p \sim q^{\gamma_{\mathrm{largeq}}}$. Bottom: Middle 95.4 per cent constraints on $p(q)$ for $0.8 < M_1/M_{\odot} < 1.2$ and $600 < s/\mathrm{au} < 1000$. We compare constraints obtained by fitting a smoothly broken power law (red hatched) to those obtained by fitting our default model, a simple double power law (i.e. forcing $\Delta = 0$). The constraints we obtain with the two models are similar, with the most significant difference being that the smoothly broken model lacks (a) the sharp break at $q_{\mathrm{break}} = 0.5$ and (b) the artificial suppression of uncertainty near a pivot point at $q = 0.3$.

A3 Smoothly broken power law

As discussed in Section 4, our choice to model the mass ratio distribution with a broken power law leads to an unphysical sharp break in the best-fitting mass ratio distribution at intermediate q. We test the sensitivity of our constraints on $p(q)$ to the assumed functional form below. As an alternative to a sharply broken power law, we fit a 'smoothly broken power law' with the functional form

$$p\,(q) \propto \left(\frac{q}{q_{\mathrm{break}}}\right)^{\gamma_{\mathrm{smallq}}} \left[1 + \left(\frac{q}{q_{\mathrm{break}}}\right)^{1/\Delta}\right]^{(\gamma_{\mathrm{largeq}} - \gamma_{\mathrm{smallq}})\Delta}. \quad (A1)$$

This function approaches $p \propto q^{\gamma_{\mathrm{smallq}}}$ at $q/q_{\mathrm{break}} \ll 1$, and $p \propto q^{\gamma_{\mathrm{largeq}}}$ at $q/q_{\mathrm{break}} \gg 1$. The parameter Δ controls the sharpness of the transition between the two regimes; in the limit of $\Delta \to 0$, it reduces to a simple double power law with a sharp transition between the two slopes.

In Fig. A3, we compare the constraints on $p(q)$ obtained for a single mass and separation bin when using the fiducial model (grey) and a smoothly broken power law (hatched red). Overall, the constraints are quite similar, but as expected, fitting a broken power-law model smooths the best-fitting profile for $p(q)$ and removes the region near $q \sim 0.3$ where the nominal uncertainty is suppressed.

We find qualitatively similar results for other bins of mass and separation. However, we note that increasing Δ shifts the peak of the distribution towards lower q (upper panel of Fig. A3). Since Δ itself is often only weakly constrained, this can lead to parameter covariances between Δ, γ_{smallq}, and q_{break}. We therefore use the simpler broken power law as our fiducial model.

APPENDIX B: SENSITIVITY TO SYSTEMATICS

The uncertainties we report on mass ratio distribution parameters represent formal fitting uncertainties due to Poisson errors, but they do not include various systematic uncertainties due to modelling choices we make that are held fixed during fitting. Here, we vary several aspects of the model to assess how sensitively our constraints depend on them. Fig. B1 shows the effects of varying the assumed IMF, stellar age distribution, and the isochrones used to generate synthetic photometry, on our parameter constraints. We show a single primary mass and separation bin with typical uncertainties and sensitivity to systematics.

Varying the IMF (blue) has very weak effects on our constraints, primarily because we fit narrow bins in primary mass independently. Our constraints are also relatively insensitive to the assumed star formation history (brown). Varying the age distribution has several effects: it changes the number of low-mass companions that are still undergoing Kelvin–Helmholtz contraction and are thus brighter, it changes the mass–luminosity relation at $M \gtrsim 0.7\,M_{\odot}$, where age effects are non-negligible; and it changes the distance distribution implied in our model because we adopt an age-dependent scale height (equation 3). Varying the model for unresolved subsystems has similarly weak effects. The green lines in Fig. B1 show constraints obtained when unresolved subsystems are ignored entirely in the model (i.e. none of the synthetic wide binary components are assigned an unresolved companion). The thus-obtained constraints are nearly indistinguishable from those obtained when subsystems are included in the model.

Changing the adopted stellar models from PARSEC to MIST (Choi et al. 2016) has the strongest effect on our constraints. Because the two models predict modestly different mass–luminosity relations, a given ΔG corresponds to different q in the two models. This has weak effects on the constraints on F_{twin} and q_{twin}: twins necessarily have similar masses, and the mapping between luminosity and mass varies less over a small range of masses than over a large one. Differences between the models are largest for low-mass stars, where isochrones are known to be more uncertain.

We note that while there clearly *are* uncertainties in our results associated with the stellar models, the PARSEC models appear to fit the *Gaia* data for low-mass stars significantly better than the MIST models. Particularly on the lower MS $M_{\mathrm{G}} \gtrsim 10$, we are unable to reproduce the morphology of the CMD unless we assume a higher-than-expected mean metallicity of $\langle [\mathrm{Fe/H}] \rangle \approx +0.2$. This is why we use the PARSEC synthetic photometry in our fiducial model. We note that the PARSEC models use surface boundary conditions that are empirically calibrated to match the observed mass–radius relation at low masses (Chen et al. 2014); without this calibration, the tension between their predictions and the observed lower MS is similar to that found for MIST models (e.g. Choi et al. 2018).
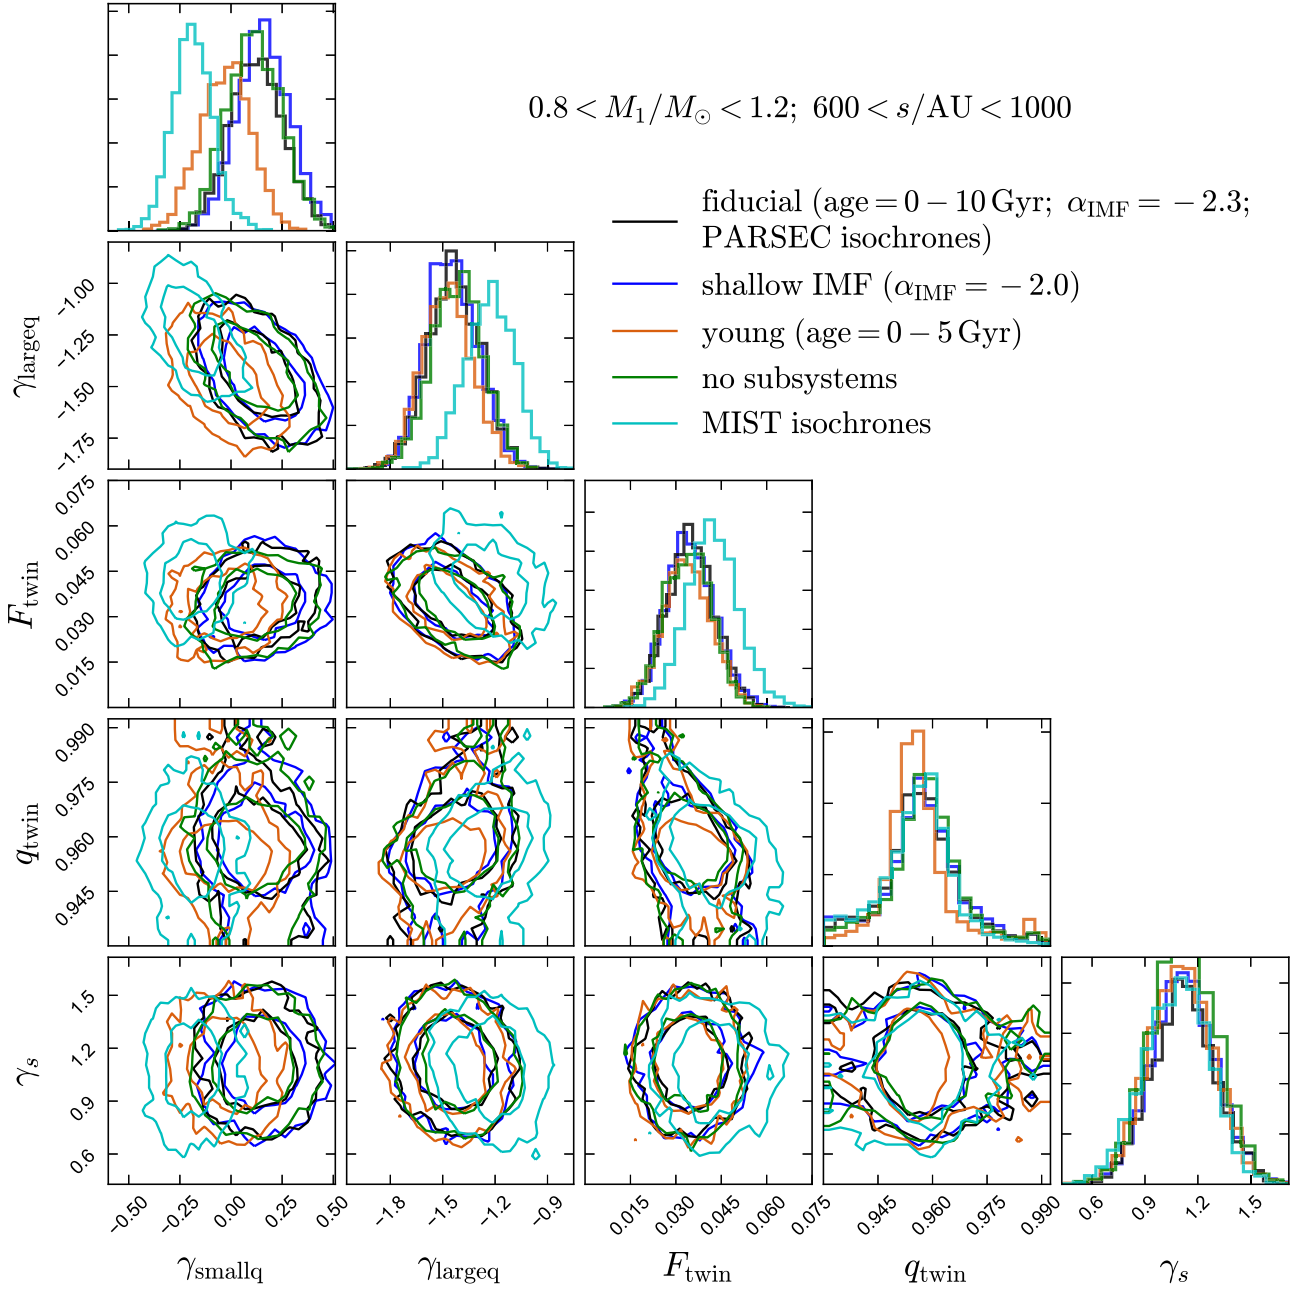
Figure B1. Sixty-eight and ninety-five per cent probability contours for one representative bin of primary mass and physical separation. Panels on the diagonal show marginalized probability distributions. Black contours shows results for fitting the fiducial model. Blue, brown, green, and cyan contours show constraints obtained when the assumed IMF, star formation history, model for unresolved binaries, and stellar models are varied. Overall, constraints are not very sensitive to the assumed IMF, SFH, or unresolved binary model; this is particularly true for $F_{\rm twin}$ and $q_{\rm twin}$. They are moderately sensitive to the choice of stellar models, particularly at low-mass ratios.

APPENDIX C: TWINS IN OTHER CATALOGUES

Another large sample of low-mass wide binaries is the SLoWPoKES-II catalogue created by Dhital et al. (2015) using SDSS photometry. This catalogue was produced without astrometry and extends to larger distances than our catalogue ($d \lesssim 1$ kpc). Despite the lack of parallaxes and proper motions Dhital et al. (2015) were able to ensure relatively low contamination by limiting their search to close angular separations ($\theta < 20$ arcsec) and using

an isochrone prior (i.e. requiring both components of a binary to have similar photometric distance). The selection function of their catalogue is quite different from ours, so we do not attempt a full probabilistic model of the mass ratio distribution. However, we check whether a twin excess is visible in the magnitude difference distributions of their catalogue and whether its strength at fixed physical separation is consistent with that found in our catalogue.

Fig. C1 (top panel) shows the distribution of Δr, the difference in apparent r-band magnitude between the two components, as
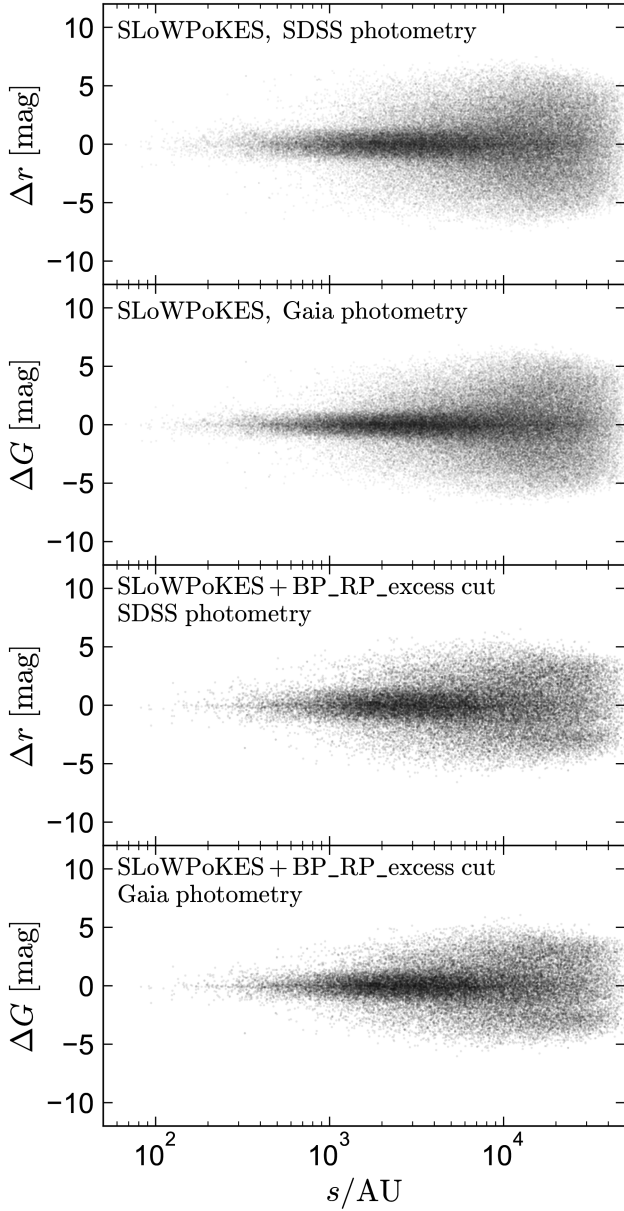
Figure C1. Excess of twins in the SLoWPoKES-II catalogue (Dhital et al. 2015). Top panel shows the full catalogue with SDSS photometry. Here there is no obvious excess of equal-brightness twins. Middle panel shows *Gaia* photometry from the full SLoWPoKES-II catalogue. The *Gaia* photometry is significantly more precise than the ground-based SDSS photometry, so a slight excess of twins is apparent, but the signal is somewhat weaker than that found in our catalogue. The bottom two panels show SDSS and *Gaia* photometry for the subset of the SLoWPoKES-II catalogue for which the *Gaia* photometry passes our cut on `phot_bp_rp_excess_factor`. This cut removes binaries in which the photometry for either component is contaminated (either by the binary companion or by a background star). Once sources with contaminated photometry have been removed, an excess of twins is visible in both the SDSS and *Gaia* photometry.

a function of physical separation for the full SLoWPoKES-II catalogue. Projected physical separations are calculated using the mean photometric distance of the two stars estimate by Dhital et al. (2015). We verified that for stars that are bright enough to have precise *Gaia* parallaxes, the photometric and geometric distances
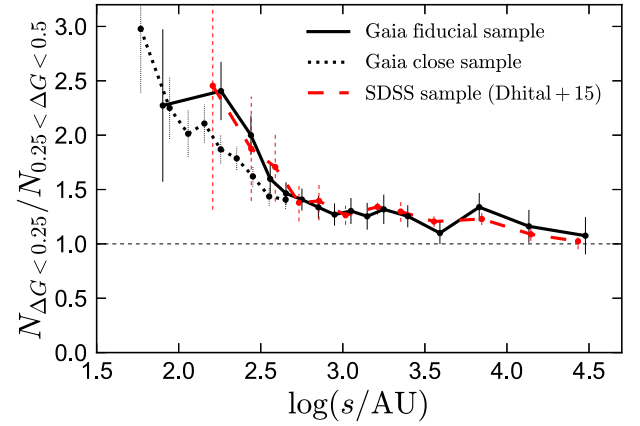
Figure C2. Ratio of the number of binaries with nearly equal magnitudes ($\Delta G < 0.25$) to the number with slightly different magnitude ($0.25 < \Delta G < 0.5$) as a function of projected physical separation (right). Error bars are 1σ. We compare the fiducial sample from this work (solid black line) to the subsample of the SLoWPoKES-II catalogue (Dhital et al. 2015) with uncontaminated *Gaia* photometry (dashed red line; bottom panel of Fig. C1) and the sample of *Gaia* binaries from El-Badry & Rix (2019), which reaches closer angular separations than our fiducial sample but passes less stringent quality cuts.

are usually in reasonably good agreement. The SLoWPoKES-II catalogue primarily contains binaries with $q \gtrsim 0.5$: there are few pairs with $\Delta r > 5$, and at physical separations $s \lesssim 2000\,\mathrm{au}$, most pairs have $\Delta r < 2$. This is primarily a consequence of the catalogue selection function. No clear excess of equal-brightness twins is apparent in the top panel of Fig. C1: the distribution of Δr at fixed separation is reasonably smooth.

To check whether the lack of obvious twin excess in the full SLoWPoKES-II catalogue is a consequence of the SDSS ground-based photometry being poorer, we cross-matched the catalogue with *Gaia* DR2. A *Gaia* source was found with 1 arcsec for both components for ∼90 per cent of the binaries in the catalogue. In the second panel of Fig. C1, we plot all of these sources (many of which do not pass the photometric and astrometric quality cuts imposed in our catalogue), showing the difference in G-band magnitude. Here, an excess of equal-brightness binaries *is* visible, but it appears somewhat weaker than in our catalogue.

Finally, in the third and fourth panels of Fig. C1, we show the SDSS and *Gaia* photometry for the subset of the SLoWPoKES-II catalogue in which the *Gaia* photometry for both components passes the quality cut on `phot_bp_rp_excess_factor` that is imposed on our catalogue. This cut removes objects in which the sum of the flux in the BP and RP bands is not consistent with the flux in the G band (see Evans et al. 2018). Because the fluxes in the BP and RP bands are dispersed over several arcsec while the G-band flux is obtained by profile fitting a narrower image, this cut efficiently selects objects in which the *Gaia* photometry (both BP/RP and G band) is contaminated by a nearby source. Once this cut is applied, a stronger excess of equal-brightness binaries is apparent in both the *Gaia* and SDSS photometry. We show below that for the subsample of the SLoWPoKES-II catalogue with uncontaminated *Gaia* photometry, the twin excess at fixed separation is consistent with that found in our catalogue.

A natural worry is that the cut on `phot_bp_rp_excess_factor` could somehow select against pairs that are close on the sky and do not have nearly identical

brightness or colour, thus erroneously producing an apparent excess of equal-brightness binaries. However, the fact that no excess of equal-brightness pairs is found for chance alignments subject to the same cut speaks against this possibility. Inspecting the SDSS images of binaries in which the *Gaia* photometry for either component does not pass the `phot_bp_rp_excess_factor` cut, we find that in most cases the contamination is quite strong: either the light from the two stars is blended, or another source is blended with one of the components. In such cases, the individual magnitudes of the components cannot be measured with high fidelity, likely leading to underestimated photometric errors.

In Fig. C2, we compare the strength of the twin excess found in the *Gaia* photometry for the SLoWPoKES-II catalogue (once the `phot_bp_rp_excess_factor` cut is applied) at fixed separation (red) to that found in our catalogue (black). The *y*-axis is similar to that in Fig. 5. The strength of the twin excess is consistent between the two catalogues at all separations. We note that at close physical separations, binaries are only spatially resolvable if they are nearby, so at $s \lesssim 500$ au, there is substantial overlap between the two catalogues. At $s \gtrsim 1000$ au, there is little overlap, because most of the SLoWPoKES-II binaries are too distant and faint to enter our catalogue.

We also show in Fig. C2 the excess of equal-brightness binaries found in the catalogue of *Gaia* binaries constructed in El-Badry & Rix (2019), which targeted closer separations. This catalogue only contains binaries with $s < 500$ au. Unlike our primary catalogue, it did not require the components to have measured G_{BP} and G_{RP} magnitudes, and it did not apply any cuts on `phot_bp_rp_excess_factor`. This makes it possible to reach binaries with a factor of ~4 closer angular separations than our primary catalogue. Fig. C2 shows that in this catalogue, the twin excess continues to increase towards smaller separations, at least to $s \approx 50$ au. The twin excess is slightly weaker at fixed separation in this catalogue, likely because the photometry is more contaminated, but it is strongly inconsistent with 0. This provides further evidence that the twin excess is not the result of the `phot_bp_rp_excess_factor` cut creating a bias against binaries with unequal brightness or colour.

In addition to the comparisons to other surveys described above, we have also verified that the twin excess we find is not an artefact of the *Gaia* photometry. Cross-matching our full binary catalogue with the Pan-STARRS1 survey (Chambers et al. 2016; Flewelling et al. 2016), we find a twin excess of similar strength in its photometry. The excess is also apparent in SDSS photometry, but it is somewhat less narrow there due to the larger typical photometric errors.

APPENDIX D: SELECTION FUNCTION

The selection function of *Gaia* DR2 is known to exhibit spatial variation on small scales as a result of the scanning law (e.g. Arenou et al. 2018). Modelling all the small-scale structure in the selection function is beyond the scope of this work, but doing so is not necessary for our purposes. First, spatial variations are almost negligible for sources within the magnitude range of our catalogue ($G \lesssim 18$). More generally, since we do not expect the *intrinsic* properties of the binary population to vary much with on-sky position (particularly on small scales), one can construct an effective selection function averaged over the whole sky without introducing biases (e.g. Bovy et al. 2016).

We model the selection function for binaries as the product of two single-star terms and a cross-term that depends on the angular

separation and magnitude difference of the two stars (equation 2). Both terms are described below.

D1 Single-star term

The single-star term (s_1 and s_2 in equation 2; here we default to s_1) is set primarily by the following cuts we imposed in creating the catalogue:

(i) 5 per cent parallax uncertainty: `parallax_over_error` > 20

(ii) Precise photometry: `phot_bp_mean_flux_over_error` > 20, `phot_rp_mean_flux_over_error` > 20, and `phot_g_mean_flux_over_error` > 20

(iii) Good astrometric model fits: $\sqrt{\chi^2/(\nu' - 5)} < 1.2 \times \max(1, \exp(-0.2(G - 19.5)))$, where χ^2 and ν' are, respectively, referred to as `astrometric_chi2_al` and `astrometric_n_good_obs_al` in the *Gaia* archive.

(iv) Uncontaminated photometry: $1.0 + 0.015(G_{BP} - G_{RP})^2 <$ `phot_bp_rp_excess_factor` $< 1.3 + 0.06(G_{BP} - G_{RP})^2$.

The motivation for these cuts is described in Lindegren et al. (2018). We discuss the effects of each cut below. We note that ER18 also required difference in total proper motion of the two stars to be precise, satisfying $\sigma_{\Delta\mu} < 1.5$ mas yr^{-1}. We find that for the subset of the binaries studied in this work, that cut has a negligible effect on the selection function, as sources that satisfy (i) and (ii) already satisfy it.

For (i), parallax error is expected to depend mainly on apparent magnitude. Fig. D1 shows that the parallax error at a given *G*-band magnitude roughly follows a lognormal distribution, the mean value of which increases for fainter stars. This allows us to calculate the fraction of stars at a given distance and magnitude that will have `parallax_over_error` > 20. In particular, the distribution of ϖ/σ_ϖ at a given distance and magnitude is

$$P(\varpi/\sigma_\varpi) = P(\log \sigma_\varpi) \left| \frac{d \log \sigma_\varpi}{d(\varpi/\sigma_\varpi)} \right|, \qquad (D1)$$

where $P(\log \sigma_\varpi)$ is the (Gaussian) distribution of log parallax error. The fraction of stars with $\varpi/\sigma_\varpi > 20$ is then

$$s_\varpi(G) = \int_{20}^{\infty} P(\varpi/\sigma_\varpi) \, d(\varpi/\sigma_\varpi) \qquad (D2)$$

$$= \frac{1}{2} \left[1 - \mathrm{erf} \left(\frac{\mu_f \ln 10 + \ln(20/\varpi)}{\sqrt{2}\sigma_f \ln 10} \right) \right]. \qquad (D3)$$

Here, μ_f and σ_f represent the mean and dispersion of the log parallax error distributions at magnitude *G*; i.e. the quantities plotted in the lower left panel of Fig. D1. Given equation (D2), we can calculate the probability that a star with given absolute magnitude at a given distance will satisfy `parallax_over_error` > 20; this is shown in the lower right panel of Fig. D1.

A similar strategy can be used to calculate the fraction of stars passing the photometric precision cuts. We find that the cut on `phot_bp_mean_flux_over_error` is the most important: once it is applied and white dwarfs are excluded, all the sources that pass it also pass the other cuts in (ii). This quantity depends primarily on the apparent magnitude in the G_{BP} band. As Fig. D2 shows, the distribution of `phot_bp_mean_flux_over_error` at a given G_{BP} is also roughly lognormal. We use the empirical dependence of the distributions of `phot_bp_mean_flux_over_error` on G_{BP} from the bottom left panel of Fig. D2 to predict the distribution of `phot_bp_mean_flux_over_error` for a hypothetical star with
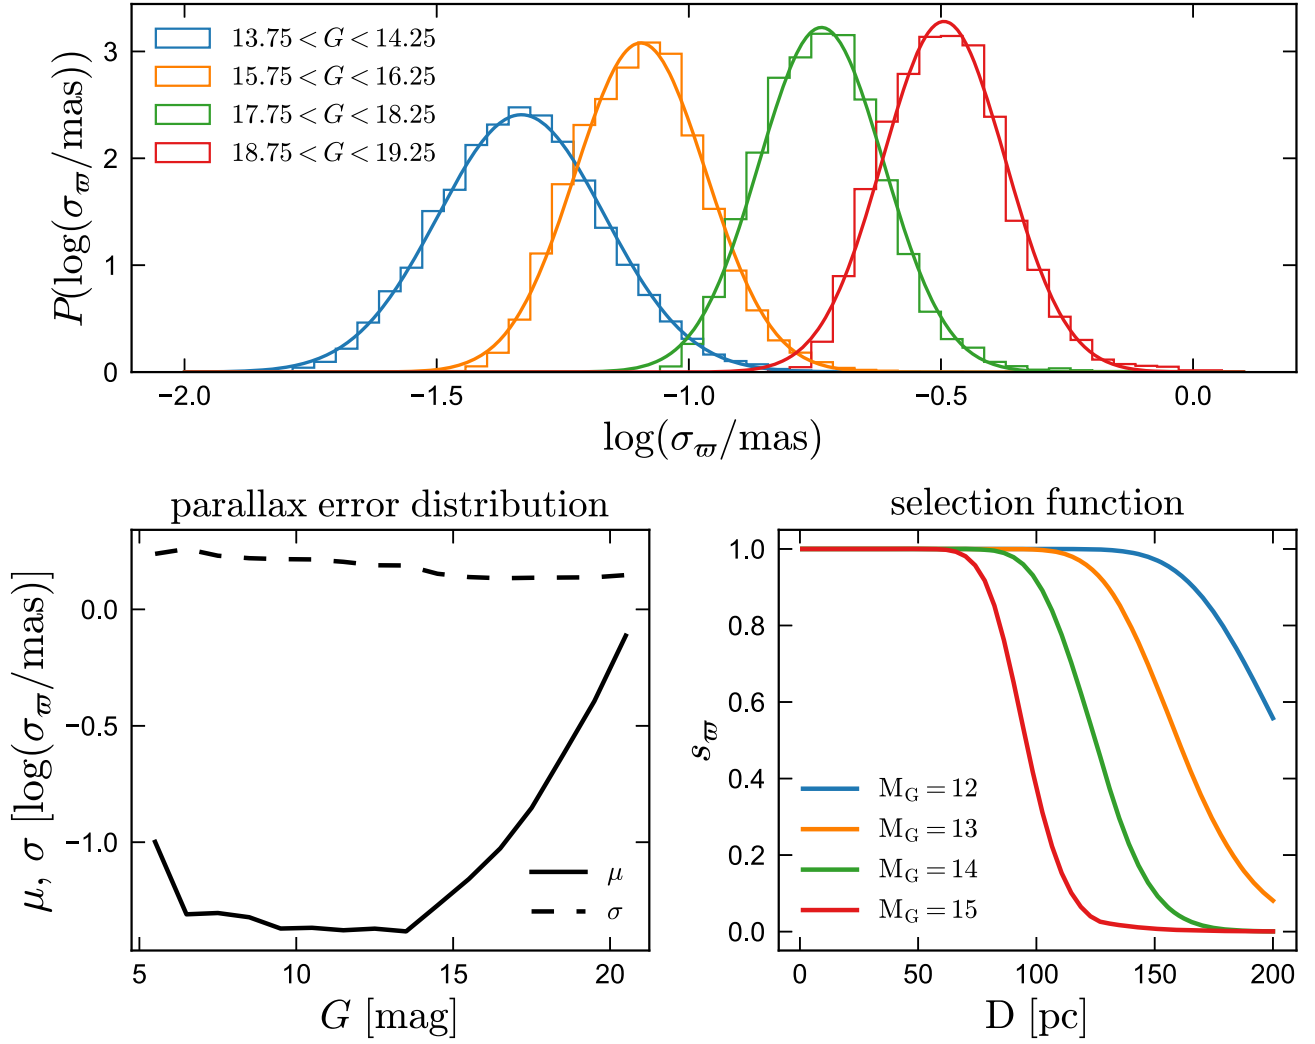
Figure D1. Top panel shows the distributions of log parallax error for nearby stars ($D < 100\,\mathrm{pc}$) in 0.5-mag wide bins of apparent magnitude. We find that the distribution of σ_ϖ at fixed apparent magnitude is roughly lognormal; solid curves show Gaussian fits to the observed histograms. Bottom left panel shows the mean and dispersion of the log parallax error distribution as a function of G-band magnitude. Modelling the distribution of parallax error as a lognormal with μ and σ following this panel allows us to calculate what fraction of stars with a given absolute magnitude and distance will pass the cut of `parallax_over_error > 20` (equation D2; bottom right panel).

absolute BP-band magnitude at a given distance, and then calculate s_{BP} as the fraction of that distribution that exceeds the adopted threshold of 20:

$$s_{\mathrm{BP}}(G_{\mathrm{BP}}) = \int_{20}^{\infty} P\left(\log\left(F_{\mathrm{BP}}/\sigma_{\mathrm{BP}}\right)\right) \mathrm{d}\log\left(F_{\mathrm{BP}}/\sigma_{\mathrm{BP}}\right) \qquad (\mathrm{D4})$$

$$= \frac{1}{2}\left[1 + \mathrm{erf}\left(\frac{\mu_f - \log 20}{\sqrt{2}\sigma_f}\right)\right]. \qquad (\mathrm{D5})$$

Here, $F_{\mathrm{BP}}/\sigma_{\mathrm{BP}}$ represents `phot_bp_mean_flux_over_error`, and μ_f and σ_f represent the mean and dispersion of the log `phot_bp_mean_flux_over_error` distributions at apparent BP-band magnitude G_{BP}; i.e. the quantities plotted in the lower left panel of Fig. D2.

It is also worth considering whether the quality cuts (iii) and (iv) change the single-star selection function. We assess the effects of each quality cut as follows. We begin with the full sample of stars within 200 pc satisfying $\varpi/\sigma_\varpi > 10$. We then apply each quality cut and monitor the effects on the sample. For this nearby sample, it is

straightforward to determine whether most of the sources removed by a particular cut are erroneous (generally distant faint stars with incorrect parallaxes) or real, because distant sources with incorrect parallaxes generally fall in a cloud below the MS (see Lindegren et al. 2018). What we wish to quantify is the fraction of real sources that are removed by each cut, and whether these cuts introduce systematic biases.

Fig. D3 shows the fraction of suspected genuine sources removed by cuts (iii) and (iv). We conclude that cut (iv) has a negligible effect on the single-star selection function. The effects of cut (iii) are also weak, but it does appear to preferentially remove nearby sources. This may be because unresolved astrometric binaries produce large deviations from linear motion of the light centroid at close distances. We tabulate s_{cut} due to cut (iv) as a function of parallax. The total single-star selection function is then

$$s_1 = s_\varpi(G)s_{\mathrm{BP}}(G_{\mathrm{BP}})s_{\mathrm{cut}}(\varpi), \qquad (\mathrm{D6})$$
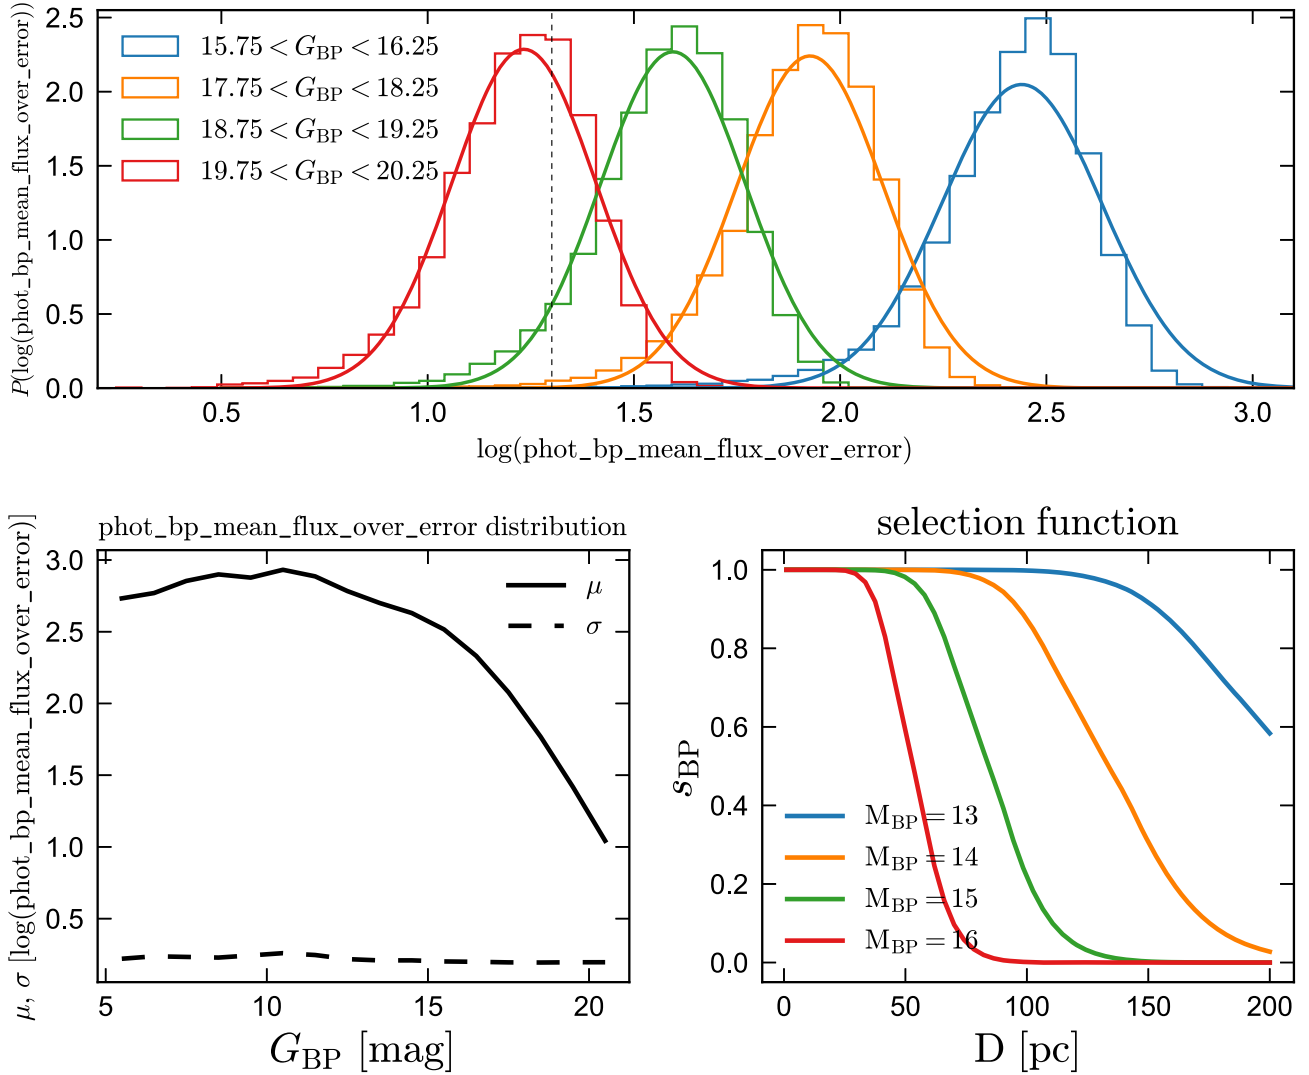
Figure D2. Top panel shows distributions of `phot_bp_mean_flux_over_error` for nearby stars ($D < 100\,\mathrm{pc}$) in different 0.5 mag wide bins of G_{BP} apparent magnitude. Solid lines show lognormal fits; vertical dashed line shows the cut of `phot_bp_mean_flux_over_error` > 20 that both components of a binary must satisfy to enter the catalogue. Bottom left panel shows μ and σ of lognormal fits to these distributions as a function of G_{BP}. From these, we calculate the probability that a star at a given distance and absolute magnitude M_{BP} will have `phot_bp_mean_flux_over_error` > 20 (bottom left panel).

where $s_\varpi(G)$ and $s_{\mathrm{BP}}(G_{\mathrm{BP}})$ are given by equations (D2) and (D4), respectively.

D2 Contrast sensitivity for close pairs

A critical aspect of the selection function for binaries is the reduction in sensitivity to a companion at close angular separations. Whether a binary enters our catalogue depends both on the angular separation of the two stars and on their flux ratio, as a secondary is more likely to be outshone or contaminated by light from the primary at fixed angular separation if the flux ratio is large than if it is small.

The sensitivity to a companion can be measured by comparing the two-point correlation function of chance alignments to what would be expected in the absence of crowding/blending effects (e.g. Arenou et al. 2018; Brandeker & Cataldi 2019). Of course, the sensitivity to companions in a particular catalogue depends on the quality cuts imposed: cuts that remove objects with somewhat contaminated photometry will lower the sensitivity.

We quantified the reduction in sensitivity to a companion as a function of angular separation θ and magnitude difference ΔG in ER18. There we found that, given our cuts on astrometric χ^2 and photometric quality (primarily `bp_rp_excess_factor`), the sensitivity to a companion goes to 0 at $\theta \ll \theta_0$, where $\theta_0 \approx 2$ arcsec for sources with similar magnitude, and θ_0 increases with ΔG. We quantified this dependence by fitting a function:

$$s_{\Delta G}(\theta) = \frac{1}{1 + (\theta/\theta_0)^{-\beta}}, \tag{D7}$$

where $\beta \approx 10$ and we fit for θ_0 as a function of ΔG. We note that the angular resolution of *Gaia* DR2 is actually significantly better than 2 arcsec: most companions are detected down 1 arcsec separations, and the detection fraction only drops to zero at $\theta <$ 0.5 arcsec (Arenou et al. 2018; Ziegler et al. 2018). The $\sim$2 arcsec limit for our catalogue is a result of our requirement that both stars have a measured `bp_rp` colour and the cut on `bp_rp_excess_factor`.
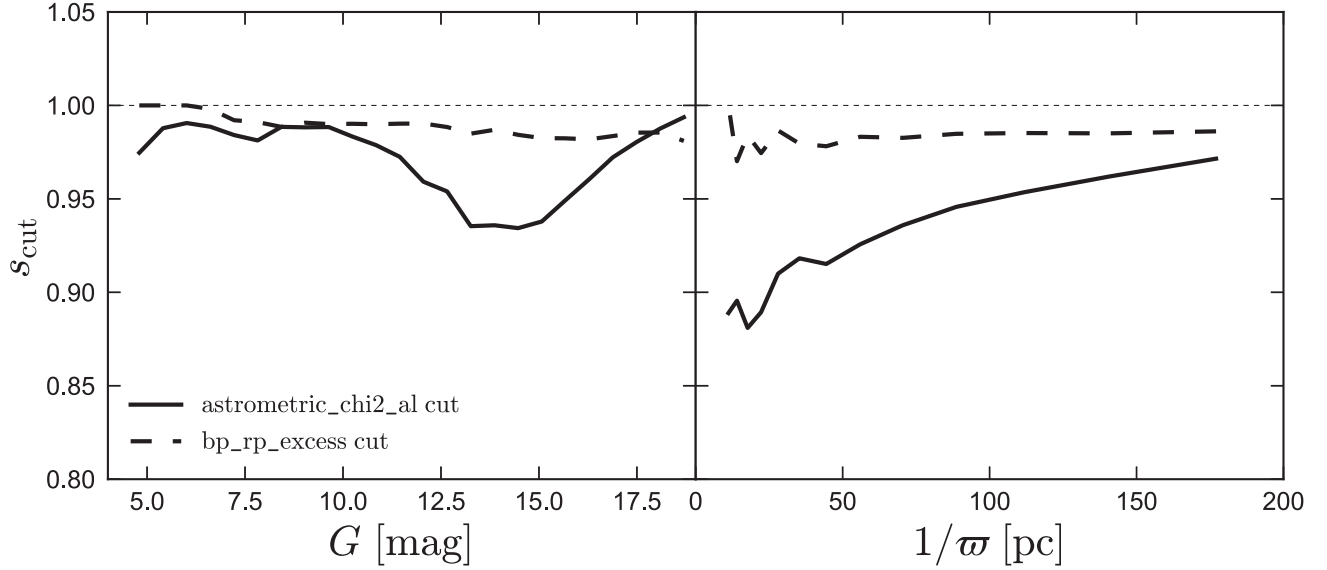
Figure D3. Fraction of suspected genuine sources within 200 pc that survive the cuts we impose on `astrometric_chi2_al` and `bp_rp_excess_factor`, as a function of apparent magnitude (left) and distance (right). Both cuts remove only a small fraction of genuine sources. The cut on `bp_rp_excess_factor` removes sources with contaminated photometry (section D2). The cut on `astrometric_chi2_al` preferentially removes nearby sources, perhaps because perturbations from a companion lead to the larger angular deviations from linear motion for nearby sources.
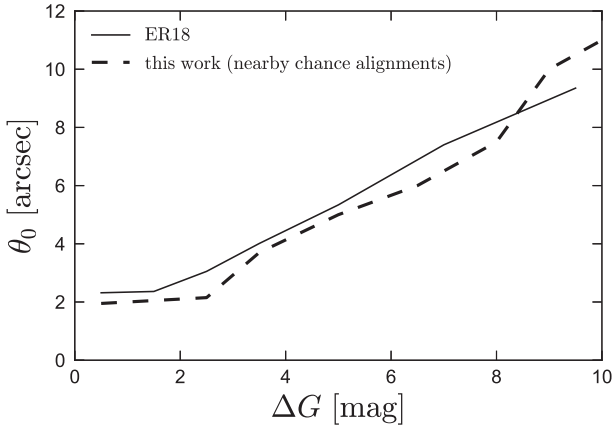
Figure D4. Contrast sensitivity given our quality cuts. θ_0 is the angular separation below which the sensitivity to a companion with magnitude difference ΔG drops off rapidly (see equation D7). Companions can be detected at closer separations when the magnitude difference is small. We compare the values of θ_0 found by ER18 (for all stars passing our quality cuts in a dense field) to those derived from our all-sky catalogue of nearby chance alignments with precise parallaxes.

The dependence of θ_0 on ΔG calculated in ER18 was derived from the source counts of all sources passing our quality cuts (iii) and (iv) in a dense field, most of which are fainter than the objects in the binary catalogue. Here, we improve slightly on the ER18 measurement by repeating their analysis, but using the sources from the chance alignment catalogue described in Section 2 instead of all sources in a dense field. The advantage of this approach is that the sources in the chance alignment catalogue, being nearby and having precise parallaxes, are more representative of objects in the binary catalogue. Unsurprisingly – since it was verified in ER18 that there are no strong trends

in contrast sensitivity with apparent magnitude or source density; see also Brandeker & Cataldi (2019) – the improved constraints are fairly similar to those derived in ER18. We compare them in Fig. D4. On average, we find that the angular resolution at fixed separation is marginally better for the bright chance alignments than was found in ER18.

APPENDIX E: MODEL VALIDATION

To test whether the assumptions of the model we use to fit for the intrinsic mass ratio distribution are valid, we use it to predict properties of the population of single stars that pass the same quality cuts as the binaries in our catalogue (Section E1) and the separation distribution of chance alignments (Section E2).

E1 Single stars

We query the *Gaia* catalogue for all stars with $\varpi > 5$ mas that pass the cuts on astrometric and photometric quality and precision that we require both members of the binary catalogue to pass. Distributions of their distance, apparent magnitude, absolute magnitude, and colour are shown in Fig. E1 with a black histogram. To compare to the model predictions, we draw masses, ages, metallicities, distances from the fiducial distributions described in Section 3, compute synthetic photometry using PARSEC isochrones, and pass the observables through the single-star selection function (Appendix 3.2). The resulting selection function weighted distributions are compared to the data in Fig. E1 (red histogram). In the upper left panel, we also show the distance distributions predicted for a uniform spatial distribution and for an exponential disc with scale height of 300 pc (comparable to the Milky Way; see Jurić et al. 2008), both assuming no incompleteness. The distance distribution of single stars is not that different from the exponential disc prediction, but it begins to deviate at $d \gtrsim 130$ pc,
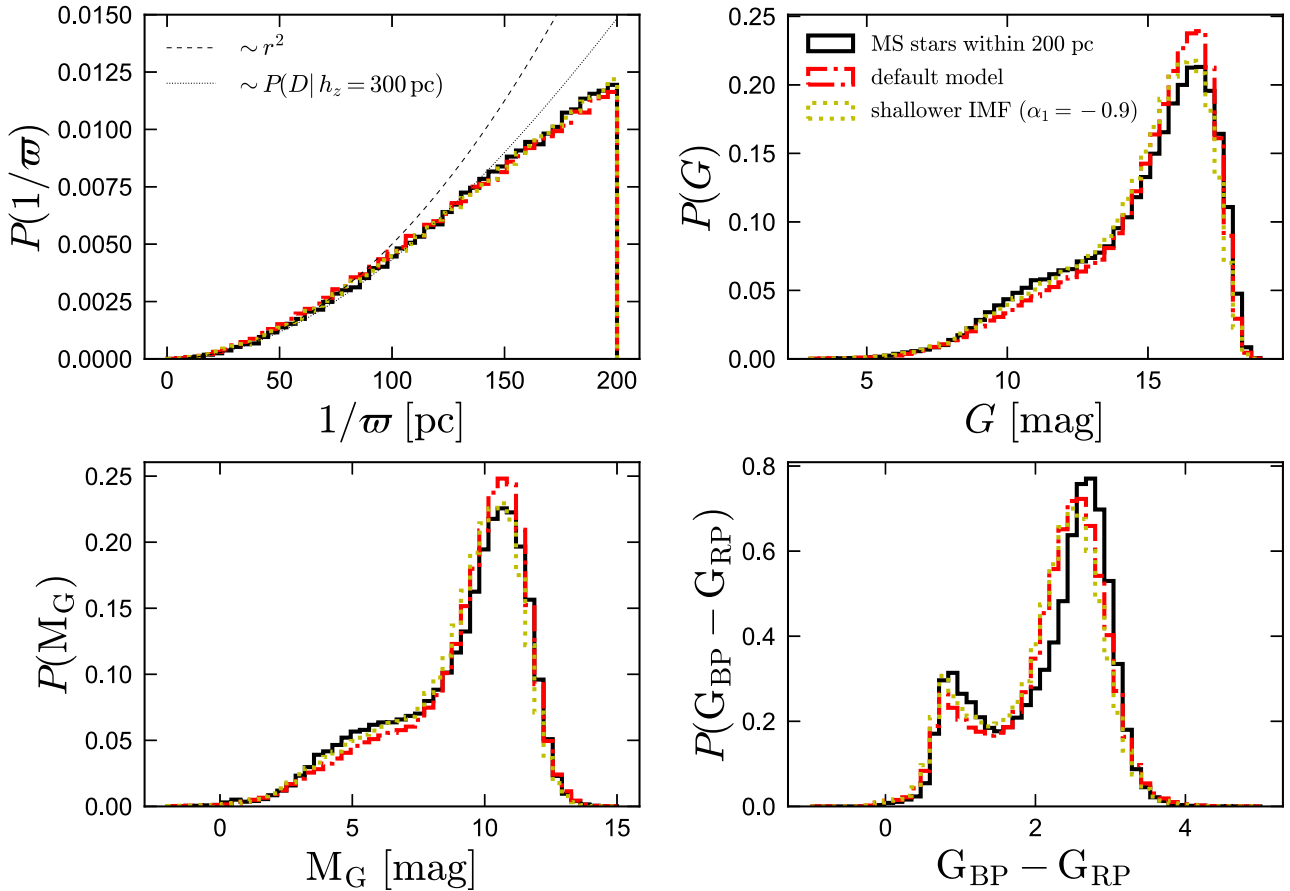
Figure E1. Validation of the single-star selection function and underlying Galactic model. Black histograms show distributions of distance, apparent magnitude, absolute magnitude, and colour for all stars (not just binaries) within 200 pc that pass the quality cuts of the binary catalogue. Black line shows the distributions predicted by our model; i.e. assuming the same single-star selection function, IMF, Galactic scale height, star formation history, and metallicity distribution that we use when fitting the binary population. The reasonably good agreement with the observed distributions suggests that our empirical selection function and Galactic model are reasonable. Our fiducial model slightly overpredicts the number of faint stars. This tension can be resolved if a shallower IMF is assumed (with a logarithmic slope of -0.9 instead of -1.3 at $M < 0.5 M_\odot$; gold). Overall, we regard the agreement between the fiducial model and observations as quite satisfactory.

which is where incompleteness effects become significant for $M_G \approx 14$ (Fig. D1).

The agreement between model and data is reasonably good. However, the fiducial model predicts slightly too many faint stars and too few bright stars. The gold histograms shows that the agreement can be improved if a slightly shallower IMF is assumed, with a logarithmic slope of -0.9 (instead of -1.3, as assumed in the fiducial model) at $M < 0.5 M_\odot$.[10] This is consistent with the recent result from Sollima (2019), who used the *Gaia* nearby star sample to measure the IMF. For the sake of this work, we are agnostic of whether Fig. E1 indicates that the low-mass IMF in the solar neighbourhood is slightly shallower than that assumed in our fiducial model or points towards a systematic in some other aspect of the model (such as the MDF or stellar models). We note that changing the assumed IMF has very little effect on our inferred mass ratio distribution (Fig. B1).

We also note that the colour distributions predicted by our model do not exactly match the observed distribution for either choice of IMF. This could be due to reddening, the adopted MDF (as the subset of stars with spectroscopic metallicity measurements is not guaranteed to be an unbiased sample), or imperfect stellar models.

E2 Chance alignments

To validate our model for the contrast sensitivity as a function of angular separation, we predict the distributions of magnitude difference at fixed separation using our model and compare to the chance alignment catalogue described in Section 2. This is accomplished as follows. We draw masses, ages, distances, and metallicities for both 'components' of a chance alignment independently, and we draw angular separations assuming $P(\theta)d\theta \propto 2\pi\theta$. Just as for true binaries, we compute the selection function for each pair by multiplying the two single-star selection functions and the angular contrast sensitivity term. The results are shown in Fig. E2 σ_0. The agreement with the real chance alignment catalogue is good.

[10]This slope was found by fitting the IMF from the CMD, as described in appendix C of El-Badry, Weisz & Quataert (2017).
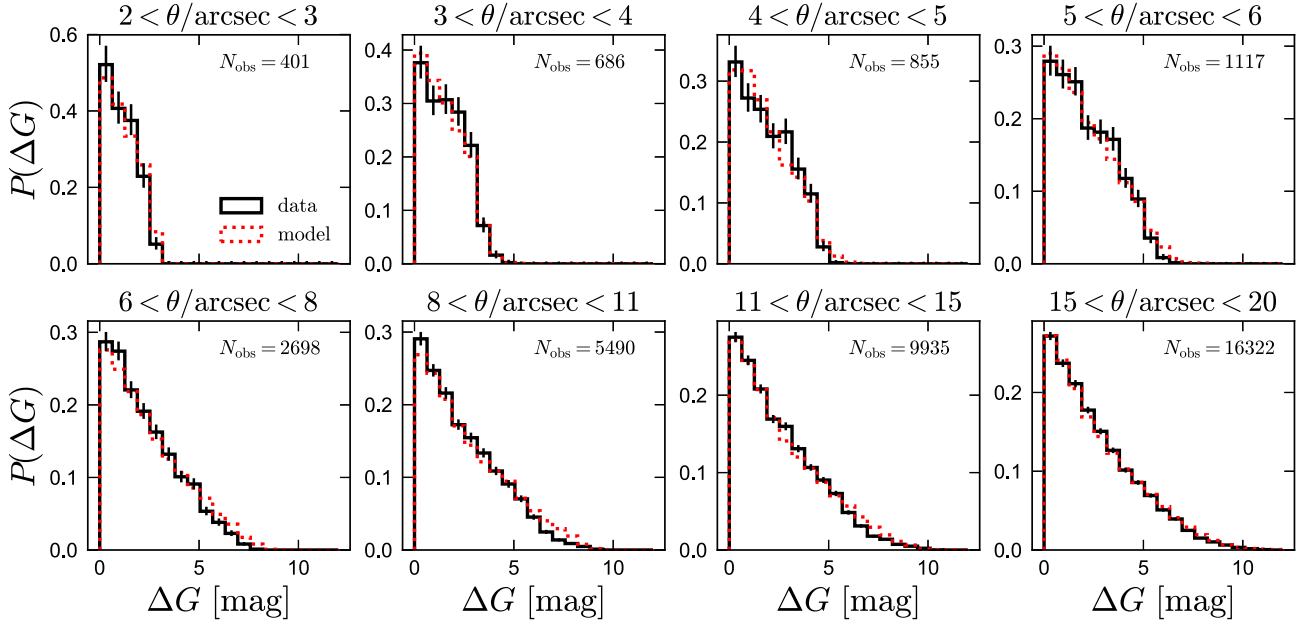
Figure E2. Distributions of apparent magnitude difference for our observed sample of chance alignments within 400 pc (black) and a simulated sample based on the same selection function and distributions of mass, age, metallicity, and distance assumed in modelling genuine binaries. Agreement between data and model validates the model assumptions, particularly the sensitivity to a companion as a function of angular separation and magnitude difference.

APPENDIX F: RV VARIABILITY

Although *Gaia* DR2 does not contain multi-epoch RV measurements, the published RV uncertainties contain information that can be exploited to detect a large fraction of RV-variable close binaries. The `radial_velocity_error` reported in the *Gaia* archive represents the uncertainty on the median of velocity measurements from several transits (see Katz et al. 2019). It is calculated as

$$\epsilon_{RV} = \left[\frac{\pi}{2 N_{obs}} \sigma_{RV}^2 + \sigma_0^2 \right]^{1/2}, \qquad (F1)$$

where $\sigma_0 = 0.11 \, \text{km s}^{-1}$ is a constant term that represents the minimum achievable RV uncertainty due to calibration issues, N_{obs} is the number of RV transits (`rv_nb_transits` in the *Gaia*

Figure F1. Median σ_{RV} (equation F2) for main-sequence stars within 200 pc that pass our quality cuts and `rv_nb_transits > 2`. Since a majority of stars are not intrinsically RV-variable, this represents the typical RV precision for stars of a given brightness and spectral type. Most stars that have σ_{RV} substantially larger than the median for their colour and magnitude are close binaries.
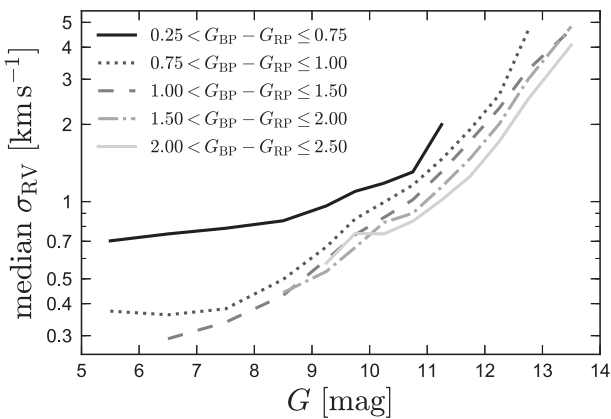
archive), and σ_{RV} is the standard deviation of the RVs measured in individual transits. The standard deviation of the measured transit RVs can thus be reconstructed as

$$\sigma_{RV} = \left[2 N_{obs} \left(\epsilon_{RV}^2 - \sigma_0^2 \right) / \pi \right]^{1/2}. \qquad (F2)$$

We expect σ_{RV} to be larger than usual if the variation in the true RV of a star between transits is large compared to the observational RV precision. This suggests that close binaries could be identified as stars with unusually large σ_{RV} for their colour and apparent magnitude.

To quantify this, we queried the *Gaia* archive for all stars within 200 pc that pass the quality cuts imposed on binary components and additionally have `rv_nb_transits > 2`. Selecting MS stars in bins of `bp_rp` colour and *G*-band magnitude, we calculated the median σ_{RV} in each bin (Fig. F1). As expected, the typical RV error increases with increasing *G* magnitude and is larger for bluer stars, which have weaker and broader absorption lines, at fixed magnitude.

We designate sources that have σ_{RV} larger than 2.5 times the median for their colour and apparent magnitude as likely close binaries. The factor of 2.5 is a practical choice to balance the number of false-positives and false-negatives. To assess the false-negative rate for this designation, we cross-matched *Gaia* DR2 with the catalogue of RV-variable MS SB1s identified by El-Badry et al. (2018b) using APOGEE spectra and calculated σ_{RV} for the subsample of that catalogue that passes our quality cuts. Among SB1s for which El-Badry et al. (2018b) found the RV to vary by at least 5 km s^{-1} between visits, 62 per cent are correctly identified as binaries based on the *Gaia* σ_{RV}. This fraction climbs to 82 per cent for sources whose APOGEE RVs varied by at least 15 km s^{-1}. On the other hand, the false-positive rate is relatively low: only 2.6 per cent of the stars classified by El-Badry et al. (2018b) as likely single are classified as binaries based on the *Gaia* σ_{RV}. This means that, although some true binaries will be missed and there will be some false-positives, the *Gaia* σ_{RV} can be used to obtain an estimate

of the close binary fraction in a population (where 'close' means $a \lesssim$ few au).

To assess whether the components of wide twins are more likely to have an unresolved close companion than the components of non-twins, we use the σ_{RV} distributions to estimate the fraction of components of twin ($\Delta G < 0.25$) and non-twin ($0.25 \leq \Delta G < 1$) wide binaries with an unresolved close companion. We only consider binaries with $100 < s/\mathrm{au} < 500$, since at wider separations, the number of 'excess' twins is subdominant relative to the underlying population (Fig. 10). We consider all components with masses in the range $0.5 < M/\mathrm{M}_\odot < 1.2$ that have *Gaia* RVs with `rv_nb_transits > 2` (i.e. we do not require both resolved components to have measured RVs). Among 546 eligible components of wide twins, 29 have σ_{RV} consistent with having unresolved close companion, implying a close companion fraction per wide binary component of

$$f_{\text{close companion, twins}} = 0.053 \pm 0.012. \tag{F3}$$

Of the 780 eligible components of non-twins, 76 have σ_{RV} consistent with having unresolved close companion. This implies

$$f_{\text{close companion, non-twins}} = 0.097 \pm 0.013. \tag{F4}$$

That is, the fraction of wide binary components that have a close unresolved companion is *higher* at the 2σ level for non-twins than for twins. This would seem to speak against a scenario in which the excess of wide twins is causally linked to hierarchical triples. Indeed, given the non-zero false-positive rate of the σ_{RV}-based close binary identification and the fact that 'excess' twins represent only about half of the population with $\Delta G < 0.25$, the data are consistent with the excess twins having no close companions at all.

APPENDIX G: FULL FITTING CONSTRAINTS

Constraints on fitting parameters for all bins of primary mass and separation are listed in Table G1. Error bars are 2σ (middle 95.4 per cent).

Table G1. Marginalized 2σ constraints on fitting parameters for all bins of primary mass and separation. For $0.1 < M_1/M_\odot < 0.4$, we fit a single power law, so $\gamma_{\mathrm{smallq}} = \gamma_{\mathrm{largeq}}$.

	$0.1 < M_1/M_\odot < 0.4$	$0.4 < M_1/M_\odot < 0.6$	$0.6 < M_1/M_\odot < 0.8$	$0.8 < M_1/M_\odot < 1.2$	$1.2 < M_1/M_\odot < 2.5$
F_{twin}					
$50 < s/\mathrm{au} < 350$	$0.189 \pm^{0.042}_{0.049}$	$0.112 \pm^{0.034}_{0.038}$	$0.059 \pm^{0.019}_{0.021}$	$0.101 \pm^{0.027}_{0.029}$	$0.086 \pm^{0.058}_{0.075}$
$350 < s/\mathrm{au} < 600$	$0.105 \pm^{0.036}_{0.035}$	$0.062 \pm^{0.018}_{0.020}$	$0.043 \pm^{0.016}_{0.019}$	$0.054 \pm^{0.014}_{0.018}$	$0.088 \pm^{0.037}_{0.047}$
$600 < s/\mathrm{au} < 1000$	$0.049 \pm^{0.032}_{0.039}$	$0.054 \pm^{0.017}_{0.020}$	$0.032 \pm^{0.019}_{0.019}$	$0.035 \pm^{0.015}_{0.016}$	$0.015 \pm^{0.024}_{0.034}$
$1000 < s/\mathrm{au} < 2500$	$0.023 \pm^{0.025}_{0.033}$	$0.049 \pm^{0.018}_{0.022}$	$0.024 \pm^{0.016}_{0.016}$	$0.013 \pm^{0.010}_{0.011}$	$0.002 \pm^{0.015}_{0.020}$
$2500 < s/\mathrm{au} < 5000$	$-0.001 \pm^{0.048}_{0.039}$	$0.047 \pm^{0.025}_{0.029}$	$0.018 \pm^{0.020}_{0.019}$	$0.023 \pm^{0.017}_{0.021}$	$-0.003 \pm^{0.023}_{0.023}$
$5000 < s/\mathrm{au} < 15\,000$	$0.005 \pm^{0.040}_{0.043}$	$0.043 \pm^{0.030}_{0.031}$	$0.014 \pm^{0.018}_{0.021}$	$0.007 \pm^{0.015}_{0.015}$	$0.003 \pm^{0.023}_{0.035}$
$15\,000 < s/\mathrm{au} < 50\,000$	$0.024 \pm^{0.094}_{0.120}$	$-0.009 \pm^{0.039}_{0.043}$	$-0.012 \pm^{0.028}_{0.032}$	$0.009 \pm^{0.024}_{0.032}$	$0.001 \pm^{0.037}_{0.057}$
q_{twin}					
$50 < s/\mathrm{au} < 350$	$0.953 \pm^{0.010}_{0.006}$	$0.959 \pm^{0.023}_{0.014}$	$0.969 \pm^{0.019}_{0.005}$	$0.962 \pm^{0.014}_{0.009}$	$0.950 \pm^{0.018}_{0.029}$
$350 < s/\mathrm{au} < 600$	$0.937 \pm^{0.006}_{0.011}$	$0.959 \pm^{0.017}_{0.009}$	$0.954 \pm^{0.021}_{0.016}$	$0.967 \pm^{0.027}_{0.010}$	$0.936 \pm^{0.006}_{0.018}$
$600 < s/\mathrm{au} < 1000$	$0.954 \pm^{0.022}_{0.025}$	$0.963 \pm^{0.025}_{0.012}$	$0.939 \pm^{0.008}_{0.018}$	$0.957 \pm^{0.023}_{0.028}$	$0.954 \pm^{0.022}_{0.040}$
$1000 < s/\mathrm{au} < 2500$	$0.975 \pm^{0.043}_{0.017}$	$0.963 \pm^{0.029}_{0.017}$	$0.941 \pm^{0.011}_{0.030}$	$0.959 \pm^{0.025}_{0.030}$	$0.957 \pm^{0.026}_{0.037}$
$2500 < s/\mathrm{au} < 5000$	$0.955 \pm^{0.025}_{0.039}$	$0.953 \pm^{0.022}_{0.019}$	$0.950 \pm^{0.019}_{0.029}$	$0.945 \pm^{0.014}_{0.039}$	$0.958 \pm^{0.027}_{0.035}$
$5000 < s/\mathrm{au} < 15\,000$	$0.959 \pm^{0.028}_{0.038}$	$0.951 \pm^{0.020}_{0.030}$	$0.965 \pm^{0.034}_{0.027}$	$0.970 \pm^{0.038}_{0.027}$	$0.960 \pm^{0.029}_{0.035}$
$15\,000 < s/\mathrm{au} < 50\,000$	$0.961 \pm^{0.030}_{0.031}$	$0.961 \pm^{0.029}_{0.033}$	$0.961 \pm^{0.029}_{0.032}$	$0.957 \pm^{0.026}_{0.038}$	$0.957 \pm^{0.026}_{0.039}$
γ_{largeq}					
$50 < s/\mathrm{au} < 350$	$0.52 \pm^{0.45}_{0.45}$	$0.17 \pm^{0.40}_{0.39}$	$-1.43 \pm^{0.48}_{0.51}$	$-1.22 \pm^{0.66}_{0.67}$	$-0.89 \pm^{0.90}_{1.00}$
$350 < s/\mathrm{au} < 600$	$-0.16 \pm^{0.43}_{0.43}$	$-0.01 \pm^{0.29}_{0.27}$	$-0.89 \pm^{0.35}_{0.33}$	$-1.83 \pm^{0.42}_{0.38}$	$-1.88 \pm^{0.54}_{0.55}$
$600 < s/\mathrm{au} < 1000$	$0.23 \pm^{0.41}_{0.45}$	$-0.44 \pm^{0.28}_{0.27}$	$-1.19 \pm^{0.30}_{0.30}$	$-1.43 \pm^{0.30}_{0.30}$	$-1.16 \pm^{0.33}_{0.28}$
$1000 < s/\mathrm{au} < 2500$	$0.39 \pm^{0.41}_{0.39}$	$-0.55 \pm^{0.28}_{0.26}$	$-0.90 \pm^{0.26}_{0.25}$	$-1.54 \pm^{0.22}_{0.22}$	$-1.55 \pm^{0.21}_{0.19}$
$2500 < s/\mathrm{au} < 5000$	$0.23 \pm^{0.52}_{0.52}$	$-0.43 \pm^{0.40}_{0.38}$	$-1.04 \pm^{0.36}_{0.34}$	$-1.52 \pm^{0.32}_{0.33}$	$-1.52 \pm^{0.29}_{0.27}$
$5000 < s/\mathrm{au} < 15\,000$	$0.22 \pm^{0.51}_{0.59}$	$-0.42 \pm^{0.44}_{0.42}$	$-0.86 \pm^{0.37}_{0.35}$	$-1.35 \pm^{0.33}_{0.31}$	$-1.31 \pm^{0.28}_{0.25}$
$15\,000 < s/\mathrm{au} < 50\,000$	$0.53 \pm^{0.86}_{0.84}$	$-0.43 \pm^{0.67}_{0.67}$	$-0.72 \pm^{0.65}_{0.63}$	$-1.39 \pm^{0.55}_{0.58}$	$-1.22 \pm^{0.46}_{0.43}$
γ_{smallq}					
$50 < s/\mathrm{au} < 350$		$0.33 \pm^{0.60}_{0.56}$	$0.20 \pm^{0.51}_{0.58}$	$0.13 \pm^{0.64}_{0.69}$	$0.13 \pm^{0.98}_{1.00}$
$350 < s/\mathrm{au} < 600$		$-0.37 \pm^{0.44}_{0.47}$	$0.24 \pm^{0.41}_{0.47}$	$0.40 \pm^{0.41}_{0.43}$	$-0.01 \pm^{0.79}_{0.86}$
$600 < s/\mathrm{au} < 1000$		$0.17 \pm^{0.44}_{0.46}$	$0.24 \pm^{0.33}_{0.36}$	$0.12 \pm^{0.26}_{0.26}$	$0.15 \pm^{0.55}_{0.72}$
$1000 < s/\mathrm{au} < 2500$		$0.53 \pm^{0.44}_{0.44}$	$0.06 \pm^{0.25}_{0.27}$	$0.12 \pm^{0.17}_{0.16}$	$0.56 \pm^{0.42}_{0.47}$
$2500 < s/\mathrm{au} < 5000$		$0.02 \pm^{0.56}_{0.58}$	$-0.15 \pm^{0.33}_{0.34}$	$-0.14 \pm^{0.21}_{0.21}$	$0.09 \pm^{0.36}_{0.43}$
$5000 < s/\mathrm{au} < 15\,000$		$0.13 \pm^{0.59}_{0.69}$	$0.12 \pm^{0.39}_{0.40}$	$-0.20 \pm^{0.21}_{0.22}$	$-0.18 \pm^{0.40}_{0.40}$
$15\,000 < s/\mathrm{au} < 50\,000$		$-0.02 \pm^{0.84}_{0.88}$	$-0.13 \pm^{0.69}_{0.60}$	$-0.24 \pm^{0.42}_{0.41}$	$-0.09 \pm^{0.64}_{0.66}$
γ_s					
$50 < s/\mathrm{au} < 350$	$-1.36 \pm^{0.22}_{0.21}$	$-1.55 \pm^{0.18}_{0.20}$	$-1.22 \pm^{0.32}_{0.38}$	$-1.02 \pm^{0.40}_{0.46}$	$-0.84 \pm^{0.77}_{0.93}$
$350 < s/\mathrm{au} < 600$	$-1.10 \pm^{0.43}_{0.42}$	$-1.20 \pm^{0.29}_{0.27}$	$-0.89 \pm^{0.39}_{0.38}$	$-0.80 \pm^{0.47}_{0.43}$	$-1.28 \pm^{1.00}_{1.00}$
$600 < s/\mathrm{au} < 1000$	$-1.48 \pm^{0.48}_{0.46}$	$-1.16 \pm^{0.31}_{0.29}$	$-1.75 \pm^{0.34}_{0.37}$	$-1.13 \pm^{0.36}_{0.36}$	$-0.03 \pm^{0.74}_{0.78}$
$1000 < s/\mathrm{au} < 2500$	$-1.67 \pm^{0.23}_{0.26}$	$-1.56 \pm^{0.14}_{0.14}$	$-1.52 \pm^{0.15}_{0.16}$	$-1.47 \pm^{0.13}_{0.13}$	$-1.41 \pm^{0.24}_{0.26}$
$2500 < s/\mathrm{au} < 5000$	$-2.05 \pm^{0.51}_{0.46}$	$-1.71 \pm^{0.30}_{0.30}$	$-1.75 \pm^{0.30}_{0.29}$	$-1.67 \pm^{0.23}_{0.26}$	$-1.40 \pm^{0.42}_{0.45}$
$5000 < s/\mathrm{au} < 15\,000$	$-1.74 \pm^{0.34}_{0.36}$	$-1.74 \pm^{0.21}_{0.21}$	$-1.62 \pm^{0.21}_{0.21}$	$-1.84 \pm^{0.16}_{0.17}$	$-1.35 \pm^{0.28}_{0.27}$
$15\,000 < s/\mathrm{au} < 50\,000$	$-1.61 \pm^{0.90}_{0.83}$	$-1.68 \pm^{0.47}_{0.49}$	$-2.21 \pm^{0.54}_{0.50}$	$-2.00 \pm^{0.37}_{0.40}$	$-1.50 \pm^{0.66}_{0.63}$

This paper has been typeset from a TEX/LATEX file prepared by the author.