

A million binaries from Gaia eDR3: sample selection and validation of Gaia parallax uncertainties

Kareem El-Badry,^{1,2*} Hans-Walter Rix,² and Tyler M. Heintz³

¹*Department of Astronomy and Theoretical Astrophysics Center, University of California Berkeley, Berkeley, CA 94720, USA*

²*Max-Planck Institute for Astronomy, Königstuhl 17, D-69117 Heidelberg, Germany*

³*Department of Astronomy, Boston University, 725 Commonwealth Ave., Boston, MA 02215, US*

Accepted to MNRAS

ABSTRACT

We construct from *Gaia* eDR3 an extensive catalog of spatially resolved binary stars within ≈ 1 kpc of the Sun, with projected separations ranging from a few au to 1 pc. We estimate the probability that each pair is a chance alignment empirically, using the *Gaia* catalog itself to calculate the rate of chance alignments as a function of observables. The catalog contains 1.3 (1.1) million binaries with $>90\%$ ($>99\%$) probability of being bound, including 16,000 white dwarf – main sequence (WD+MS) binaries and 1,400 WD+WD binaries. We make the full catalog publicly available, as well as the queries and code to produce it. We then use this sample to calibrate the published *Gaia* DR3 parallax *uncertainties*, making use of the binary components’ near-identical parallaxes. We show that these uncertainties are generally reliable for faint stars ($G \gtrsim 18$), but are underestimated significantly for brighter stars. The underestimates are generally $\leq 30\%$ for isolated sources with well-behaved astrometry, but are larger (up to $\sim 80\%$) for apparently well-behaved sources with a companion within $\lesssim 4$ arcsec, and much larger for sources with poor astrometric fits. We provide an empirical fitting function to inflate published σ_{ϖ} values for isolated sources. The public catalog offers wide ranging follow-up opportunities: from calibrating spectroscopic surveys, to precisely constraining ages of field stars, to the masses and the initial-final mass relation of white dwarfs, to dynamically probing the Galactic tidal field.

Key words: binaries: visual – stars: evolution – methods: statistical – catalogues – parallaxes

1 INTRODUCTION

About half of all solar-type stars are members of binary systems, and a majority of these are so widely separated that the two components never interact (e.g. Moe & Di Stefano 2017). With orbital periods ranging from ~ 10 to $\sim 10^8$ years, most of these binaries can in some sense be viewed as clusters of two: the components formed from the same gas cloud and have orbited one another ever since. They thus have essentially the same age, initial composition, and distance, but generally different masses and occasionally different evolutionary phases. This makes wide binaries useful for calibrating stellar models as well as spectroscopic and astrometric surveys.

At angular separations greater than about one arcsecond, wide binaries are easily resolvable as two point sources. Distinguishing physically bound binary stars from chance alignments (“optical doubles”) has been a long-standing challenge for binary star astronomy. Indeed, the first systematic binary star catalog was constructed under the assumption that all close pairs were chance alignments (Herschel 1782), an assumption that was only shown to be incorrect two decades later (Herschel 1803).

For bright binaries at close angular separations, chance alignments can be excluded probabilistically. However, the contamination rate from chance alignments increases at wider separations and fainter magnitudes. Inclusion of proper motion data can aid the selection of

genuine binaries, in which the two stars have nearly identical proper motions (e.g. Luyten 1971, 1979; Salim & Gould 2003; Chanamé & Gould 2004; Dhital et al. 2010). Many wide binary searches have specifically targeted high-proper motion stars, which have fewer phase-space neighbors that can be mistaken for binary companions. If available, parallaxes and radial velocities are also useful for distinguishing binaries from chance alignments (e.g. Close et al. 1990; Andrews et al. 2017).

Prior to the *Gaia* mission (Gaia Collaboration et al. 2016), useful parallaxes for this purpose were only available for (relatively) small samples of nearby and bright stars. *Gaia* DR2 (Gaia Collaboration et al. 2018) dramatically expanded the sample of stars with well-measured parallaxes and proper motions, enabling the construction of unprecedentedly pure and extensive wide binary samples. El-Badry & Rix (2018, hereafter ER18) searched *Gaia* DR2 for pairs of stars within 200 pc of the Sun with parallaxes and proper motions consistent with being gravitationally bound, and projected separations of up to 50,000 AU (0.24 pc). Their catalog prioritized purity over completeness, and thus imposed relatively strict cuts on astrometric and photometric quality and signal-to-noise ratio (SNR). This resulted in a catalog of $\sim 55,000$ binaries with an estimated contamination rate of $\sim 0.1\%$.

Using the same basic strategy but less stringent cuts on astrometric SNR, Tian et al. (2020) extended the ER18 binary search to larger distances ($d < 4$ kpc) and wider separations ($s < 1$ pc). This produced a substantially larger sample of $\sim 800,000$ binary candidates,

* E-mail: kelbadry@berkeley.edu

but with a higher contamination rate: chance-alignments dominate their catalog at $s \gtrsim 20,000$ au, though higher-purity subsamples can be selected by imposing stricter cuts on astrometric SNR. Their catalog contained 325,000 binaries with $s < 20,000$ au; it is expected to be reasonably pure in this separation regime.

Another binary catalog was produced by [Hartman & Lépine \(2020\)](#), who combined *Gaia* DR2 astrometry with a catalog of high-proper motion stars not contained in *Gaia* DR2. In order to reduce contamination from chance alignments, they limited their search to binaries with proper motions larger than 40 mas yr^{-1} ; this translates to a distance limit of order 200 pc for typical stars in the Galactic disk, but to a larger search volume for stars on halo-like orbits with large tangential velocities. Their primary binary catalog contains $\sim 100,000$ binary candidates, with projected separations as large as 10 pc.

In this paper, we use *Gaia* eDR3 data ([Gaia Collaboration et al. 2020a](#)) to further expand the sample of known wide binaries. Compared to DR2, eDR3 astrometry is based on a $\sim 1.5 \times$ longer time baseline. This yields significant improvements in parallax and especially proper motion uncertainties. For example, the median uncertainties in `parallax` and `pmra` (one-dimensional proper motion) at $G = 18$ have respectively improved from 0.165 mas to 0.120 mas, and from $0.280 \text{ mas yr}^{-1}$ to $0.123 \text{ mas yr}^{-1}$ ([Lindegren et al. 2018, 2020b](#)). Improvements at the bright end are more significant, due to better handling of systematics; i.e., at $G = 13$, the median parallax uncertainty decreased from 0.029 mas to 0.015 mas. This improved astrometric precision allows us to distinguish bound binaries from chance alignments to larger distances and wider separations than was possible with DR2. Our approach is a compromise between the strategies adopted by [Tian et al. \(2020\)](#) and [ER18](#). Like [Tian et al. \(2020\)](#), we search out to wide separations and relatively low astrometric SNR, so that at wide separations and faint magnitudes, the full catalog is dominated by chance alignments. However, we also empirically estimate and assign each binary a probability that it is a chance alignment, making it straightforward to select pure subsets of the catalog.

The remainder of the paper is organized as follows. In Section 2, we describe how our binary candidate sample is selected and cleaned. Section 3 describes how we quantify the contamination rate from chance alignments and estimate the probability that each binary candidate is bound. Section 4 details basic properties of the catalog and a cross-match with the LAMOST survey. In Section 5, we use the binary sample to validate the *Gaia* eDR3 parallax uncertainties. We summarize and discuss our results in Section 6. The public catalog is described in Section 7. Details about the calculation of chance-alignment probabilities are provided in the Appendix.

2 SAMPLE SELECTION

To reduce contamination from chance alignments, we limit our sample to pairs in which both components have moderately precise astrometry. We retrieved from the *Gaia* archive all sources with parallaxes greater than 1 mas (corresponding to a nominal distance limit of 1 kpc, which in reality is blurred due to parallax errors), fractional parallax uncertainties less than 20%, absolute parallax uncertainties less than 2 mas, and non-missing G -band magnitudes. This was achieved with the following ADQL query:

```
select *
from gaiaedr3.gaia_source
where parallax > 1
```

```
and parallax_over_error > 5
and parallax_error < 2
and phot_g_mean_mag is not null
```

The query returns a total of $N = 64,407,853$ sources, corresponding to $N(N+1)/2 \approx 2 \times 10^{15}$ possible pairs. Of these, we consider as initial binary candidates all pairs that satisfy the following:

- *Projected separation less than 1 parsec*: the angular separation between the two stars, θ , must satisfy

$$\frac{\theta}{\text{arcsec}} \leq 206.265 \times \frac{\varpi_1}{\text{mas}}, \quad (1)$$

where ϖ_1 is the parallax of the star with the brighter G magnitude. The maximum search radius of 1 pc (corresponding to an orbital period of $\sim 10^8$ years) is chosen because a vanishingly small number of bound binaries are expected to exist at separations wider than this, where the Galactic tidal field becomes comparable to the gravitational attraction of the two stars (e.g. [Binney & Tremaine 2008](#)). The separation beyond which the Galactic tidal field dominates a binary's internal acceleration is called the Jacobi radius. In the Solar neighborhood, it is given by $r_J \approx 1.35 \text{ pc} \times (M_{\text{tot}}/M_{\odot})$, where M_{tot} is the total mass of the binary ([Jiang & Tremaine 2010](#)). At separations slightly below r_J , binaries are efficiently disrupted by gravitational perturbations from objects such as other stars and molecular clouds (e.g. [Weinberg et al. 1987](#)).

- *Parallaxes consistent within 3 (or 6) sigma*: the parallaxes of the two components, ϖ_1 and ϖ_2 , must satisfy

$$|\varpi_1 - \varpi_2| < b \sqrt{\sigma_{\varpi,1}^2 + \sigma_{\varpi,2}^2}, \quad (2)$$

where $\sigma_{\varpi,i}$ is the parallax uncertainty of the i -th component, and $b = 3$ for pairs with $\theta > 4$ arcsec, or $b = 6$ for pairs with $\theta < 4$ arcsec. The less stringent cut at $\theta < 4$ arcsec is adopted because the chance alignment rate there is low and parallax uncertainties are significantly underestimated at close angular separations (Section 5).

- *Proper motions consistent with a Keplerian orbit*: The two stars in a wide binary will have proper motions that are similar, but, due to orbital motion, not identical. We require

$$\Delta\mu \leq \Delta\mu_{\text{orbit}} + 2\sigma_{\Delta\mu}, \quad (3)$$

where $\Delta\mu$ is the observed scalar proper motion difference, $\sigma_{\Delta\mu}$ its uncertainty, and $\Delta\mu_{\text{orbit}}$ the maximum proper motion difference expected due to orbital motion. The first two quantities are calculated as

$$\Delta\mu = \left[(\mu_{\alpha,1}^* - \mu_{\alpha,2}^*)^2 + (\mu_{\delta,1} - \mu_{\delta,2})^2 \right]^{1/2}, \quad (4)$$

$$\sigma_{\Delta\mu} = \frac{1}{\Delta\mu} \left[\left(\sigma_{\mu_{\alpha,1}^*}^2 + \sigma_{\mu_{\alpha,2}^*}^2 \right) \Delta\mu_{\alpha}^2 + \left(\sigma_{\mu_{\delta,1}}^2 + \sigma_{\mu_{\delta,2}}^2 \right) \Delta\mu_{\delta}^2 \right]^{1/2}, \quad (5)$$

where $\Delta\mu_{\alpha}^2 = (\mu_{\alpha,1}^* - \mu_{\alpha,2}^*)^2$ and $\Delta\mu_{\delta}^2 = (\mu_{\delta,1} - \mu_{\delta,2})^2$. Here $\mu_{\alpha,i}^* \equiv \mu_{\alpha,i} \cos \delta_i$, α and δ denote right ascension and declination, and μ_{α} and μ_{δ} , the proper motion in the right ascension and declination directions. Following [ER18](#), we take

$$\Delta\mu_{\text{orbit}} = 0.44 \text{ mas yr}^{-1} \times \left(\frac{\varpi}{\text{mas}} \right)^{3/2} \left(\frac{\theta}{\text{arcsec}} \right)^{-1/2}, \quad (6)$$

which is the maximum proper motion difference expected for a circular orbit of total mass $5M_{\odot}$; it corresponds to a projected physical velocity difference

$$\Delta V_{\text{orb}} = 2.1 \text{ km s}^{-1} \times \left(\frac{s}{1000 \text{ au}} \right)^{-1/2}, \quad (7)$$

where $s = 1000 \text{ au} \times (\theta/\text{arcsec}) (\varpi/\text{mas})^{-1}$ is the projected separation. This quantity is of course not equal to the full 3D separation, or to the semimajor axis, a . For randomly oriented orbits with a realistic eccentricity distribution, s and a usually agree within a factor of 2 (see ER18, their Appendix B). For our purposes, it is generally an acceptable approximation to assume $a \sim s$. We note that the cut on proper motion difference removes from the sample a significant fraction of unresolved hierarchical triples and higher-order multiples (see Section 4.4).

We apply the cuts on projected separation, parallax difference, and proper motion difference to all possible pairs. The projected separation cut reduces the list of possible binaries to 10^{10} ; additionally requiring consistent parallaxes and proper motions reduces it to 10^8 initial candidate pairs. A random sample of 1% of these is plotted in the top panel of Figure 1, and their projected separation distribution is shown in gray in Figure 2. A large majority of initial candidate pairs are chance alignments, not genuine binaries. This is evident both from the fact that many candidate pairs are in the Galactic bulge, LMC, and SMC, and from the fact that their separation distribution peaks at the widest separations (Figure 2), where true binaries are rare.

Stars in the bulge, LMC, and SMC would ideally be excluded by the requirement of $\varpi > 1$ and $\varpi/\sigma_\varpi > 5$ in the initial query, but a significant fraction of sources in crowded fields have spurious parallaxes (e.g. Fabricius et al. 2020). Most of these spurious background pairs can be excluded by imposing astrometric quality cuts – for example, we find that adding the requirement of `astrometric_sigma5d_max` < 1 to our initial query reduces the number of initial candidates by a factor of 10 while only removing a minority of genuine binaries¹. We opted against applying such cuts because they do remove some real binaries, and we find that chance alignments with spurious parallaxes can be efficiently filtered out by the cleaning described in Section 2.1.

In addition to chance alignments of background sources with spurious parallaxes, our initial selection also efficiently selects members of star clusters and moving groups, about 100 of which can be seen in the top panel of Figure 1. These are not spurious, in the sense that they really do contain many pairs of stars within our search volume that are close in phase space and in some cases mutually bound (e.g. Oh et al. 2017). However, most of them are not binaries and will become unbound when the clusters dissolve.

2.1 Cleaning clusters, background pairs, and triples

We clean the list of initial binary candidates in several passes. First, beginning with all the sources returned by our initial ADQL query, we count for each source the number of phase-space neighbors that are brighter than $G = 18$ and consistent with the size and velocity dispersion of a typical cluster. We define neighboring sources as those that satisfy the following:

- Projected separation less than 5 pc; i.e., $\theta \leq 17.19 \text{ arcmin} \times (\varpi/\text{mas})$.
- Proper motions within 5 km s^{-1} ; this translates to a proper motion difference $\Delta\mu \leq 1.05 \text{ mas yr}^{-1} \times (\varpi/\text{mas})$, with a 2σ tolerance.

¹ `astrometric_sigma5d_max` is the longest principal axis in the 5-dimensional error ellipsoid, in mas. A large value indicates that at least one of the astrometric parameters is poorly determined in the 5-parameter solution. More information on this and other *Gaia* flags can be found in the [Gaia eDR3 data model](#).

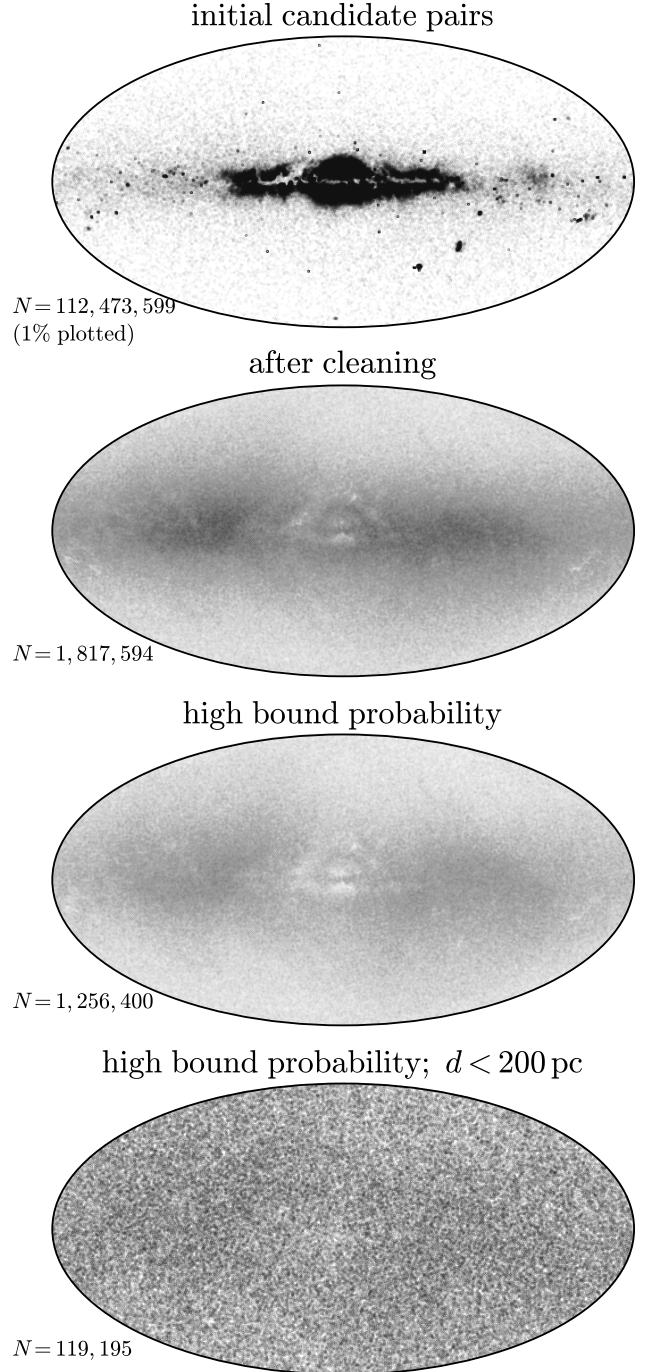


Figure 1. Sky distribution of binary candidates in Galactic coordinates. Top panel shows a random subset of all pairs with consistent parallaxes and proper motions. The vast majority of these are not genuine binaries. Many clusters are visible, as well as the Magellanic clouds and inner Galaxy. These are beyond the nominal 1 kpc search limit but enter the dataset due to spurious parallaxes. 2nd panel shows the sky distribution after clusters, moving groups, and resolved triples have been filtered (Section 2.1). This removes the majority of spurious background pairs. 3rd panel shows the pairs with $\mathcal{R} < 0.1$ (corresponding approximately to 90% bound probability; see Section 3.2). Some structure remains, primarily tracing dust. Bottom panel shows sources within 200 pc, which are distributed almost uniformly on the sky.




Figure 2. Separation distribution of binary candidates at three stages of the selection procedure, corresponding to the top three panels of Figure 1. Most initial candidates (gray) are chance alignments, which dominate at $s \gtrsim 3,000$ au. Cleaning resolved triples, clusters, and moving groups removes a large fraction of chance alignments (black), but they still dominate at $s \gtrsim 30,000$ au. The separation distribution of binaries with high bound probability ($\gtrsim 90\%$; red; see Section 3.2) falls off steeply at wide separations.

- Parallaxes consistent within 2 sigma; i.e., $\Delta\varpi \leq 2\sqrt{\sigma_{\varpi,1}^2 + \sigma_{\varpi,2}^2}$.

We remove from our binary candidate list all pairs in which either component has more than 30 neighbors as defined above. Only 6.5 million of the 64 million sources in the search sample have more than 30 neighbors, and inspection reveals that a majority of these are not in the search volume at all, but rather are spurious sources in the Galactic bulge, LMC and SMC. Removing candidates containing these sources shrinks the candidate list from 112,473,599 pairs to 2,881,543 and removes most of the obvious structure seen in the top panel of Figure 1.

Next, we remove all overlapping pairs. That is, if either component of a binary candidate is a member of another binary candidate, we remove both pairs. This removes genuine resolved triples, which are efficiently identified by our initial search (e.g. [Perpinya-Vallès et al. 2019](#)) and are not rare ([Tokovinin 2014](#)). It also removes some additional chance alignments. This cut shrinks the candidate sample from 2,881,543 to 1,918,362.

Finally, we search for members of small clusters or moving groups not removed in the first pass. Adopting the phase-space coordinates of the brighter component of each pair as representing the pair, we count the number of neighboring pairs for each candidate, defined using the same three criteria used when counting neighboring sources (without any magnitude cut). We reject all candidates that have more than 1 neighboring pair, shrinking the sample from 1,918,362 to 1,817,594.

It is important to note that some real binaries will be removed during the filtering of resolved triples, clusters, and moving groups. In regions of high stellar density, a distant tertiary candidate that is really a chance alignment can cause a genuine binary to be rejected as a triple. Similarly, some bound binaries *do* exist within clusters, and these will all be rejected. An upper limit of $\sim 15\%$ can be set on the fraction of true binaries lost during cleaning by comparing the gray and black histograms in Figure 2 at close separations; this is an upper limit because some of the pairs removed at close separations are genuinely members of resolved triples, moving groups, and clusters.

The binary candidate sample after removal of resolved triples and suspected cluster members is shown in the 2nd panel of Figure 1 and with the black histogram in Figure 2. This is the catalog published along with this work; it contains both high-confidence binaries and pairs that are very likely chance alignments. The contamination rate from chance alignments increases rapidly with separation (Section 3), while the true binary separation distribution decreases monotonically over the range of separations to which we are sensitive (e.g., [ER18](#)). The separation at which the number of binary candidates per dex of separation begins to increase ($\log(s/\text{au}) \approx 4.5$ for the black histogram in Figure 2) thus marks the separation at which chance alignments begin to dominate the full sample. That is, when considering the full catalog, a majority of binary candidates with $s \gtrsim 30,000$ au are chance alignments. It is, however, possible to select subsets of the catalog that are free of chance alignments out to much wider separations; see Section 3.

We define the “primary” and “secondary” components, denoted “1” and “2”, as the component with the brighter and fainter G magnitude, respectively. Both components are on the main sequence in a majority of binaries; in these cases, the “primary” is generally also the more massive component. For binaries containing white dwarfs, the secondary will often be more massive than the primary.

3 CHANCE ALIGNMENTS

The contamination rate from chance alignments depends on a variety of factors, including angular separation, parallax and proper motion, their respective uncertainties, and the local source density. We use two complementary approaches to constrain the chance alignment rate for different subsets of the catalog.

First, we repeat our binary selection procedure on the GeDR3mock catalog produced by [Rybizki et al. \(2020\)](#). This catalog is built on a realization of the Besançon model of the Milky Way ([Robin et al. 2003](#)) produced with *Galaxia* ([Sharma et al. 2011](#)). It contains no binaries, so by construction all binary candidates selected from the mock catalog are chance alignments. It does, however, contain a realistic population of open clusters, a variety of stellar populations, an approximation of the *Gaia* eDR3 selection function, and realistic astrometric uncertainties. We repeat the full binary selection process described in Section 2 on the mock catalog, including filtering of clusters and resolved triples. We remove pairs with angular separations $\theta < 0.5$ arcsec by hand, because the *Gaia* eDR3 sensitivity drops precipitously at closer separations ([Fabricius et al. 2020](#)), and this is not accounted for in the mock catalog. The separation distribution of candidates selected from the mock catalog is shown in the left panel of Figure 3.

Second, we produce an empirical chance alignment sample based on the actual *Gaia* eDR3 catalog, following the method introduced by [Lépine & Bongiorno \(2007\)](#). Prior to selecting potential binary companions to each star, we artificially shift it from its true position by ≈ 0.5 degrees, increasing its reported RA by $0.5 \text{ sec}(\delta)$ degrees. We then repeat the binary search, treating each star’s shifted coordinates as its true coordinates when searching for possible companions. This process avoids selecting real binaries, since stars are shifted away from their true companions, but preserves chance alignment statistics, because the source density within our 1 kpc search volume does not vary much on 0.5 degree scales. Copying and shifting the catalog effectively doubles the number of possible chance alignments, increasing the total number of pairs from $N(N+1)/2$ to $N(N+1)$, so we retain members of the shifted chance alignment catalog with 50% probability. We again remove pairs with $\theta < 0.5$ arcsec. The




Figure 3. Two methods for estimating the contamination rate from chance alignments. Left panel compares the separation distribution of binary candidates (black) to that of candidates selected from the eDR3 mock catalog (Rybizki et al. 2020; cyan), which contains no true binaries. Right panel compares the same binary candidates (black) to the separation distribution of candidates when stars are artificially shifted 0.5 degrees from their true positions when selecting candidate companions (orange; by construction, these are not true binaries). Both methods show that chance alignments dominate the full catalog at wide separations ($s \gtrsim 30,000$ au, where the separation distribution begins to increase), but true binaries dominate at closer separations.

separation distribution of chance alignments produced in this way is shown in the right panel of Figure 3.

The two methods predict similar chance alignment rates. At wide separations ($\log(s/\text{au}) \gtrsim 4.5$ for the full catalog), the number of chance alignments is similar to the number of pairs in the binary candidates catalog. In this regime, most candidates are chance alignments. Not surprisingly, chance alignments begin to dominate at the separation where the binary candidate separation distribution begins to increase. The number of chance alignments per logarithmic separation interval increases steeply with separation: there are 100 times more chance alignments with $5 < \log(s/\text{au}) < 5.1$ than there are with $4 < \log(s/\text{au}) < 4.1$. This is a consequence of the larger available area for background stars to be found in at wider separations. At fixed distance, the area in which chance alignments can appear scales as $\sim 2\pi s ds$. At wide separations, the chance alignment distributions thus scale as $dN/d \log s \sim s^2$. This scaling does not hold exactly, in part due to small-scale clustering and in part due to the less stringent parallax consistency required at small angular separations, but it provides a good approximation to the chance alignment rate at wide separations ($\log(s/\text{au}) \gtrsim 4$).

We use the “shifted” chance alignment catalog in the rest of our analysis, because we find that it reproduces the separation distribution of binary candidates in the large-separation limit somewhat more reliably than the mock catalog.

3.1 Chance alignment rate for subsets of the catalog

Although chance alignments dominate the full sample at $s \gtrsim 30,000$ au, it is possible to select subsets of the catalog that are free of chance alignments out to wider separations. This is illustrated in Figure 4, which shows the separation distributions of various subsets of the binary candidate and the shifted chance alignment catalogs. As expected, the chance alignment rate at fixed separation is lower for samples with high galactic latitude, bright component stars, small fractional parallax errors, large space velocities, or large proper motions. Nevertheless, it is challenging to select any subset

that remains pure beyond $s \sim 10^5$ au (~ 0.5 pc) without also dramatically reducing the sample size. We make the catalog of shifted chance alignments publicly available in order to facilitate estimation of the chance alignment rate in various subsamples of the catalog.

3.2 Estimating chance alignment probabilities

As illustrated in Figure 4, chance alignments and true binaries are found in different, but overlapping, regions of parameter space. We estimate the probability that a particular binary candidate is bound by comparing, at its location in parameter space, the local density of binary candidates and that of chance alignments from the shifted catalog. The process is described in detail in Appendix A. The “densities” are evaluated in a seven-dimensional space using a Gaussian kernel density estimate (KDE). The dimensions are (1) angular separation, (2) distance, (3) parallax difference uncertainty, (4) local sky density, (5) tangential velocity, (6) parallax difference over error, and (7) proper motion difference over error. We re-scale these quantities to all have similar, order-unity dynamic range before fitting KDEs to both the binary candidate and the chance alignment distributions.

We denote the KDE-estimated density of chance alignments at a point \vec{x} in the 7-dimensional parameter space as $\mathcal{N}_{\text{chance align}}(\vec{x})$, and that of binary candidates as $\mathcal{N}_{\text{candidates}}(\vec{x})$. The latter quantity is expected to be the sum of the chance-alignment and true binary densities. We then calculate the ratio of these two quantities,

$$\mathcal{R}(\vec{x}) = \mathcal{N}_{\text{chance align}}(\vec{x}) / \mathcal{N}_{\text{candidates}}(\vec{x}). \quad (8)$$

This ratio approximately represents the probability that a binary candidate at position \vec{x} is a chance alignment, so selecting only candidates with small \mathcal{R} is an efficient method for eliminating chance alignments. \mathcal{R} is not strictly a probability – for example, it is not strictly less than one (Figure 5) – but it is a serviceable approximation for one. We calculate \mathcal{R} values for all members of the binary candidate and chance alignment catalogs.

Figure 5 (left) shows the distribution of \mathcal{R} values for both catalogs. There is a narrow population of binary candidates with \mathcal{R} near zero;

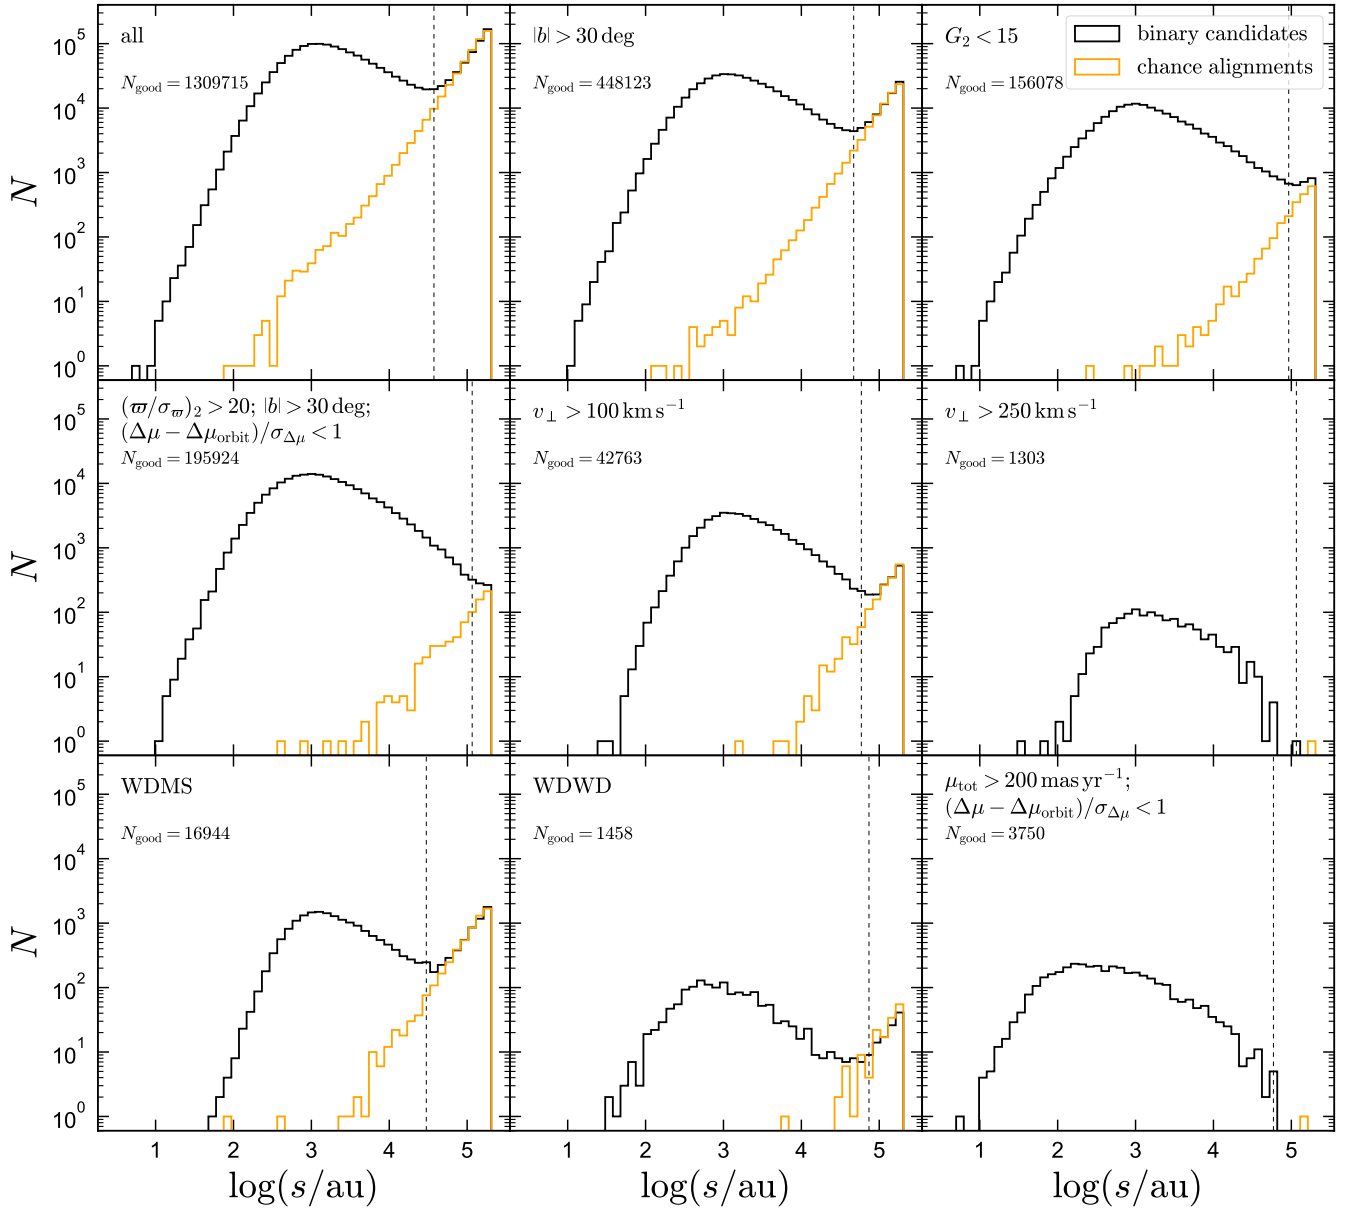


Figure 4. Contamination rate from chance alignments in various subsets of the catalog. Black histograms show binary candidates. Orange histograms show a catalog of chance alignments constructed by shifting stars 0.5 degrees when selecting candidate binaries. It contains no real binaries, but it has similar chance alignment statistics to the true binary candidates catalog. For the set of cuts illustrated in each panel, the vertical dashed line shows the widest separation at which there are more than $2\times$ the number of binary candidates as chance alignments. N_{good} is the number of binary candidates in that subset closer than this separation. Panels show the full catalog, high galactic latitudes, binaries with both components brighter than $G = 15$, precise parallaxes at high Galactic latitudes, high tangential velocities (e.g. halo-like orbits), WD+MS and WD+WD binaries, and high-proper motion pairs.

these are objects that have a high probability of being bound. There is a second population of candidates with $\mathcal{R} \sim 1$; these objects are likely chance alignments. The separation distributions of binary candidates with \mathcal{R} below several thresholds are shown in the right panel of Figure 5. As expected, the vast majority of binary candidates with close separations ($\log(s/\text{au}) \lesssim 4$) have low \mathcal{R} values, indicating a high bound probability. At wider separations, the separation distribution of high-probability binaries falls off precipitously. We emphasize that this drop-off is steeper than that of the separation distribution of all binaries, because for a low \mathcal{R} threshold, more true binaries will be

excluded from the high-confidence sample than chance alignments will be included in it.

We validate the use of \mathcal{R} as a proxy for the chance-alignment probability in Figure 6, which compares the *Gaia* DR2 radial velocities (Sartoretti et al. 2018; Katz et al. 2019) of the two components of candidates in which both stars have measured RVs. We only plot candidates in which both components have $\sigma_{\text{RV}} < 10 \text{ km s}^{-1}$ and the separation is wider than 50,000 au, where the full catalog is dominated by chance alignments. One expects the RVs of the two components to be similar for genuine wide binaries. For chance alignments, the RVs of the two components should be drawn from a broad

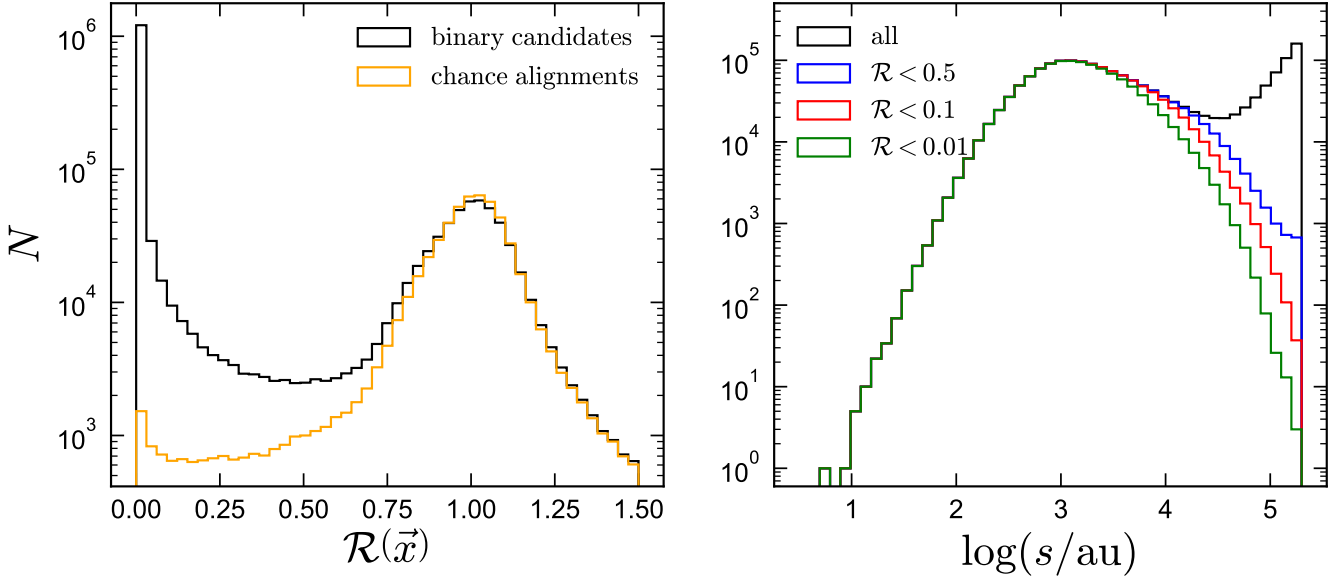


Figure 5. Left: ratio of the local density of chance alignments to binary candidates (Equation 8), for all pairs in the binary candidate sample (black) and the shifted chance alignment sample (orange). Likely chance alignments have $\mathcal{R} \sim 1$; objects with high probability of being bound have $\mathcal{R} \sim 0$. Right: separation distribution of binary candidates below different \mathcal{R} thresholds. The chance-alignment rate increases at wide separations, so the separation distribution of low- \mathcal{R} candidates falls off more steeply for progressively smaller \mathcal{R} thresholds.

distribution with width comparable to the local velocity dispersion of the Galactic disk, and thus they will generally be inconsistent. For the full sample, there are indeed plenty of pairs with obviously inconsistent RVs, as chance alignments dominate at wide separations, even for bright pairs (e.g. Figure 4). Red points in Figure 6 show binaries with $\mathcal{R} < 0.1$. As expected, these pairs all have RVs close to the one-to-one line and are likely all bound. We stress that RVs are not used in creating the catalog or calculating \mathcal{R} values, so this result bolsters our confidence in the chance alignment ratios calculated from the shifted catalog.

We note that only bright stars ($G \lesssim 13.5$) had RVs published in *Gaia* DR2 (with no new RVs added in eDR3), so the fraction of chance alignments among binaries where both components have RVs is lower than in the full sample at the same separation. RVs for fainter stars are compared in Section 4.5, where we take RVs from the LAMOST survey.

In the rest of the paper, we define the “high bound probability” or “high confidence” subset of the catalog as the subset with $\mathcal{R} < 0.1$; this corresponds approximately to $> 90\%$ probability of being bound. This does *not* mean that 10% of the pairs in this subset are chance alignments: most candidates in it have $\mathcal{R} \ll 0.1$ (Figure 5). Interpreting \mathcal{R} as the probability that a given pair is a chance alignment, we estimate that 4,600 of the 1.26 million candidates with $\mathcal{R} < 0.1$ are chance alignments (0.4%). For $\mathcal{R} < 0.01$, the same fraction is 870 out of 1.15 million (0.08%). We make the full candidate catalog and \mathcal{R} values available, including pairs that are likely chance alignments.

4 BASIC PROPERTIES OF THE CATALOG

Basic properties of the binary catalog are shown in Figure 7 and listed in Table 1. Only the high-confidence pairs ($\mathcal{R} < 0.1$) are shown. Following ER18, we classify stars as white dwarfs (WDs) or main-sequence (MS) based on their position in the *Gaia* color-absolute

Classification	$N_{\text{candidates}}$	$N_{\mathcal{R} < 0.1}$	Description
MSMS	1,412,903	877,416	both MS
WDMS	22,563	16,156	one WD, one MS
WDWD	1,565	1,390	both WD
MS??	378,877	360,180	one MS, one no colors
WD??	646	547	one WD, one no colors
????	1,040	711	both no colors
Total	1,817,594	1,256,400	

Table 1. Contents of the binary candidate catalog. All stars with measured $G_{\text{BP}} - G_{\text{RP}}$ colors are classified as “MS” or “WD” depending whether they fall above or below the dashed lines in Figure 8. $N_{\text{candidates}}$ is the number of candidate binaries with separations up to 1 pc (black histogram in Figure 2); $N_{\mathcal{R} < 0.1}$ is the number with high bound probability (red histogram in Figure 2).

magnitude diagram (CMD): defining $M_G = G + 5 \log(\varpi/100)$, we classify as WDs objects with $M_G > 3.25(G_{\text{BP}} - G_{\text{RP}}) + 9.625$; all other stars with measured $G_{\text{BP}} - G_{\text{RP}}$ colors are classified as MS stars. Under the ansatz that the two components have the same distance, we use the (usually more precise) parallax of the primary, (the brighter star) for both components when calculating M_G .

The adopted WD/MS boundary in the CMD is shown with a dashed line in Figures 7 and 8. The boundary is not entirely unambiguous – particularly for the WD+MS binaries, there are a few objects near the boundary that may be misclassified – but a majority of objects do fall clearly on the WD or MS sequences. We note that the “MS” classification serves only to exclude WDs. The CMDs in Figure 7 show that while most non-WD stars are indeed on the main sequence, the “MS” class also includes some giants, subgiants, pre-main sequence stars, and brown dwarfs. The number of binary candidates and high-confidence binaries in each class is summarized in Table 1.

The data acquisition window for BP/RP spectra is 2.1×3.5 arcsec wide, preventing colors from being measured for most close pairs




Figure 6. *Gaia* RVS velocities of the components of binary candidates wider than 50,000 au. These are not used in constructing the catalog or calculating bound probabilities, but are useful for validation. Black points show all candidate binaries for which both components have $\sigma_{RV} < 10 \text{ km s}^{-1}$; red points show the subset of these with $\mathcal{R} < 0.1$, corresponding approximately to a 90% bound probability. These all fall close to the one-to-one line, suggesting that they are indeed bound.

(Arenou et al. 2018). The majority of sources with a comparably bright companion within 2 arcsec thus do not have measured $G_{BP} - G_{RP}$ colors and cannot be classified as WD or MS stars based on *Gaia* data alone. Components lacking a color measurement are denoted “??” in Table 1 and in the catalog. About 30% of all high-confidence binaries have one component with unknown color; 0.05% lack colors for both components. In 98% of cases where only one component has a color, it is the brighter component. In many cases where no *Gaia* color is available, colors from other surveys (e.g. Pan-STARRS) should be sufficient to distinguish WD and MS components.

The catalog contains 8.8×10^5 high-confidence MS+MS binaries, more than 16,000 high-confidence WD+MS binaries, and 1,390 high-confidence WD+WD binaries. Angular separations range from 0.2 arcsec to one degree. The peak of the angular separation distribution is at 1.2 arcsec. This is simply a result of the *Gaia* eDR3 angular resolution, since the intrinsic separation distribution falls off monotonically with increasing separation over all separations that are well-represented in the catalog (e.g. Duchêne & Kraus 2013). There are 271 pairs with separations between 0.2 and 0.4 arcsec, including 24 below 0.3 arcsec. Fabricius et al. (2020, their Figure 7) found some indication that the *Gaia* eDR3 catalog may contain spurious duplicated sources at separations below 0.4 arcsec (that is, single sources that were erroneously classified as two sources), so it is possible that a small fraction of the closest pairs in the catalog are spurious. We do not, however, find any increase in the angular separation distribution at close separations, as might be expected to arise from a population of duplicated sources.

The median magnitude of high-confidence primaries is $G = 15.2$, and that of secondaries is $G = 17.7$. Most WDs in the catalog are significantly fainter: the median magnitudes of primaries and

secondaries in the WD+WD sample are 19.1 and 19.8. For WD+MS binaries, the same values are 15.5 for the primaries (in most cases, the MS star) and 19.4 for the secondaries. The median parallax of the full high-confidence sample is 2.05 mas ($1/\varpi \approx 485$ pc); the median distances for WD+WD and WD+MS binaries are 148 and 212 parsecs. Because of their closer distances, the WD+WD and WD+MS binaries are distributed roughly uniformly on the sky. The MS+MS sample, which extends to larger distances, bears clear imprint of the stratification of the Galactic disk. Most of the binaries in the sample are part of the kinematic “disk” population, with a median tangential velocity $v_{\perp} = 4.74 \text{ km s}^{-1} \times (\mu_{\text{tot}}/\text{mas yr}^{-1}) (\varpi/\text{mas})^{-1} \approx 35 \text{ km s}^{-1}$. There is also evidence of a kinematic “halo” population with $v_{\perp} \gtrsim 200 \text{ km s}^{-1}$ that contains a few thousand binaries.

The WD+MS and WD+WD binaries are shown separately on the CMD in Figure 8. On top of the WD+WD binaries, we plot WD cooling tracks for carbon-oxygen cores with hydrogen atmospheres (Holberg & Bergeron 2006; Kowalski & Saumon 2006; Tremblay et al. 2011; Bergeron et al. 2011; Bédard et al. 2020). WDs cool as they age, moving from the upper left to the lower right of the CMD. Cooling ages are indicated with triangular symbols along the tracks, which mark intervals of 1 Gyr. The faintest WDs in the catalog have implied cooling ages of about 10 Gyr. 31 WD+WD binaries have one component that falls below the $1.2M_{\odot}$ cooling track; i.e., with a photometrically-implied mass $M > 1.2M_{\odot}$. For WD+MS binaries, 75 WDs fall below the $1.2M_{\odot}$ cooling track. Hydrogen-atmosphere cooling tracks are not appropriate for WDs with non-DA spectral types, so spectroscopic classification must be obtained before masses and ages of individual WDs can be inferred with high fidelity.

The catalog also contains about 10,000 high-confidence binaries in which the primary is a giant (about half these giants are in the red clump), including about 130 giant-giant binaries. These are all quite bright (both components have $G < 11$). They, along with the WD+MS binaries and the $\approx 13,000$ binaries in which one component is a subgiant, can serve as useful calibrators for stellar ages.

Massive stars are not well-represented in the catalog; there are 351 high-confidence primaries with $G_{BP} - G_{RP} < 1$ and $M_G < 0$. This cut corresponds roughly to $M \gtrsim 3M_{\odot}$, though extinction complicates the mapping between M_G and mass. 75% of the high-confidence MS+MS binaries have primaries with M_G between 9.5 and 3.8, corresponding approximately to $0.4 < M_1/M_{\odot} < 1.3$. For secondaries, the 75% range is $M_G = 11.3 - 6.4$, corresponding to $0.2 < M_2/M_{\odot} < 0.8$. The catalog also contains about 80 binaries in which one component is likely a brown dwarf. We identify these on the CMD as objects with $M_G > 16.5$ and $G - G_{RP} > 1.3$; they are all within 80 pc of the Sun and therefore are also found in the binary catalog produced by Gaia Collaboration et al. (2020b).

A small but noticeable fraction of sources, particularly secondaries, are scattered below the main sequence in the CMD, between the WDs and MS stars (bottom right panel of Figure 7). The majority of these sources are separated from a brighter companion by only a few arcsec; the most likely explanation for their anomalous CMD position is thus that they are MS stars with contaminated BP/RP photometry. The majority of such sources can be filtered out using cuts on `bp_rp_excess_factor` and `phot_bp/rp_n_blended_transits` (e.g. Riello et al. 2020), but we refrain from employing such cuts since they also remove a significant fraction of sources with acceptable photometry and astrometry. Some sources below the main sequence may also be stars with spurious parallaxes or biased colors, and a few are likely real astrophysical sources, primarily cataclysmic variables and detached but unresolved WD+MS binaries (e.g. Abrahams et al. 2020; Belokurov et al. 2020).

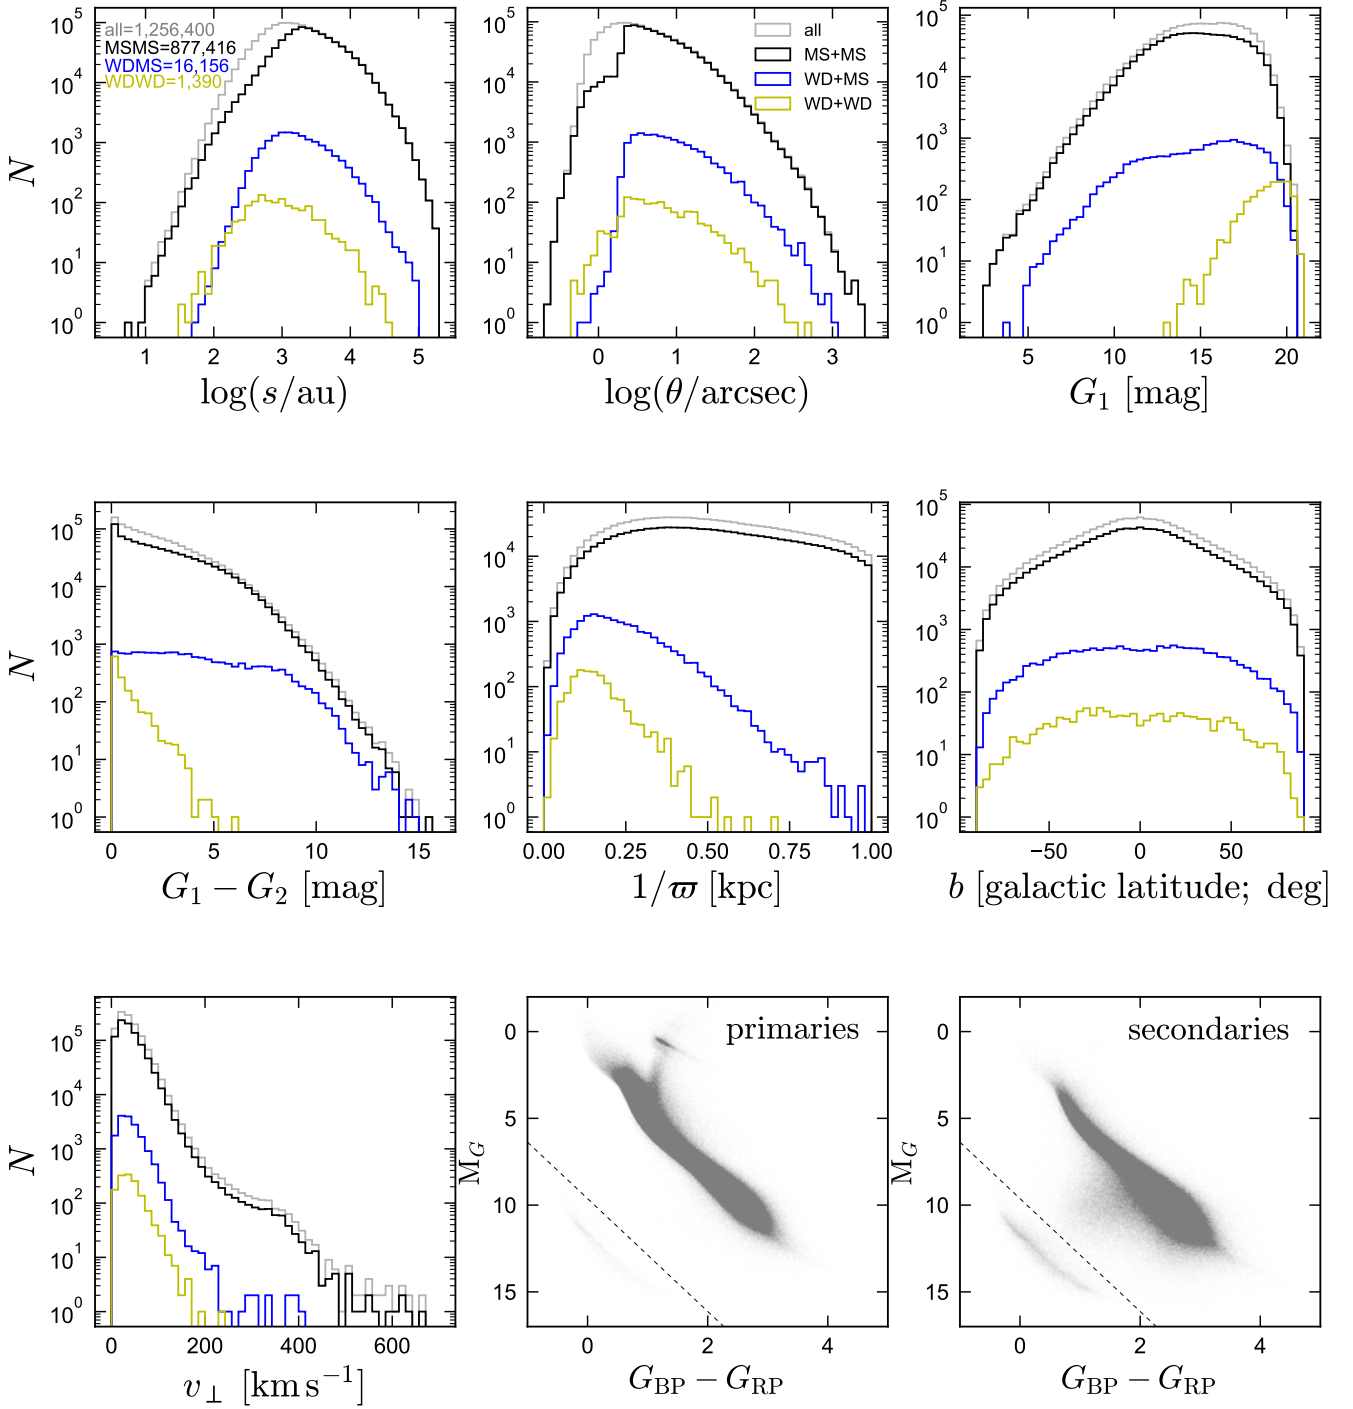


Figure 7. Basic properties of the high-bound probability binary sample: physical and angular separation, apparent magnitude of the primary, magnitude difference, distance, galactic latitude, plane-of-the-sky velocity, and color-magnitude diagrams. We designate stars as white dwarfs (WDs) or main sequence (MS; including giants) based on whether they fall below or above the dashed line in the bottom right panels. Black, blue, and yellow histograms show properties of MS+MS, WD+MS, and WD+WD binaries. We only classify binaries as MS+MS, WD+MS, or WD+WD if both components have a $G_{\text{BP}} - G_{\text{RP}}$ color. A majority of binaries with separation $\theta < 2$ arcsec do not have *Gaia* colors for a least one component and are therefore not classified; this accounts for most of the difference between the black and gray histograms (see also Table 1).

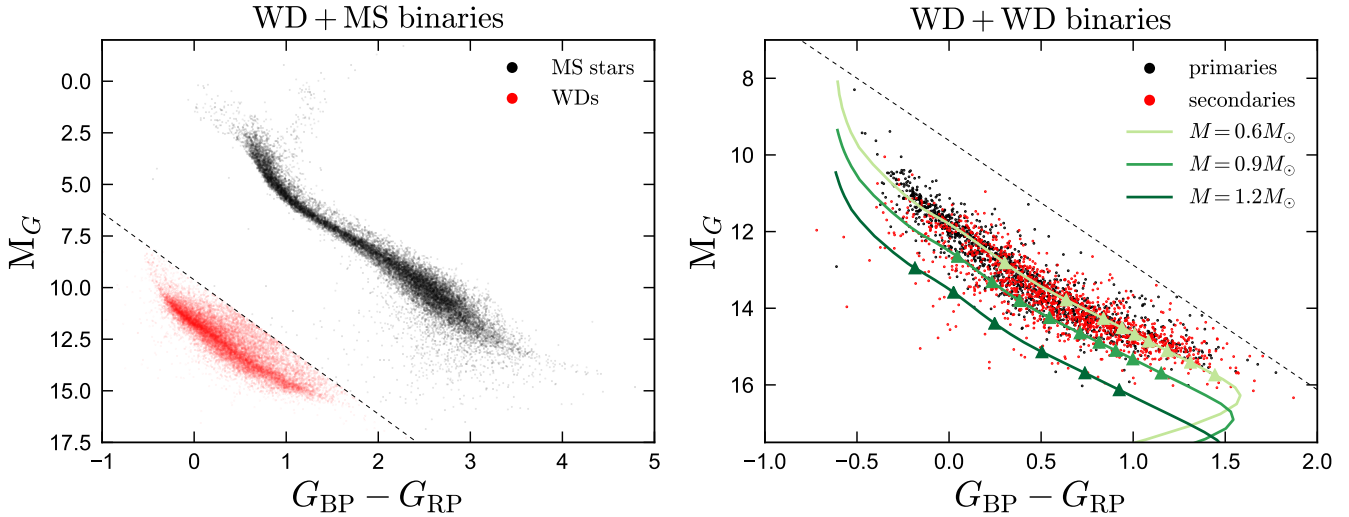


Figure 8. Color–absolute magnitude diagrams of high-confidence WD+MS and WD+WD binaries. Left: black points show MS components, which are usually the primaries; WD components are shown in red. Right: primary and secondary WDs are shown in black and red. We overplot cooling models for hydrogen-atmosphere WDs with masses between 0.6 and 1.2 M_{\odot} . Symbols mark 1 Gyr intervals of cooling age, with the first (leftmost) symbol at 1 Gyr.




Figure 9. Magnitude difference between the two stars in chance alignments (top) and bound wide binaries (bottom) as a function of angular separation. The sign of ΔG is randomized. Chance alignments are selected following the same procedure used for real binaries, but with the requirement that the two stars’ parallaxes be *inconsistent* rather than consistent. The chance alignments illustrate the contrast sensitivity of the *Gaia* eDR3

catalog: at close separations, pairs with large magnitude difference are not detected. In the bottom panel, a narrow excess population of binaries with $\Delta G \approx 0$ is visible, which is absent in the top panel. This highlights the excess population of equal-mass “twin” binaries.

4.1 Twin binaries and contrast sensitivity

An excess population of equal-brightness (and presumably, equal-mass) “twin” binaries is also found in the catalog. Its existence is most obvious in the distribution of magnitude difference, ΔG , as a function of separation, which is shown in Figure 9 and compared to chance alignments. Unlike the chance alignment catalogs constructed from the mock catalog and shifted catalogs (Section 3), these are selected in the same way as true binaries, but with the requirement that the parallaxes and proper motions of the two components be *inconsistent*. For easier visualization of the distribution of magnitude difference near $\Delta G = 0$, the sign of ΔG is randomized. The top panel illustrates the separation-dependent contrast sensitivity of the *Gaia* eDR3 catalog: at close angular separations, sources with significantly brighter companions are outshone. This leads to a contrast limit of $\Delta G \approx 4$ mag at $\theta = 1$ arcsec and $\Delta G \approx 7$ mag at $\theta = 2$ arcsec. The contrast limit at a given separation is not “sharp”, but is manifest as a smooth drop in sensitivity with increasing ΔG (e.g. Brandeker & Cataldi 2019).² The contrast sensitivity is significantly improved in the binary catalog produced in this work compared to the one produced by ER18: that work required both components to have relatively uncontaminated BP/RP colors and thus contained basically no binaries closer than 2 arcsec.

The bottom panel of Figure 9 shows true binary candidates, in which the two components *do* have consistent parallaxes and proper motions. Unlike with the chance alignments, here there is a narrow excess population with $\Delta G \approx 0$. The extent and provenance of this population was studied by El-Badry et al. (2019); here we simply note that it is also clearly apparent in our catalog. Because the twin excess is most prominent at close physical separations, it is somewhat more obvious in our catalog, which extends to closer angular and physical separations.

² At very close separations ($\theta \lesssim 0.7$ arcsec; not shown in Figure 9), the *Gaia* eDR3 catalog contains only equal-brightness pairs (Lindegren et al. 2020b, their Figure 6). The contrast sensitivity is relatively smooth at $\theta > 1$ arcsec.

4.2 Comparison to other catalogs

Figure 10 compares the distributions of projected physical separation, angular separation, and distance of the catalog produced here to other wide binary catalogs in the literature:

- The ER18 catalog (magenta) was produced from *Gaia* DR2 using a similar strategy to this work, but it was limited to binaries within 200 pc ($\varpi > 5$ mas) and used more stringent cuts on both photometric and astrometric quality and SNR. In order to maintain high purity, it only contains binary candidates with $s < 50,000$ au.

- Tian et al. (2020) expanded the ER18 search strategy to a larger volume (4 kpc) and used less stringent quality cuts, while still using *Gaia* DR2 data. Unlike this work, which uses a cut of $\varpi/\sigma_\varpi > 5$ for both components, they required $\varpi/\sigma_\varpi > 20$ for the primary and $\varpi/\sigma_\varpi > 2$ for the secondary. This results in somewhat different contamination properties and completeness. They also searched out to $s = 1$ pc. In addition to their full catalog of all candidates, Tian et al. (2020) published 3 smaller catalogs with high purity, which are not shown in Figure 10.

- Hartman & Lépine (2020) did not use a strict distance cut, but limited their search to high-proper-motion pairs with $\mu > 40$ mas yr⁻¹. This preferentially selects nearby stars, since proper motion is inversely proportional to distance at fixed transverse velocity. For a “typical” tangential velocity of $v_\perp = 35$ km s⁻¹, their proper motion cut corresponds to $d < 185$ pc. However, stars on halo-like orbits with larger tangential velocities are included to larger distances; e.g., a binary with $v_\perp = 200$ km s⁻¹ will have $\mu > 40$ mas yr⁻¹ out to a distance of 1.05 kpc. Rather than employing strict cuts on parallax and proper motion consistency, Hartman & Lépine (2020) used empirical estimates of the chance alignment rate as a function of position and proper motion difference from a shifted catalog (similar to our approach in Section 3) to distinguish true binaries from chance alignments. Their approach has the advantage of not requiring specific cuts in parallax or proper motion difference, which are always somewhat arbitrary. A disadvantage is that it does not account for the heteroskedasticity of parallax and proper motion uncertainties – i.e., the chance alignment probability is higher for pairs with large astrometric uncertainties, and this is not accounted for in their analysis.

One difference between the catalog produced in this work and the other catalogs is obvious in Figure 10: our sample extends to smaller angular separations, and thus also physical separations. This is partly a result of the improved angular resolution of *Gaia* eDR3 (e.g., Fabricius et al. 2020) but is primarily due to a change in search strategy. ER18 and Tian et al. (2020) required both components of candidate binaries to have $G_{BP} - G_{RP}$ colors and to pass photometric quality cuts related to the `bp_rp_excess_factor` reported in *Gaia* DR2 (Evans et al. 2018). This set a soft resolution limit of ~ 2 arcsec, with a wider effective limit for pairs with large brightness contrast. We do not require colors or employ a photometric quality cut in this work; this adds an additional $\sim 400,000$ binary candidates to the sample that would be excluded if we did (Table 1).

Both the catalog produced in this work and the one from Tian et al. (2020) become dominated by chance alignments at $s \gtrsim 30,000$ au. The “high bound probability” subset of our catalog does not, but its separation distribution falls off steeply at wide separations. This decline at wide separations is steeper than that of the intrinsic separation distribution, since a decreasing fraction of binaries at wide separations can be identified as bound with high confidence. This can be seen in comparing the red separation distribution to the magenta one from ER18, which tracks the intrinsic separation distribution

over $5000 < s/\text{au} < 50,000$. The separation distribution from Hartman & Lépine (2020) has a similar logarithmic slope to the one from ER18 in this separation range and likely tracks the intrinsic separation distribution out to wider separations. At the very widest separations represented in that catalog (5+ pc, exceeding the local Jacobi radius and corresponding to an orbital period of about a Gyr), it is unlikely that pairs are actually bound. This may reflect the fact that the search strategy employed by Hartman & Lépine (2020) is sensitive to any pairs that are closer in phase space than chance alignments from the shifted catalog, without explicit consideration of the expected orbital velocities. That is, their search does not distinguish between bound binaries, moving groups, or stellar streams.

In terms of absolute numbers, the catalog represents a factor of ~ 4 increase in the number of high-confidence binaries over the one from Tian et al. (2020). To our knowledge, it is the largest published catalog of high-confidence binaries of any type. The sample could likely be expanded by a further factor of a few by loosening the distance and parallax uncertainty limits, or dropping the parallax cut entirely, while focusing on close angular separations (e.g. Dhital et al. 2015). However, the cuts we use in this paper provide a reasonable compromise between sample size, purity, and data quality.

A wide binary catalog based on *Gaia* eDR3 was also produced by Gaia Collaboration et al. (2020b), which contains pairs within 100 pc. It is not shown in Figure 10, but we find that within 100 pc, it is almost identical to ours.

4.3 Space density

Figure 11 compares the separation distributions of binaries in different (cumulative) distance bins. We scale these by the effective stellar volume corresponding to the distance cut, i.e., by a factor proportional to the total number of stars expected in each distance sample. Were it not for the finite scale height of the Galactic disk, this factor would simply be the search volume $V = 4\pi d_{\text{max}}^3/3$, where d_{max} is the distance limit. We approximate the total stellar density within our search volume as a plane-parallel exponential distribution with the Sun at the midplane and a scale height $h_z = 300$ pc (Juric et al. 2008). We then define an effective volume \tilde{V} , which represents the number of stars in a sphere of radius d_{max} divided by the stellar density at the disk midplane:

$$\tilde{V} = 2\pi \int_0^{d_{\text{max}}} e^{-z/h_z} (d_{\text{max}}^2 - z^2) dz \quad (9)$$

$$= 2\pi \left[2h_z^3 \left(e^{-d_{\text{max}}/h_z} - 1 \right) + 2h_z^2 d_{\text{max}} e^{-d_{\text{max}}/h_z} + h_z d_{\text{max}}^2 \right]. \quad (10)$$

As expected, this expression asymptotes to $4\pi d_{\text{max}}^3/3$ in the limit of $d_{\text{max}} \ll h_z$. In Figure 11, the separation distributions for each value of d_{max} are divided by the appropriate value of \tilde{V} .

Dashed vertical lines in Figure 11 mark a separation of $s_{\text{res limit}} = (2 \text{ arcsec}) \times d_{\text{max}}$. At separations $s < s_{\text{res limit}}$, incompleteness due to the *Gaia* eDR3 angular resolution starts to become severe. The figure shows that for $d_{\text{max}} \lesssim 200$ pc, incompleteness is due primarily to the angular resolution limit: at $s > s_{\text{res limit}}$, the separation distributions in different distance bins overlap. However, for $d_{\text{max}} = 500$ pc or 1000 pc, the catalog contains fewer binaries per effective volume than at closer distances, even at $s > s_{\text{res limit}}$. This reflects the fact that at sufficiently large distances, some binaries will have components that are too faint to pass the `parallax_over_error > 5` limit, or to be detected at all. At $d = 200$ pc, `parallax_over_error > 5` implies $\sigma_\varpi < 1$ mas. This is satisfied by most sources with $G \lesssim 20.5$ (Lindgren et al. 2020b), corresponding to $M_G = 14.0$, near the




Figure 10. Comparison of the catalog produced in this work (black and red histograms) to the wide binary catalogs published by ER18, Tian et al. (2020), and Hartman & Lépine (2020), all of which are based on *Gaia* DR2 data. This catalog expands the one produced by Tian et al. (2020) by a factor of ~ 4 . The full catalog (black) has lower purity than the El-Badry & Rix (2018) and Hartman & Lépine (2020) catalogs at wide separations. The subset with high bound probability (red) maintains high purity at wide separations, but does not purport to trace the intrinsic separation distribution there.




Figure 11. Midplane number density of binaries in different cumulative distance bins. Vertical lines mark $\theta = 2$ arcsec at the distance limit of each bin, roughly the angular separation at which incompleteness due to blending becomes significant. At all distances, the turnover at close separations is driven primarily by the resolution limit, but the 15 and 25 pc samples show a relatively flat separation distribution below a few hundred au. Beyond 200 pc, there is significant incompleteness due to stars being too faint to pass our fractional parallax error threshold – i.e., the calculated number densities are lower than in the nearby samples even at wide separations, where blending should not be important.

bottom of the main sequence (e.g. Figure 8). That is, the sample is expected to be almost complete for $d_{\max} = 200$ pc, except for crowding/blending effects at close separations. On the other hand, at $d = 1$ kpc, $\text{parallax_over_error} > 5$ implies $\sigma_{\varpi} < 0.2$ mas, $G \lesssim 18.8$, and $M_G \lesssim 8.8$, meaning that most of the lower main sequence will be excluded.

All of the separation distributions in Figure 11 increase toward smaller separations at $s > s_{\text{res limit}}$, but the distributions for $d_{\max} = 15$ pc and $d_{\max} = 25$ pc are similar and do appear to flatten above the resolution limit, at $s \sim 30$ au. For a typical binary in the catalog with total mass $1 M_{\odot}$, this corresponds to a period of order 200 years, which is indeed near the peak of the approximately

lognormal separation distribution for solar-type binaries (Raghavan et al. 2010). We caution that effects of astrometric acceleration also become important in this regime (Section 6.3), potentially leading to spurious parallaxes and preventing pairs from being recognized as binaries by our search.

Integrating over separation, the distributions in Figure 11 imply a total space density of (0.006 ± 0.001) wide binaries with $s > 30$ au per cubic parsec in the solar neighborhood. For context, the space density of all unresolved *Gaia* eDR3 sources in the solar neighborhood is 0.07 pc^{-3} (Gaia Collaboration et al. 2020b), about 10 times higher. When all members of multiple systems are counted individually, the total stellar space density in the solar neighborhood is about 0.10 pc^{-3} (e.g. Winters et al. 2020).

4.4 Orbital velocities from proper motion differences

Precise parallaxes and proper motions make it possible to estimate orbital velocities (projected onto the plane of the sky) from the proper motion difference between the two stars. The plane-of-the-sky velocity difference ΔV can be calculated as

$$\Delta V = 4.74 \text{ km s}^{-1} \times \left(\frac{\Delta\mu}{\text{mas yr}^{-1}} \right) \left(\frac{\varpi}{\text{mas}} \right)^{-1}. \quad (11)$$

Here $\Delta\mu$ is the scalar proper motion difference (Equation 4), and ϖ is the parallax of the binary, for which we take the parallax of the brighter component. The corresponding uncertainty is

$$\sigma_{\Delta V} = 4.74 \text{ km s}^{-1} \sqrt{\frac{(\Delta\mu)^2}{\varpi^4} \sigma_{\varpi}^2 + \frac{\sigma_{\Delta\mu}^2}{\varpi^2}}, \quad (12)$$

with $\sigma_{\Delta\mu}$ calculated from Equation 5, $\Delta\mu$ and $\sigma_{\Delta\mu}$ in mas yr^{-1} , and ϖ and σ_{ϖ} in mas. We implicitly assume here that the two stars have the same parallax. Equation 12 is almost always dominated by the first term under the radical; i.e., parallax errors dominate over proper motion errors. The median value of $\sigma_{\Delta V}$ for all high-confidence binaries in the catalog is 0.33 km s^{-1} ; 195,601 have $\sigma_{\Delta V} < 0.1 \text{ km s}^{-1}$.

Figure 12 explores the ΔV values of binary candidates in the catalog, and the effects of the ΔV cuts we employ on its purity and completeness. The upper left panel shows all pairs from our initial

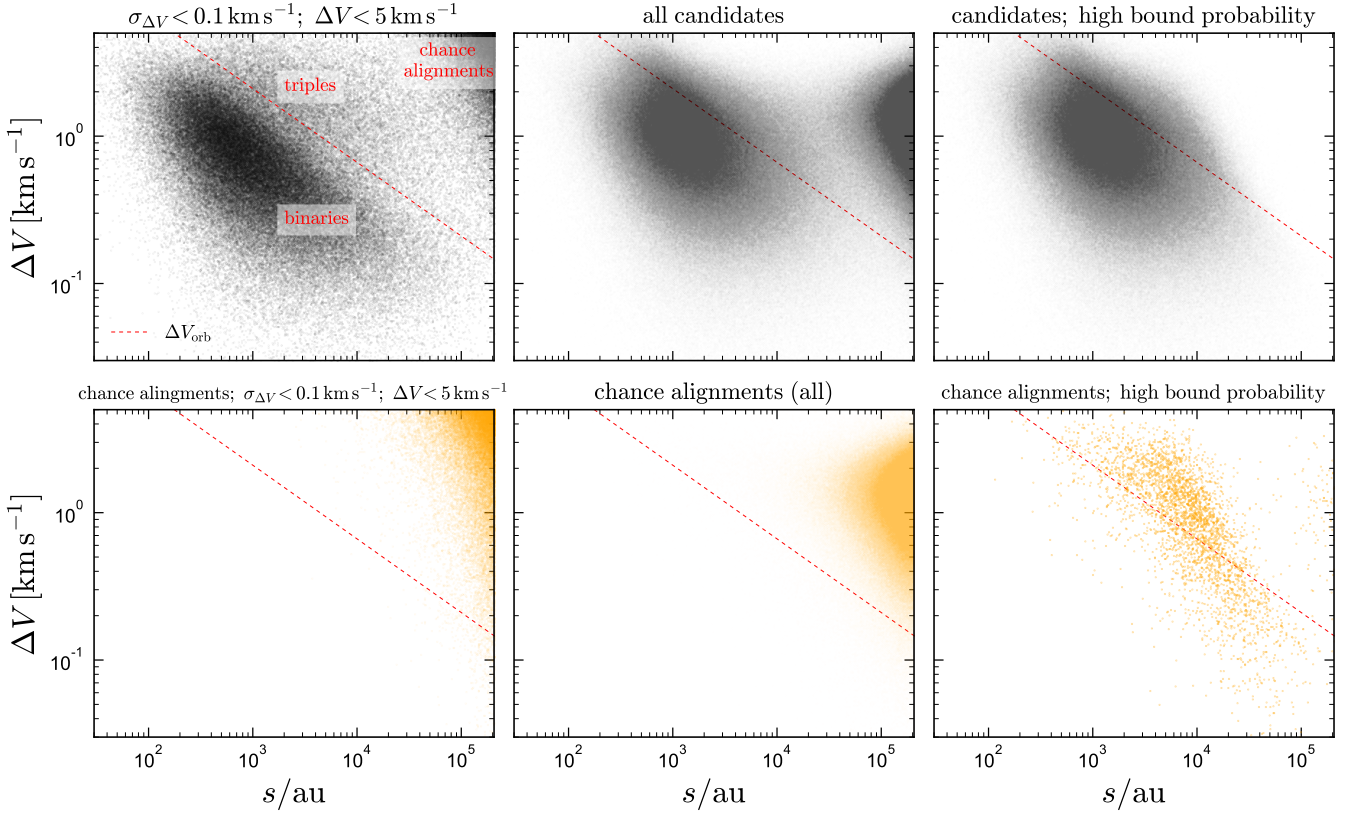


Figure 12. Upper left: plane-of-the-sky velocity difference between the components of wide binary candidates selected with $\Delta V < 5 \text{ km s}^{-1}$ and $\sigma_{\Delta V} < 0.1 \text{ km s}^{-1}$. Dashed red line shows the maximum expected velocity difference for a binary with total mass $5M_{\odot}$ (Equation 7). Binary candidates in our catalog (top center) are required to be consistent, within 2σ , with falling below this line, but pairs with large proper motion uncertainties can scatter well above it. The clouds of points at large separation and ΔV are chance alignments; bottom panels show chance alignments from the shifted catalogs for the same selection. Points at large separation and ΔV are excluded by the requirement of high bound probability ($\mathcal{R} < 0.1$; upper right). In the upper left panel, there is a population of triples and higher-order multiples above the dashed line, with no corresponding population in the shifted catalog (lower left). These have increased ΔV due to the gravitational effects of an unresolved close companion. They are generally gravitationally bound, but are excluded from our primary catalog.

query that have $s < 1 \text{ pc}$, consistent parallaxes, $\Delta V < 5 \text{ km s}^{-1}$, and $\sigma_{\Delta V} < 0.1$. Note that the cut of $\Delta V < 5 \text{ km s}^{-1}$ is generally less strict than the one we adopt in constructing our primary catalog (Equation 3), which is equivalent to $\Delta V < \Delta V_{\text{orb}} + 2\sigma_{\Delta V}$, with ΔV_{orb} given by Equation 7.

In the upper left panel, there is a clear ridgeline of binaries with $\Delta V \sim s^{-1/2}$, as expected from Kepler’s laws. This population largely falls below ΔV_{orb} (dashed red line) because most of the binaries in the catalog have total masses less than $5M_{\odot}$. There is not, however, a sharp drop-off at $\Delta V > \Delta V_{\text{orb}}$. The population with $\Delta V > \Delta V_{\text{orb}}$ likely consists primarily of triples and higher-order multiples, in which the plane-of-the-sky velocity of one component is affected by a close, unresolved companion (e.g. Clarke 2020; Belokurov et al. 2020). It is also possible that some of these pairs are moving groups that are not actually bound but remain close in phase space (e.g. Pittordis & Sutherland 2019; Coronado et al. 2020); however, we find that the population exists even among binaries on halo-like orbits, favoring multiplicity as the primary explanation for it. We also find that most pairs with $\Delta V > \Delta V_{\text{orb}}$ have unusually large ruwe^3 for at

least one component (Figure 13), suggesting that these components are unresolved binaries (e.g. Belokurov et al. 2020).

A large fraction of these unresolved triples and higher-order multiples are excluded from our actual catalog by the requirement of $\Delta V < \Delta V_{\text{orb}} + 2\sigma_{\Delta V}$ rather than e.g., $\Delta V < 5 \text{ km s}^{-1}$; this should be kept in mind when using the catalog for applications involving higher-order multiplicity. We do not use a constant ΔV cut (e.g. $\Delta V < 5 \text{ km s}^{-1}$) in constructing the full catalog because this would result in a much higher contamination rate from chance alignments. For $5,000 < s/\text{au} < 10,000$, 22% of pairs with $\Delta V < 5 \text{ km s}^{-1}$ and $\sigma_{\Delta V} < 0.1 \text{ km s}^{-1}$ are excluded from the catalog by Equation 3. If we assume that large ΔV values are due mainly to subsystems, this implies that about 12% of wide binary components in this separation range have an unresolved subsystem that imparts a large enough photocenter perturbation to significantly increase ΔV . This corresponds to a subsystem separation range of $\sim 1 - 100 \text{ au}$, since photocenter wobbles will average out over the 34-month Gaia eDR3 baseline at closer separations, and the subsystem-induced ΔV will be small at wider separations. About 25% of wide binary components have a subsystem in this separation range (Tokovinin & Smekhov 2002;

³ ruwe , the re-normalised unit-weight error, is a measure of astrometric goodness-of-fit that corrects for global trends in the other reported goodness-

of-fit indicators with magnitude and color. Values above about 1.4 indicate potential problems.

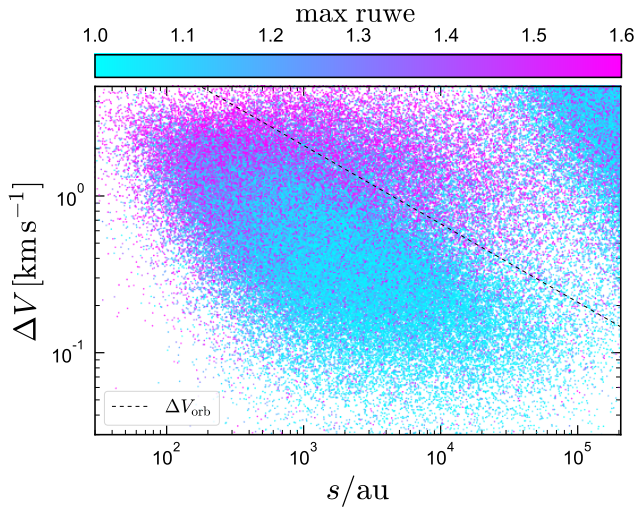


Figure 13. Binary candidates from the upper left panel of Figure 12, now colored by the maximum ruwe of the two components. Most pairs with $\Delta V > \Delta V_{\text{orb}}$ and $s \ll 10^5$ au (those labeled “triples” in Figure 12) have $\text{ruwe} \gtrsim 1.4$ for at least one component, in most cases because that component is an unresolved binary. At $s \lesssim 500$ au, most pairs have large ruwe due to centroiding errors at close angular separations (see also Figure 18).

Tokovinin et al. 2010), so it is quite plausible that the large- ΔV pairs are mostly unresolved triples and higher-order multiples.

The clouds of points at large separation and ΔV in Figure 12 are chance alignments. These are primarily pairs with large $\sigma_{\Delta V}$; otherwise they would be excluded by the requirement of $\Delta V < \Delta V_{\text{orb}} + 2\sigma_{\Delta V}$. The bottom panels show the distribution of chance alignments from the shifted catalog (Section 3). These are distributed similarly to the large-separation cloud among binary candidates. No chance-alignment cloud is visible in the binary candidates with high bound probability (upper right), but a few chance alignments do scatter into the bound binary cloud, mostly at $\Delta V > \Delta V_{\text{orb}}$ (lower right).

We note that the interpretation of ΔV calculated from Equation 11 as a physical velocity difference between the two components of a binary breaks down at large angular separations, where projection effects become important (e.g. El-Badry 2019). Indeed, two stars in an ultra-wide binary can have *identical* space velocities but substantially different plane-of-the-sky proper motions. The magnitude of the apparent proper motion difference depends primarily on angular separation and is therefore largest for nearby binaries, which are also the binaries with the smallest $\sigma_{\Delta V}$. For the sample shown in the upper left panel of Figure 12, projection effects become important beyond about 20,000 au, which is – perhaps not coincidentally – the separation beyond which the trend of $\Delta V \sim s^{-1/2}$ appears to flatten. These projection effects can be corrected if the RVs of at least one component are known (El-Badry 2019).

4.5 Cross-match with LAMOST

We cross-matched the binary catalog with the LAMOST survey (Cui et al. 2012, DR6 v2), the currently most extensive spectroscopic survey providing stellar parameters and abundances. We began with the LAMOST low-resolution “A, F, G and K Star” catalog, which contains atmospheric parameters and metallicities for 5,773,552 spectra (including some duplicate observations). We cross-matched the cat-

alog with *Gaia* eDR3 using the CDS Xmatch service,⁴ which uses *Gaia* proper motions to propagate source positions to epoch J2000. We matched each LAMOST observation to the nearest *Gaia* source within 1 arcsec. For sources with more than one LAMOST observation, we retained only the observation with the highest g -band SNR. This left us with LAMOST data for 4,306,131 sources, which we then matched to our catalog using *Gaia* source ids.

This yielded 91,477 binaries in which at least one component has a LAMOST spectrum. This sample will be useful for a variety of applications, such as studying the dependence of the binary fraction on metallicity (e.g. El-Badry & Rix 2019; Hwang et al. 2020). Here, we focus on a subset of the cross-match: those binaries in which *both* components have a LAMOST spectrum, and the angular separation is at least 3 arcseconds. The latter cut is to avoid cases where both stars fall inside a single fiber, leading to potentially biased stellar parameters and abundances (e.g. El-Badry et al. 2018). The RVs of these binaries can be used to verify whether most binary candidates are bound. This test is similar to that shown in Figure 6 with *Gaia* RVs, but LAMOST spectra extend to fainter magnitudes than RVs from *Gaia*, which are currently only available at $G \lesssim 13$.

Figure 14 compares the LAMOSTs RVs and metallicities (i.e., $[\text{Fe}/\text{H}]$) of the components of these binaries. These are generally expected to be consistent for genuine wide binaries. The RVs for most binaries do indeed fall close to the one-to-one line, but there are some outliers: 91 of the 2840 binaries in the sample have RVs that are more than 3-sigma discrepant. A potential worry is that these pairs are not binaries at all, but chance alignments. To assess whether this is likely to be the case, we plot the separation distribution of the discrepant pairs in the bottom right panel of Figure 14. This shows that the separation distribution of pairs with large RV differences is similar to that of all pairs. Chance alignments are much more common at wide separations (Figure 4), so if chance alignments were the root of the discrepant RVs, one would expect these pairs to be clustered at large separations. The primary reason for the discrepant RVs is likely again higher-order multiplicity. This is likely to affect the LAMOST RVs more than it does those from *Gaia* (Figure 6), because the LAMOST RVs and their uncertainties are based on a single epoch. Unresolved short-period binaries are usually filtered out of the *Gaia* RV sample we consider, because σ_{RV} for that sample is calculated from the epoch-to-epoch RV dispersion, and we required $\sigma_{\text{RV}} < 10 \text{ km s}^{-1}$ for both components.

The lower left panel of Figure 14 shows the distribution of uncertainty-normalized RV difference between the two components of binaries; i.e., $\Delta \text{RV} / \sigma_{\Delta \text{RV}} = (\text{RV}_1 - \text{RV}_2) / \sqrt{\sigma_{\text{RV},1}^2 + \sigma_{\text{RV},2}^2}$. The median $\sigma_{\Delta \text{RV}}$ for this sample is 7 km s^{-1} , which is larger than the typical $\sim 1 \text{ km s}^{-1}$ orbital velocity for these binaries, so we expect the width of the main distribution to be dominated by measurement uncertainties, with the tails dominated by higher-order multiples. If the reported σ_{RV} values are accurate, $\Delta \text{RV} / \sigma_{\Delta \text{RV}}$ should be distributed as a Gaussian with $\sigma \approx 1$, with some outliers at higher velocity difference. The distribution is indeed approximately Gaussian, but it is narrower than $\sigma = 1$; the bulk of the distribution is better approximated by $\sigma \approx 0.7$. This suggests that the LAMOST RV uncertainties are typically *overestimated* by $\sim 30\%$.

The distribution of uncertainty-normalized $[\text{Fe}/\text{H}]$ differences tells a different story. It is not well-described by a single Gaussian, but has a narrow component with $\sigma \lesssim 1$, and broad, asymmetric tails. This suggests that $\sigma_{[\text{Fe}/\text{H}]}$ values are considerably underestimated

⁴ <http://cdsxmatch.u-strasbg.fr/>

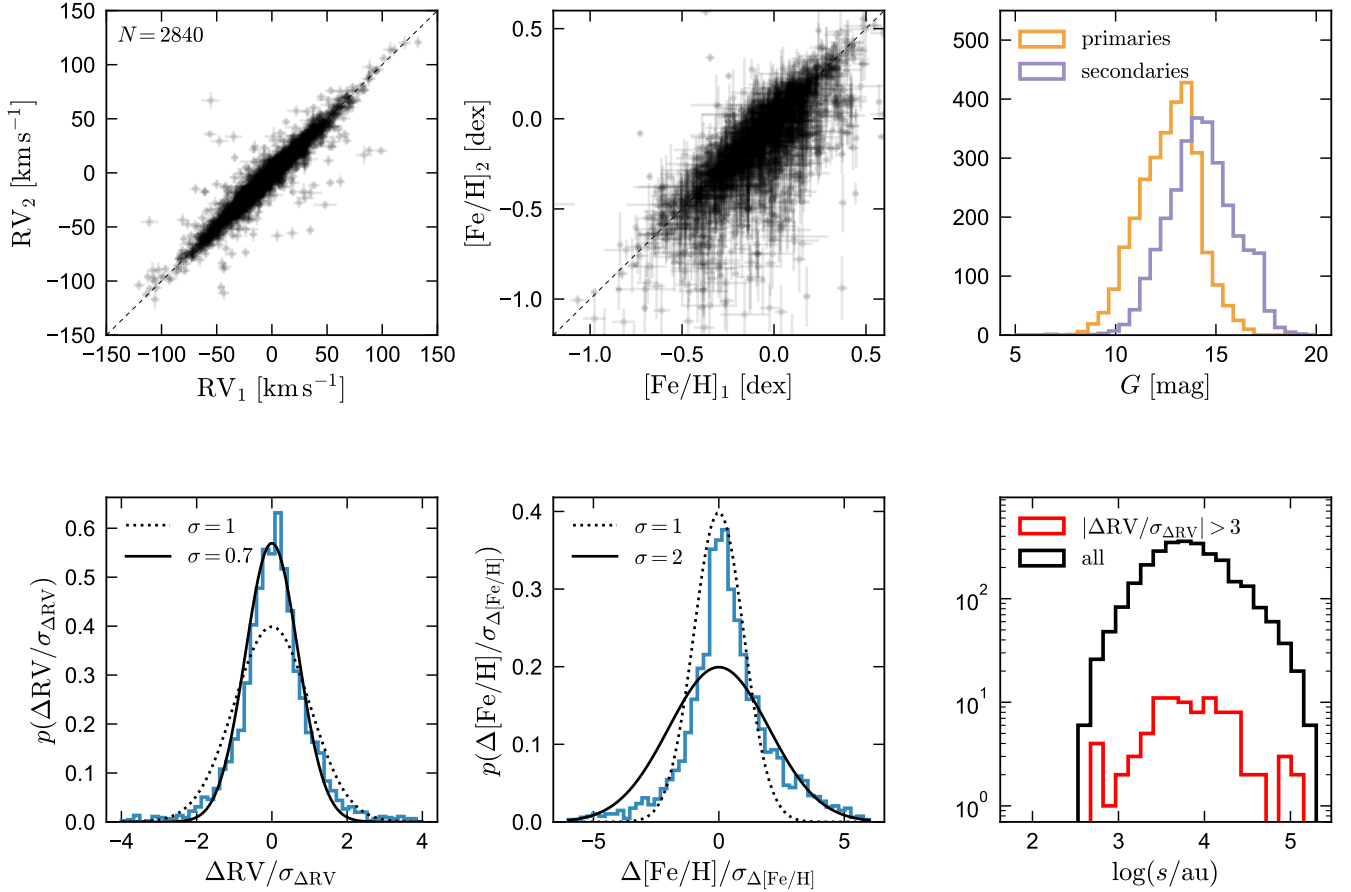


Figure 14. High-confidence binaries in which both components were observed by the LAMOST survey. We only include pairs with angular separation $\theta > 3$ arcsec to avoid blending. Top left and center: comparison of the RVs and metallicities ($[\text{Fe}/\text{H}]$) of the primary (x -axis) and secondary (y -axis). In most binaries, the two components have very similar RV and $[\text{Fe}/\text{H}]$, as expected. There is evidence of a systematic bias toward lower $[\text{Fe}/\text{H}]$ for fainter targets. Upper right: G magnitudes of both components. This sample includes fainter stars than the *Gaia* RVs sample. Bottom left and center: distributions of uncertainty-normalized RV and $[\text{Fe}/\text{H}]$ difference, compared to Gaussians with the listed σ . These suggest that the LAMOST RV uncertainties are overestimated by $\sim 30\%$ on average, while $[\text{Fe}/\text{H}]$ uncertainties are underestimated and subject to temperature systematics. Bottom right: separation distributions. Pairs in which RVs are not consistent within 3σ are shown in red. If these pairs were chance-alignments, one would expect them to be concentrated at the widest separations. They are not, suggesting that they are either higher-order multiples with the RV of one component biased by an unresolved close companion, or pairs in which one component has a catastrophically wrong RV.

for a significant fraction of the catalog. We find (not shown in the figure) that the distribution becomes narrower and more Gaussian when we only consider binaries in which the two components have similar magnitude and effective temperature. This suggests that the larger-than-expected metallicity differences are due primarily to temperature systematics in the abundance pipeline. This is particularly evident in the upper middle panel of Figure 14, which shows that the secondaries (which have lower T_{eff}) systematically have lower reported $[\text{Fe}/\text{H}]$. It is of course possible that some binaries really have inconsistent surface abundances, but work with higher-quality spectra than those which underlie the LAMOST metallicities suggests such abundance anomalies are rare (Hawkins et al. 2020).

Many of the binaries in our catalog were also observed by other spectroscopic surveys. We defer analysis of these data to future work, but comment that the type of analysis shown in Figure 14 will be useful in calibrating the abundances derived by surveys (and their uncertainties). For example, an earlier version of our binary catalog constructed from *Gaia* DR2 was recently fruitfully used by Buder

et al. (2020) to assess the reliability of abundances derived by the GALAH survey.

5 CALIBRATING GAIA DR3 PARALLAX UNCERTAINTIES

Because the two stars in a wide binary have very nearly the same distance, our catalog provides a straightforward method of validating the *Gaia* eDR3 parallax uncertainties. We do this by calculating the uncertainty-normalized parallax difference between the two components of each binary, $\Delta\varpi/\sigma_{\Delta\varpi} = (\varpi_1 - \varpi_2) / \sqrt{\sigma_{\varpi,1}^2 + \sigma_{\varpi,2}^2}$, just as we did for the LAMOST RVs. In the limit of accurate parallax uncertainties (and small differences in the true distance to the two components), this quantity should be distributed as a Gaussian with $\sigma = 1$. If the reported parallax uncertainties are underestimated, one expects a wider distribution, and possibly deviations from Gaussianity.

Figure 15 shows distributions of $\Delta\varpi/\sigma_{\Delta\varpi}$ for binaries at a range of angular separations and magnitudes. To isolate underestimated

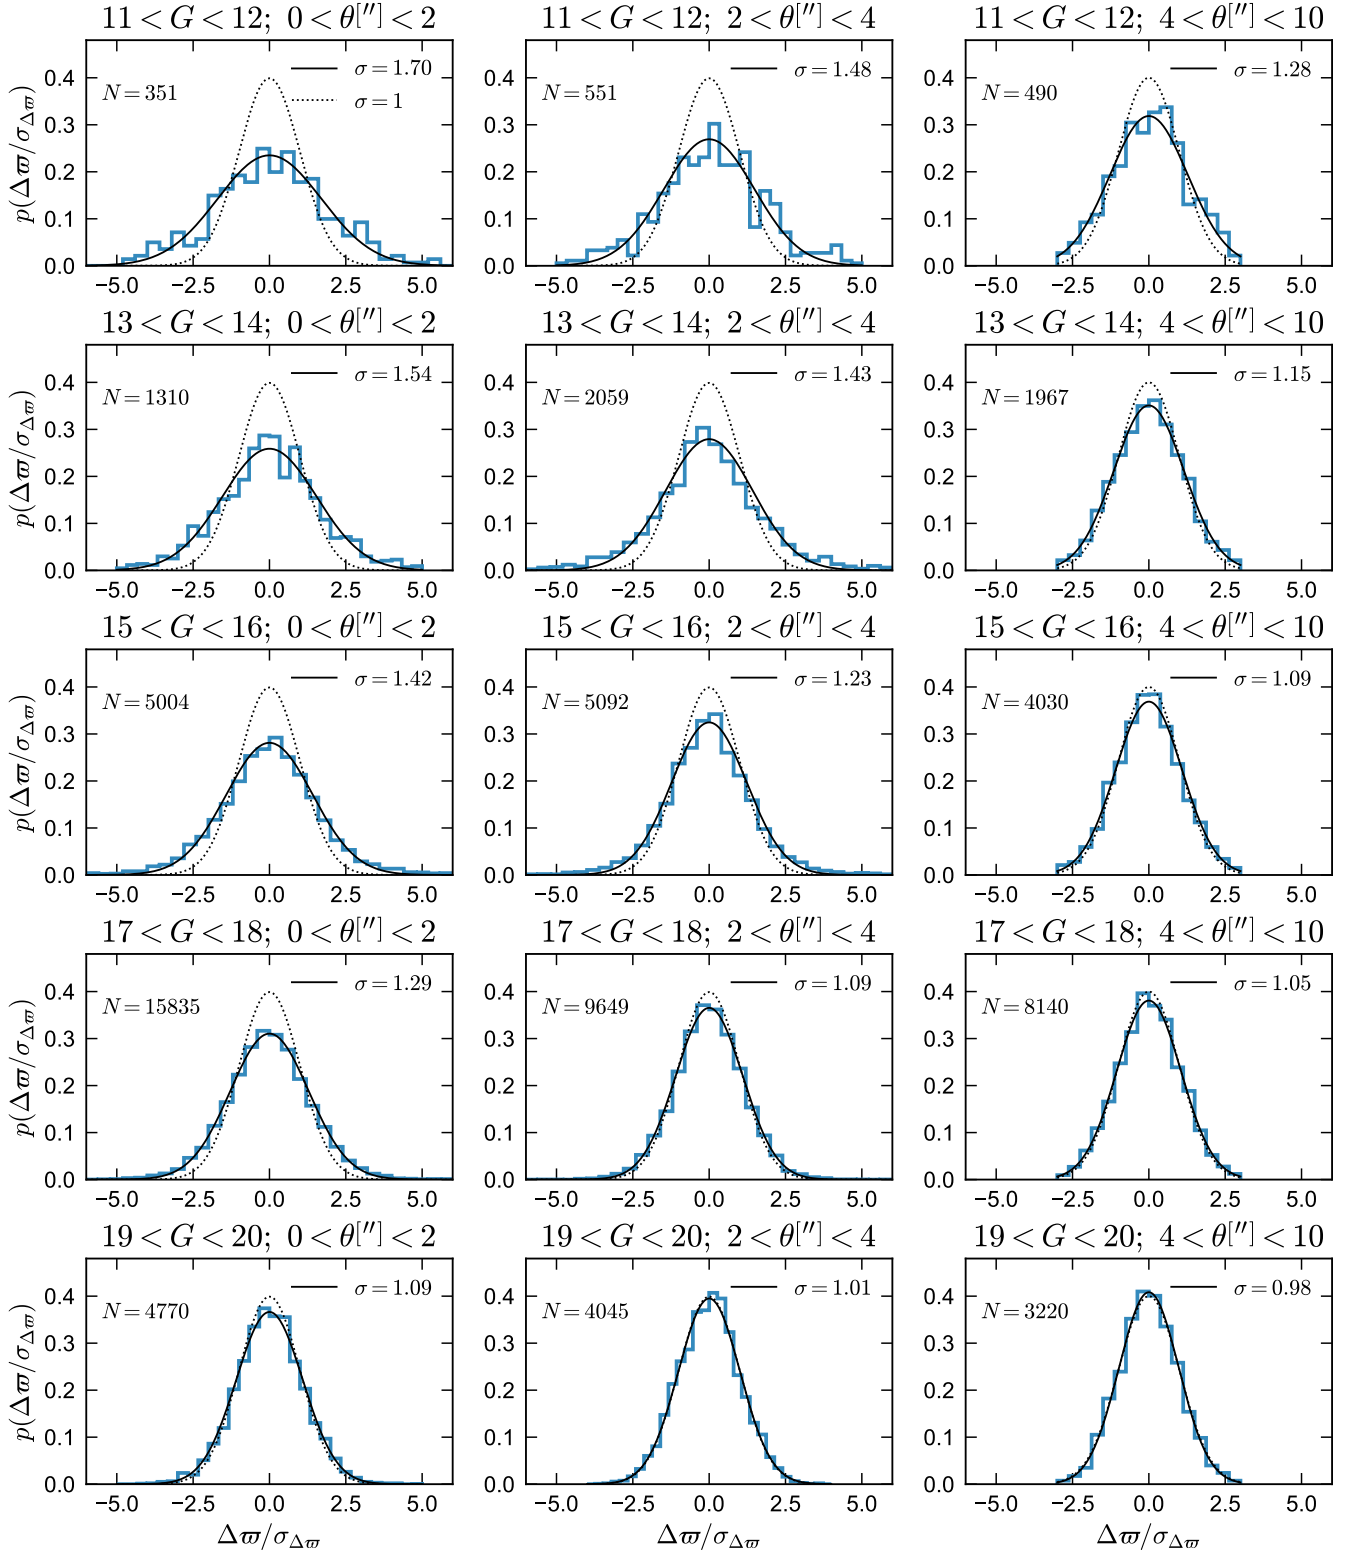


Figure 15. Distributions of uncertainty-normalized parallax difference between the two components of binaries, $\Delta\pi/\sigma_{\Delta\pi} = (\varpi_1 - \varpi_2)/\sqrt{\sigma_{\varpi,1}^2 + \sigma_{\varpi,2}^2}$. Each panel shows a different bin of angular separation and magnitude, with both components falling in the quoted magnitude range. Because the two components of a binary have essentially the same distance, this quantity would be expected to follow a Gaussian distribution with $\sigma = 1$ (dotted lines) if the formal parallax uncertainties were accurate. $\sigma > 1$ points toward underestimated parallax uncertainties. Solid black lines show Gaussian fits, with σ noted in the legends. At all magnitudes, the best-fit σ is larger at close separations, implying that σ_{ϖ} is more severely underestimated for sources with nearby companions. At fixed separations, the fractional underestimate of σ_{ϖ} is larger for bright stars.

random errors (as opposed to systematic shifts due to e.g. variations in the parallax zeropoint), we only show binaries in which both components have magnitudes in the quoted range (corresponding to a magnitude difference of $\Delta G < 1$). This figure only shows high-confidence binaries ($\mathcal{R} < 0.1$) in which both components have $\text{ruwe} < 1.4$, indicating an apparently well-behaved astrometric solution. In this and all validation of the parallax errors, we exclude binaries that are wide and nearby enough that their physical size might measurably contribute to the parallax difference. Under the ansatz that the projected physical separation s is comparable to the line-of-sight distance difference and is much smaller than the distance, we expect the true parallax difference between the two stars to scale as $\Delta\varpi_{\text{true}} \sim 1/d - 1/(s+d) \approx s/d^2$, or

$$\Delta\varpi_{\text{true}} = \frac{1}{206265} \text{ mas} \times \left(\frac{\theta}{\text{arcsec}} \right) \left(\frac{\varpi}{\text{mas}} \right). \quad (13)$$

This quantity is negligible compared to $\sigma_{\Delta\varpi}$ for most binaries in the catalog, but not for the nearest and widest binaries. We therefore exclude all binaries in which $\Delta\varpi_{\text{true}}/\sigma_{\Delta\varpi} > 0.05$.

Blue histograms in Figure 15 show the observed distributions. We show 5 bins of G magnitude, each 1 mag wide, and 3 bins of angular separation. At separations closer than 4 arcsec, the catalog contains binaries with $|\Delta\varpi|/\sigma_{\Delta\varpi} < 6$; at wider separations, $|\Delta\varpi|/\sigma_{\Delta\varpi} < 3$ (Section 2). Also shown in Figure 15 are Gaussian fits to the data. Because pairs with $|\Delta\varpi|/\sigma_{\Delta\varpi} < 3$ (or 6) do not enter the catalog, it is necessary to account for this truncation of the distribution. Particularly when parallax uncertainties are underestimated, simply calculating the sample standard deviation would underestimate the best-fit σ . We assume the observed values of $\Delta\varpi/\sigma_{\Delta\varpi}$ are drawn from a distribution defined as

$$p(\Delta\varpi/\sigma_{\Delta\varpi}) = \begin{cases} A \exp\left[-\frac{(\Delta\varpi/\sigma_{\Delta\varpi})^2}{2\sigma^2}\right], & |\Delta\varpi/\sigma_{\Delta\varpi}| < b \\ 0, & |\Delta\varpi/\sigma_{\Delta\varpi}| > b \end{cases} \quad (14)$$

where $A = \frac{1}{\sqrt{2\pi}\sigma \times \text{erf}(b/\sqrt{2}\sigma)}$, erf is the error function, and $b = 3$ or 6 is the sigma-threshold above which binaries are rejected. The log-likelihood for a set of uncertainty-normalized parallax differences is then

$$\ln L = \sum_i \ln p((\Delta\varpi/\sigma_{\Delta\varpi})_i), \quad (15)$$

where the sum is calculated over all binaries in the set. For each panel of Figure 15, we maximize Equation 15 to find the value of σ that best describes the truncated Gaussian. This is plotted with a solid line, and the value of σ is shown in the legend. For comparison, we also plot a Gaussian with $\sigma = 1$, the distribution expected in the limit of accurate parallax uncertainties. To avoid having a few outliers with strongly underestimated parallax uncertainties bias the fits, we exclude binaries with $|\Delta\varpi|/\sigma_{\Delta\varpi} > 3$ and set $b = 3$ in all separation bins. The figure shows that the distributions of $\Delta\varpi/\sigma_{\Delta\varpi}$ are indeed approximately Gaussian, and the fits are reasonably good representations of the data.

Two trends are clear in Figure 15: (a) at fixed apparent magnitude, the width of the observed distributions increases at closer separations, and (b) at fixed separation, their width increases toward brighter magnitudes. That is, parallaxes are more strongly underestimated for bright sources with close companions. This is illustrated more clearly in Figure 16, in which black points with error bars show the best-fit Gaussian σ values for each bin of magnitude and angular separation, considering only sources with $\text{ruwe} < 1.4$ for both components. At all separations, parallaxes are most severely underestimated at $G \approx 13$. This is likely related to the fact that the window class (i.e., the pixel

sampling scheme around detected sources; see Rowell et al. 2020) changes at $G \approx 13$. The largest σ_{ϖ} underestimate factors are also accompanied by abrupt changes in the zeropoint (Lindgren et al. 2020a). For sources with $G < 13$, a 2D window is used and is fitted with a point spread function. For $G > 13$, only the collapsed 1D scan, which is fitted with a line spread function, is available. Sources near $G = 13$ have a mix of 2D and 1D windows and are likely more affected by any calibration issues.

Figure 16 also shows results for binaries in which at least one component has $\text{ruwe} > 1.4$, indicative of a potentially problematic astrometric solution. As expected, the best-fit σ values are significantly larger for these binaries, at all separations and magnitudes.

A natural question is whether the broadened distributions of uncertainty-normalized parallax difference could be a result of contamination from chance alignments that are not actually bound, rather than underestimated parallax uncertainties. This hypothesis can be ruled out for two reasons. First, we can empirically estimate the chance-alignment rate for different subsets of the catalog (e.g. Figures 3 and 4), and we find it to be extremely low for the samples we use for parallax error validation. Second, our analysis suggests that σ_{ϖ} is most severely underestimated for bright binaries at close separations, and this is precisely the region of parameter space where the chance alignment rate is lowest (e.g. Figure 4).

Besides ruwe , *Gaia* eDR3 contains other diagnostics of potentially problematic astrometric fits. In particular, the parameter `ipd_gof_harmonic_amplitude` quantifies how much the image parameter determination goodness-of-fit varies with scan angle; a large value is likely indicative of a marginally-resolved binary. The related `ipd_frac_multi_peak` diagnostic quantifies in what fraction of scans multiple peaks are detected. Figure 17 separately plots the inferred σ_{ϖ} underestimate for sources in which both components have `ipd_gof_harmonic_amplitude` < 0.1 and `ipd_frac_multi_peak` < 10 (green), and those in which at least one component fails one of these cuts (blue). These thresholds are motivated by the experiments performed in Gaia Collaboration et al. (2020b). At all separations, the implied σ_{ϖ} underestimate is larger for sources that do not pass one of the IPD cuts; the difference is largest for bright pairs at close separations.

Figure 18 shows how the prevalence of problematic sources according to the ruwe and IPD diagnostics depends on separation. Considering all binaries in the catalog, we plot the fraction of pairs at a given separation in which at least one component does not pass the cut listed in the legend. At separations larger than a few arcseconds, this fraction is $\sim 10\%$ for the ruwe cut, and $\sim 3\%$ for both the IPD cuts. However, all three diagnostics of problematic solutions increase steeply at $\theta \lesssim 2$ arcsec, with the fraction approaching unity for the ruwe and `ipd_frac_multi_peak` cuts, and $\sim 40\%$ for the `ipd_gof_harmonic_amplitude` cut. This is not unexpected: for binaries at close separations, there will necessarily be two peaks in the image. This will unavoidably lead to biases in the image parameter determination, particularly for sources with $G > 13$, where the images are collapsed to 1D. There is also a danger of misattributing some scans to the wrong component in close pairs, leading to problems in the astrometric solution. In some cases, poor astrometric fits may also be due to astrometric acceleration (Section 6.3). Figures 16 and 17 show that the ruwe and IPD cuts are indeed useful for identifying sources with potentially problematic astrometric solutions. However, they likely do not catch all problematic sources: parallax uncertainties are underestimated somewhat even for sources that pass all cuts, and by a larger factor at close separations.




Figure 16. Fractional parallax uncertainty underestimate (i.e., the best-fit σ in Figure 15) as a function of magnitude and separation. Black symbols show binaries in which both components have $\text{ruwe} < 1.4$; red symbols show those in which at least one component has $\text{ruwe} > 1.4$, indicating a potentially problematic astrometric solution. From left to right, panels show increasing angular separation. As expected, sources with $\text{ruwe} > 1.4$ have larger σ_{ϖ} underestimates at all magnitudes and separations. Sources with $\text{ruwe} < 1.4$ have more severe underestimates of σ_{ϖ} at close separations and for bright stars, particularly at $11 < G < 13$.




Figure 17. Fractional parallax uncertainty underestimate (i.e., the best-fit σ reported in Figure 15) as a function of magnitude. All binaries considered have $\text{ruwe} < 1.4$ for both components. Green points additionally have $\text{ipd_gof_harmonic_amplitude} < 0.1$ and $\text{ipd_frac_multi_peak} < 10$ for both components; blue points have at least one component that does not pass these cuts. Sources that do not pass the IPD cuts have more strongly underestimated uncertainties at all separations and magnitudes.

5.1 Comparison to Gaia DR2

Figure 19 compares the inferred underestimate factor of σ_{ϖ} for DR2 and eDR3 astrometry. We attempt to match all the binary candidates in the catalog with *Gaia* DR2 using the `dr2_neighbourhood` catalog in the *Gaia* archive. For each component of each binary, we identify the likely corresponding DR2 source as the source within 100 mas that has the smallest magnitude difference compared to DR3. There are 1,894 primaries and 15,274 secondaries in the catalog for which no corresponding DR2 source could be identified. There are also 17,514 primaries and 62,424 secondaries for which there is a corresponding source in DR2 that only has a 2-parameter solution. Still, 96% of candidates have a corresponding DR2 source with a 5-parameter solution for both components. Figure 19 considers the

subset of these binaries in which both components have $\text{ruwe} < 1.4$ in both DR2 and eDR3.

Overall trends with magnitude and separation are similar in DR2 and DR3. At $G > 13$, the inferred σ_{ϖ} underestimates are smaller in DR3 at all separations, though they are still modest in DR2. Even at $G \sim 19$ and wide separations, where the inferred σ_{ϖ} in DR3 are consistent with being accurate or very slightly overestimated, those in DR2 are underestimated by $\sim 5\%$ on average. However, at $11 < G < 13$, the DR3 uncertainties are overestimated more than those in DR2. Given that the reported values of σ_{ϖ} decreased by a factor of two on average between DR2 and DR3 in this magnitude range, the DR3 parallaxes are still “better” on average than those from DR2, but Figure 19 implies that the true gains are somewhat more modest than those reported.

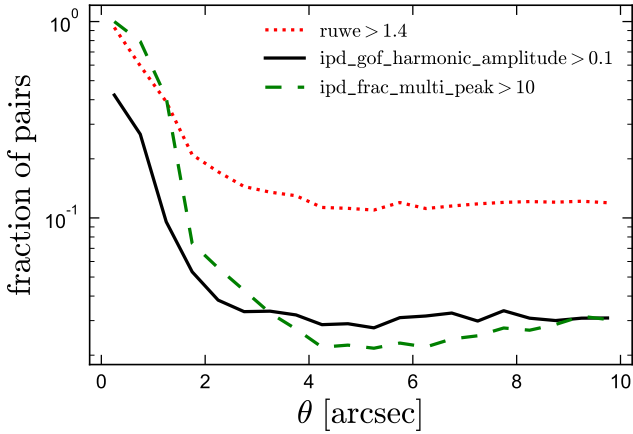


Figure 18. Fraction of pairs in which at least one component fails the indicated quality cut (larger values are indicative of a potentially problematic astrometric solution). All three indicators are significantly enhanced within $\theta \lesssim 2$ arcsec.

This is shown explicitly in Figure 20, which shows the reported and true (inferred) median σ_{ϖ} for DR2 and DR3 as a function of apparent G - magnitude and angular separation. Dashed lines show the median value of σ_{ϖ} in the sample, considering both primaries and secondaries. Solid lines show the result of multiplying these values by the appropriate factors from Figure 19. The corrected median σ_{ϖ} values are at least 30% smaller in DR3 than in DR2 at all magnitudes and separations. For bright stars ($G \lesssim 13$), the gains are generally more than a factor of 2.

5.2 Parallax zeropoint corrections and 5-parameter vs 6-parameter solutions

The *Gaia* eDR3 parallax zeropoint is known to vary with apparent magnitude, color, and ecliptic latitude. We do not attempt to account for this variation when constructing the binary catalog. Because all the sources in the catalog have $\varpi > 1$ mas, the effects of zeropoint corrections, which are typically on the order of 0.02 mas, are modest. Lindegren et al. (2020a) derived an empirical zeropoint for eDR3 using quasars, stars in the LMC, and binaries. Here we investigate whether “correcting” the parallaxes using the prescriptions they provide can reduce the inferred σ_{ϖ} underestimate factors.

Figure 21 shows the inferred σ_{ϖ} underestimate factors with and without the zeropoint correction. We also separately plot binaries in which both components have a 5-parameter astrometric solution and those in which both have a 6-parameter solution. Details about the differences between 5- and 6-parameter solutions are discussed in Lindegren et al. (2020b); 5-parameter solutions are generally more reliable. Binaries with one 5- and one 6-parameter solution are excluded. At $G \lesssim 18$, most sources have 5-parameter solutions. The exception is sources with a close companion ($\theta \lesssim 2$ arcsec), which usually lack both reliable colors and 5-parameter solutions.

At fixed separation and magnitude, the effects of applying the zeropoint correction are encouraging but modest: for widely-separated binaries with $G \approx 13$ and 5-parameter solutions, the inferred σ_{ϖ} underestimate factor decreases from ~ 1.30 to 1.25. Improvements are generally smaller at closer separations and for 6-parameter solutions. The small effect is not unexpected: because the binaries we consider all have small magnitude differences, the zeropoint corrections are

similar for both components, and the parallax *difference* does not change much when the correction is applied. At close separations, where most stars have 6-parameter solutions, the underestimate factors are similar for 5- and 6-parameter solutions. This suggests that the increased uncertainties at close separations are not primarily due to the transition from 5- to 6-parameter solutions.

5.3 Color dependence

The parallax zeropoint is also known to depend somewhat on color, likely due to the color-dependence of the PSF (Lindegren et al. 2020a,b). We therefore investigate how the σ_{ϖ} underestimate factors depend on color in Figure 22, where we separately consider binaries in which both components are blue and those in which both components are red. The color boundary we use, $\nu_{\text{eff}} = 1.4 \mu\text{m}^{-1}$, corresponds to $G_{\text{BP}} - G_{\text{RP}} \approx 1.6$, which is approximately the red limit of the *Gaia* quasar sample. Figure 22 shows that for sources with $G > 16$, the inferred σ_{ϖ} underestimates are somewhat larger for red sources. We verified that there are not significant angular separation differences between the red and blue pairs at fixed magnitude, so color is the most likely driving variable.

5.4 Fitting function to inflate σ_{ϖ}

We fit a function to our inferred σ_{ϖ} inflation factors as a function of G magnitude, which can be used to empirically correct σ_{ϖ} values reported for isolated sources. To derive a correction appropriate for single sources with well-behaved astrometry, we consider binaries with $\theta > 5$ arcsec and $\text{ruwe} < 1.4$ for both components, and we apply the zeropoint correction from Lindegren et al. (2020a) to both components’ parallaxes. Because we observe a general decline in the inflation factor with increasing G and a peak at $G \approx 13$ (Figure 23), we fit a polynomial plus a Gaussian bump:

$$f(G) = A \exp \left[-\frac{(G - G_0)^2}{b^2} \right] + p_0 + p_1 G + p_2 G^2 \quad (16)$$

We find $A = 0.21$, $G_0 = 12.65$, $b = 0.90$, $p_0 = 1.141$, $p_1 = 0.0040$, and $p_2 = -0.00062$. This is also plotted in Figure 23. Multiplying by f will – on average – correct reported σ_{ϖ} values for *Gaia* eDR3 and single-source solutions in DR3. The correction is appropriate for sources with $7 < G < 21$ that have $\text{ruwe} < 1.4$, have no comparably bright sources within a few arcsec, and have already had their parallaxes corrected by the zeropoint from Lindegren et al. (2020a). The effects of having a close companion on σ_{ϖ} likely depend on the brightness contrast and a variety of other factors; a rough estimate of the magnitude of the inflation can be obtained from Figure 16-17. The correction can be reasonably applied to both 5- and 6-parameter astrometric solutions. We do not fit separate corrections for red and blue sources, but we note that at the faint end, the inflation factors are generally somewhat smaller for blue sources.

5.5 Angular correlations in parallaxes

The σ_{ϖ} inflation factors inferred in this work and predicted by Equation 16 should be interpreted as lower limits. *Gaia* eDR3 parallaxes are subject to systematic trends on degree scales (and larger) due to the scanning law (e.g. Fabricius et al. 2020; Lindegren et al. 2020b). The angular separations of most of the binaries in the catalog are significantly smaller than this (Figure 7), so the “local” parallax zeropoint for the two stars is usually very similar. These local positional




Figure 19. Same as Figures 16 and 17, but now comparing results obtained with *Gaia* DR2 vs DR3 parallaxes. All binaries considered have $\text{ruwe} < 1.4$ for both components in both DR2 and DR3. At $G \gtrsim 14$, the uncertainties are underestimated less in DR3 than in DR2. DR3 uncertainties are overestimated somewhat more severely at $11 < G < 13$.




Figure 20. Median reported (dashed) and corrected (solid) parallax uncertainty as a function of G magnitude, for *Gaia* DR3 (black) and DR2 (red). The corrected values are obtained by multiplying the reported values by the factors shown in Figure 19, which are empirically determined from the reported parallax differences of the components of wide binaries. Panels show different bins of angular separation. Only stars with $\text{ruwe} < 1.4$ in both data releases are considered. The rightmost panel is appropriate for single stars outside of crowded fields.

variations in the zeropoint are not accounted for in the correction from [Lindegren et al. \(2020a\)](#) and also inflate the effective parallax uncertainties. With a typical scale of about $10 \mu\text{as}$, these may contribute significantly to the uncertainties at $G \lesssim 13$.

A contemporaneous study by [Zinn \(2021\)](#) validated *Gaia* eDR3 parallaxes and their uncertainties using bright giants in the Kepler field ($9 \lesssim G \lesssim 13$) with independent distance estimates from asteroseismology. They tested our parallax uncertainty inflation function (Equation 16) and found it to perform well; i.e., no further uncertainty inflation was required after it was applied. The angular size of the Kepler field (~ 10 deg) is larger than the scale on which the strongest angular correlations in *Gaia* eDR3 are manifest (e.g. [Lindegren et al. 2020b](#), their Figure 14). This suggests that any additional uncertainty inflation required due to angular correlations in the zeropoint is modest.

5.6 Comparison to other work

The reliability of parallax uncertainties reported in *Gaia* eDR3 has been investigated by several other works. [Fabricius et al. \(2020, their Figure 19\)](#) used the dispersion in parallaxes reported for distant objects – quasars, stars in the LMC, and stars in dwarf spheroidal (dSph) galaxies, which should all have negligible true parallaxes – to estimate σ_π inflation factors as a function of G magnitude. Their results from quasars and dSph stars, which are only available at the faint end ($G \gtrsim 16$) are broadly consistent with our results. In the LMC, they infer inflation factors that are larger at fixed magnitude than our results at wide separations, or their results for quasars and dSph at the same magnitude. This discrepancy is very likely a result of crowding: the source density in the LMC is large enough that a significant fraction of sources have another source within a few arcsec, and as we have shown (e.g. Figure 19), sources with companions within a few arcsec have more severely underestimated parallaxes. This likely also

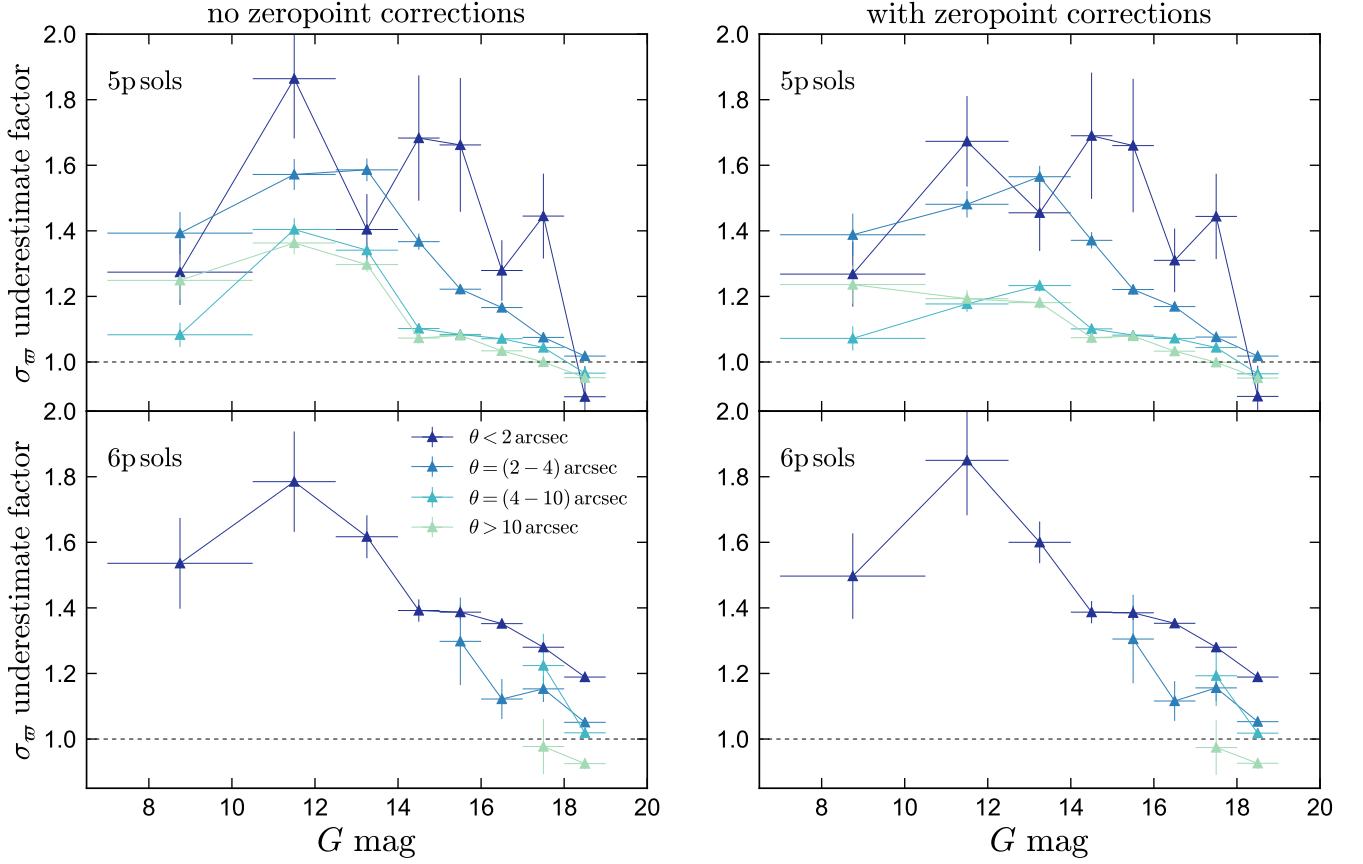


Figure 21. σ_w underestimate factors (similar to Figure 16) for binaries in which both components have five-parameter solutions (top) or both have six-parameter solutions (bottom). Left column uses the *Gaia* eDR3 parallaxes directly; right column applies the magnitude-, color-, and position-dependent parallax zeropoint correction from [Lindegren et al. \(2020a\)](#). For bright sources, only binaries with close angular separations have six-parameter solutions. The inferred σ_w underestimate factors depend mainly on separation and magnitude; they are similar for 5- vs 6-parameter solutions. Applying the zeropoint correction reduces the inferred overestimate factor slightly for bright sources.




Figure 22. Similar to Figure 16, but with binaries divided by color. Only binaries in which both components have $\text{ruwe} < 1.4$ and the two components fall in the same color bin are included. ν_{eff} refers to the parameters `nu_eff_used_in_astrometry` and `pseudocolour`, respectively, for 5- and 6-parameter solutions. The boundary of $\nu_{\text{eff}} = 1.4 \mu\text{m}^{-1}$ corresponds approximately to $G_{\text{BP}} - G_{\text{RP}} = 1.6$. Most quasars, white dwarfs, and MS stars with $M \geq 0.6 M_{\odot}$ are bluer than this; lower-mass MS stars are redder. Left and middle panels show all separations; right panel shows wide separations. At the faint end, σ_w overestimates are somewhat larger for red sources.

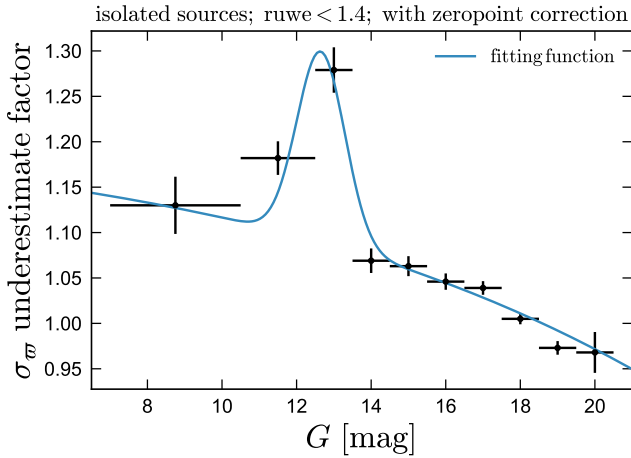


Figure 23. Parallax uncertainty inflation factors inferred from widely-separated pairs with $\text{ruwe} < 1.4$ and parallaxes corrected by the zeropoint from Lindegren et al. (2020a). This is the σ_w underestimate factor for typical single sources with well-behaved astrometry. A fitting function (Equation 16) is provided.

explains why Maíz Apellániz et al. (2021) found somewhat larger uncertainty inflation factors in globular clusters than we do with widely separated binaries. In the Kepler field, Zinn (2021) found a mean uncertainty estimate of $(22 \pm 6)\%$ for sources with $9 \lesssim G \lesssim 13$ (with most sources at the faint end of this range); this is consistent with our results.

6 SUMMARY AND DISCUSSION

We have constructed a catalog of 1.2 million high-confidence, spatially resolved wide binaries using *Gaia* eDR3 and have used them to empirically validate the reported parallax uncertainties in *Gaia* eDR3. Overall, the results are very encouraging: outside of crowded regions (i.e., for stars with no comparably bright companion within a few arcseconds), parallax uncertainties for sources with well-behaved astrometric fits are underestimated by *at most* 30% (at $G \approx 13$), and by considerably less at fainter magnitudes. Our main results are as follows:

(i) : *Catalog description*: The full catalog contains 1.2 million high-confidence binaries, including 15,982 WD+MS binaries and 1362 WD+WD binaries (Figures 7 and 8), and $\sim 20,000$ binaries containing giants and subgiants. The binaries span a projected separation range of a few au to 1 pc, have heliocentric distance up to 1 kpc, and include both (kinematic) disk and halo binaries. All binaries have reasonably precise astrometry, with $\varpi/\sigma_\varpi > 5$ for both components. The full catalog of 1.8 million binary candidates become dominated by chance alignments at $s \gtrsim 30,000$ au, but high-quality subsets can be selected that are relatively pure out to separations as large as 1 pc (Figure 4). The catalog builds on previous efforts to identify binaries using *Gaia* DR2, expanding the sample of known high-confidence binaries by a factor of 4 (Figure 10). This increase in sample size owes partly to the higher astrometric precision provided by *Gaia* eDR3, and partly to improvements in the binary identification and vetting strategy.

(ii) : *Quantifying and controlling chance alignments*: We estimate the contamination rate from chance alignments using two approaches: a mock catalog that does not contain any true binaries,

and a version of the *Gaia* catalog in which stars have been artificially shifted from their true positions, removing real binaries but preserving chance-alignment statistics (Figure 3). Both approaches show that the full catalog has high purity at $s < 10,000$ but becomes dominated by chance alignments at $s \gtrsim 30,000$ au. Using the shifted chance alignment catalog, we show how one can select subsamples that have lower contamination rates, including some that are pure out to 1 pc (Figure 4). We also include in the catalog an estimate of the probability that each binary candidate is a chance alignment; this is constructed empirically from the distribution of known chance alignments in a seven-dimensional space of observables (Figure 5 and Appendix A). We use radial velocities from *Gaia* (Figure 6) and LAMOST (Figure 14) to validate these probabilities.

(iii) *Orbital velocities*: The high-precision of *Gaia* astrometry makes obvious the plane-of-the-sky velocity difference of the components of binaries due to orbital motion (Figure 12). About 200,000 binaries in the catalog have sufficiently accurate astrometry that the plane-of-the-sky velocity difference between the components can be measured with accuracy $\sigma_{\Delta V} < 100 \text{ m s}^{-1}$. A Keplerian decline in the velocity difference, $\Delta V \propto s^{-1/2}$, is visible out to $s \sim 20,000$ au, where projection effects become important. The requirement of proper motions consistent with Keplerian orbits excludes a significant fraction of hierarchical triples and higher-order multiples from the catalog (Figure 13).

(iv) *Validation of Gaia DR3 parallax uncertainties*: We use the sample of high-confidence binaries to validate the published parallax uncertainties included in *Gaia* eDR3. This analysis makes use of the fact that the two stars in a binary have essentially the same distance and thus should generally have reported parallaxes that are consistent within their uncertainties (Figure 15). We find that the published uncertainties are accurate for faint stars ($G \gtrsim 18$) that have well-behaved astrometric solutions and do not have a companion within a few arcseconds (Figure 16). They are underestimated somewhat for brighter stars, particularly in the range of $11 < G < 13$, where the published uncertainties should be multiplied by a factor of 1.3 on average. The degree to which uncertainties are underestimated is larger for sources with large ruwe , $\text{ipd_gof_harmonic_amplitude}$, and $\text{ipd_frac_multi_peak}$ (Figure 16 and 17), and is larger for red sources than blue sources (Figure 22). The reported parallax uncertainties are generally more reliable in *Gaia* eDR3 than they were in DR2, except at $11 < G < 13$ (Figure 19). In an absolute sense, the σ_w values improved by at least 30% from DR2 to eDR3, at all magnitudes and separations (Figure 20). We provide an empirical fitting function to correct reported σ_w values (Figure 23).

Parallax uncertainties are underestimated more for binaries with angular separations less than a few arcseconds. Sources with resolved close companions are more likely to have high ruwe and IPD diagnostics related to binarity (Figure 18), but the underestimates of σ_w is enhanced at close separations even for pairs in which both components have low ruwe and IPD flags. This is true both for sources with 5- and 6-parameter solutions (Figure 21).

6.1 Scientific uses for the catalog

This paper was primarily concerned with assembling the wide binary catalog. Here we note a few possible uses for the sample, which will be pursued in future work.

- *Calibrating stellar ages* : A useful property of wide binaries is that the two stars have basically the same total age, but can have different evolutionary states. If the age of one component can be constrained (e.g., because it is a WD, a subgiant, or a giant with as-

teroseismic mass constraints), that age constraint can be transferred to the companion (e.g. Chanamé & Ramírez 2012; Fouesneau et al. 2019; Qiu et al. 2020). This in turn can be used to calibrate more poorly-understood age indicators for MS stars, such as gyrochronology, stellar activity, and X-ray luminosity (e.g. Soderblom 2010).

- *The initial-final mass relation:* The WD+WD sample will be useful for constraining the initial-final mass relation (IFMR) for WDs: the masses and cooling ages of both WDs can (often) be well-constrained from photometry, and the IFMR can be constrained by the fact that both WDs have the same total age (and presumably, follow the same IFMR; e.g., Andrews et al. 2015). The subset of the WD+MS sample in which the age of the non-WD component can be constrained independently will also be useful for constraining the IFMR (e.g. Catalán et al. 2008).

- *WD masses from gravitational redshift:* Because the WD and MS star in a wide binary have essentially the same RV, the difference in their apparent RVs is due primarily to the WD’s gravitational redshift (which is typically 20–100 km s⁻¹). This provides a useful way of measuring WD mass that are essentially model-independent if the distance is well-constrained (e.g. Koester 1987). Masses from gravitational redshift will be particularly useful for measuring the mass distributions of WDs with rare spectral types and poorly-understood formation histories, such as the “Q-branch” WDs revealed by the *Gaia* CMD (Bergeron et al. 2019; Cheng et al. 2020).

- *Abundances for WD progenitors:* a MS companion provides a window into the chemical abundances of a WD’s progenitor, which are otherwise inaccessible. Among other applications, this provides an avenue to compare the primordial and final abundances of disintegrating planets around polluted WDs, which represent a significant fraction of the WD population (Koester et al. 2014; Farihi 2016).

- *Wide binary spin alignment:* An open question in binary star formation is how aligned the spins vectors of binaries are, and how this varies with separation (e.g. Justesen & Albrecht 2020). Of particular interest is whether the spins of excess “twin” binaries, which may have formed in circumbinary disks, are more aligned than those of non-twins at the same separation (El-Badry et al. 2019). Spin inclinations can be measured by combining a spectroscopic measurement of $v \sin i$ with a rotation period measured from spots (which requires a light curve) and a radius from parallax and temperature. About 6,000 of the high-confidence binaries in the catalog have both components with $G < 14$ and angular separations $\theta > 30$ arcseconds. These are ideal for follow-up study with TESS, because they are separated widely enough that high-quality light curves are available for both components separately. About 39,000 binaries in the catalog are in the K2 fields, and 5,000 are in the Kepler field.

- *Calibration of spectroscopic surveys:* The surface abundances of stars in wide binaries are generally very similar (e.g. Hawkins et al. 2020). This enables diagnosis of systematics in the abundances reported by surveys (e.g. Figure 14). Similar analyses can be done with stars in clusters, but binaries are more abundant and populate abundance space more densely than clusters.

- *Dynamical probes:* At separations wider than about 10,000 au, wide binaries are susceptible to dynamical disruption through gravitational encounters with other stars, compact objects, or molecular clouds (e.g. Weinberg et al. 1987). This makes the wide binary separation distribution a sensitive probe of the population of possible perturbers (e.g. Yoo et al. 2004; Tian et al. 2020).

6.2 Constraining the parallax zeropoint with binaries

A further application of the *Gaia* binary sample, which we have not explored in this work, is calibration of the parallax zeropoint.

We have intentionally limited our analysis to binaries in which both stars have almost the same magnitude, (and, because almost all the stars are on the main sequence, the same color). This avoids complication arising from the magnitude- and color-dependence of the parallax zeropoint, which should be nearly the same for both stars. If we were to consider binaries with substantially different primary and secondary magnitudes, magnitude-dependence of the zeropoint would be manifest as a shift in the *mean* signed parallax difference: that is, the distributions in Figure 15 would no longer be centered on zero. If the absolute zeropoint at a particular magnitude and color can be pinned down from external data (e.g. quasars at the faint/blue end), binaries then allow for determination of the zeropoint at all other magnitudes and colors. Some analysis along these lines was carried out by Lindegren et al. (2020a) and Fabricius et al. (2020).

A challenge to carrying out this type of calibration with our current catalog is Lutz-Kelker bias (Lutz & Kelker 1973): because the fainter secondaries have larger parallax errors, their parallaxes in our sample will on average be overestimated more than those of the primaries. That is, in the absence of any magnitude-dependence of the zeropoint, the mean value of $\varpi_1 - \varpi_2$ would be negative. This bias must be eliminated or accounted for in order for reliable determination of the zeropoint from binaries to be feasible. The most straightforward path forward is likely to select binaries without explicit cuts on parallax.

6.3 Astrometric acceleration

We have shown that parallax uncertainties are more severely underestimated at close angular separations (e.g. Figure 16), and that a larger fraction of close binaries have $\text{ruwe} > 1.4$ for at least one component (Figure 18). Here we consider whether this is likely due to actual astrometric acceleration, or other issues.

Whether orbital acceleration of a binary is detectable depends on a variety of factors, including the eccentricity, orientation, and phase of the orbit. Here we derive a crude estimate. We consider a face-on circular orbit with $M_1 + M_2 = 1M_\odot$ and $M_1 \gg M_2$, with an angular separation θ viewed at a distance d . The semi-major axis is $a = 1 \text{ au} \times (d/\text{pc}) (\theta/\text{arcsec})$, and the orbital period is $P = 1 \text{ yr} \times (d/\text{pc})^{3/2} (\theta/\text{arcsec})^{3/2}$. During the 34-month baseline of *Gaia* DR3, the azimuthal angle swept out by the secondary is

$$\phi = 2\pi \times ((34 \text{ months}) / P) \quad (17)$$

$$= 0.018 \text{ radians} \times (d/(100 \text{ pc}))^{-3/2} (\theta/\text{arcsec})^{-3/2}, \quad (18)$$

where $\phi = 2\pi$ would signify a full orbit.

In the limit of small ϕ , the total orbital motion that is perpendicular to the instantaneous proper motion vector at the first observation is $\vartheta_\perp \approx \frac{1}{2}\theta\phi^2$, from Taylor expanding $x = \theta \cos \phi$. The total perpendicular deviation from linear motion is thus

$$\vartheta_\perp \approx 0.16 \text{ mas} \times \left(\frac{d}{100 \text{ pc}}\right)^{-3} \left(\frac{\theta}{\text{arcsec}}\right)^{-2} \quad (19)$$

$$\approx 0.16 \text{ mas} \times \left(\frac{d}{100 \text{ pc}}\right) \left(\frac{s}{100 \text{ au}}\right)^{-2}. \quad (20)$$

The deviation from the *best-fit* single-star orbit will likely be a factor of a few smaller than this.

To determine whether orbital acceleration is plausibly detectable, this quantity can be compared to the typical astrometric precision (e.g. Figure 20). At the typical distance of binaries in the catalog, $d \approx 500 \text{ pc}$, the predicted deviation for $\theta = 1 \text{ arcsec}$ is on the order of 0.001 mas, well below the sensitivity of *Gaia* eDR3. Astrometric acceleration due to orbital motion is thus not expected to be detectable for the large majority of binaries in our catalog, and it is therefore

likely that the more strongly underestimated σ_{ϖ} at close separations is primarily due to other issues, such as centroiding errors or some scans being attributed to the wrong component.

Astrometric acceleration should, in principle, be non-negligible for the nearest and closest binaries in the catalog. Considering only binaries with $s < 100$ au, the median deviation predicted by Equation 20 is 0.3 mas, which is larger than σ_{ϖ} for the majority of that sample. To investigate whether there is evidence of acceleration in our sample at close separations, we compared the DR2 and eDR3 proper motions of both components, under the assumption that acceleration should manifest as a change in mean proper motion from epoch 2015.5 to 2016.0 (e.g. Kervella et al. 2019). It is important to note that the coordinate systems of DR2 and eDR3 are not identical. An ad-hoc correction was applied to the eDR3 coordinate frame to remove a ~ 0.1 mas yr⁻¹ rotation that was present in the coordinate system for bright stars in DR2 (Lindegren et al. 2020b). Properly aligning the coordinate systems between the two releases is non-trivial (e.g. Brandt 2018). For this reason, and because proper motion uncertainties, like parallax uncertainties, are likely underestimated somewhat, it is beyond the scope of our investigation to determine which proper motion differences are significant. Instead, we simply consider how the fraction of sources with proper motion differences above a particular threshold depends on separation. We find that the fraction of binaries with inconsistent proper motions for one or both components is strongly enhanced at close angular separations, as would be expected in the presence of accelerations. For the full binary catalog, the fraction of pairs that have at least one component with DR2 and eDR3 proper motions inconsistent within 3σ is 18% at $\theta > 4$ arcsec, but 50% at $\theta < 1$ arcsec, and 85% at $\theta < 0.5$ arcsec.

However, the fraction of sources with inconsistent proper motions is significantly enhanced even at $d > 500$ pc, where Equation 20 suggests that any perpendicular acceleration should be negligible. Indeed, at fixed magnitude, the fraction of sources with inconsistent proper motions depends primarily on angular, not physical, separation. This suggests spurious astrometry (due to bias in the image parameter determination or source misidentification) for sources with close companions is the primary cause for the apparent acceleration. Although Equation 20 suggests that acceleration should often be detectable in our sample at $s < 50$ au, the expected sensitivity is not yet realized there due to problematic astrometry for barely-resolved sources. The detectability of accelerations with *Gaia* was also investigated by Belokurov et al. (2020). They found that while ruwe is often enhanced in close binaries, the enhancement can be reliably tied to orbital motion only in the regime where a significant (order unity) fraction of the orbit is covered by the *Gaia* time baseline.

7 CATALOG DESCRIPTION

The full binary catalog will be hosted at CDS. It can also be accessed at <https://zenodo.org/record/4435257>. All columns in the `gaiaedr3.gaia_source` catalog are copied over for both components. We also include the columns `source_id`, `parallax`, `parallax_error`, `pmra`, `pmdec`, `pmra_error`, `pmdec_error`, and `ruwe` from *Gaia* DR2 for both components; these have the prefix `dr2_`.

Columns ending in “1” and “2” refer to the primary and secondary component, respectively. The primary is always the component with the brighter *G* magnitude. We also include columns `pairdistance` (angular separation θ , in degrees), `sep_AU` (projected separation s , in au), `R_chance_align` (\mathcal{R} ; Equation 8), and `binary_type` (e.g. MSMS, WDMS, etc.; see Table 1). The ordering in `binary_type`

does not account for primary/secondary designations; i.e., all binaries containing a WD and a MS star are designated WDMS, irrespective of whether it is the WD or MS component that is brighter.

The shifted chance alignment catalog is also available. It contains the same columns as the binary candidate catalog, except the *Gaia* DR2 columns. Because one component of each pair has been shifted from its true position in the `gaia_source` catalog, the `ra` and `dec` columns in it do not match those reported in the `gaia_source` catalog.

ACKNOWLEDGEMENTS

We thank the anonymous referee for a constructive report, Jackie Faherty for help creating visualizations of the catalog, and Eliot Quataert, Dan Weisz, Jan Rybizki, and Anthony Brown for helpful comments. We acknowledge earlier discussions with Tim Brandt that proved seminal for this paper. We are grateful to Geoff Tabin and In-Hei Hahn for their continued hospitality during the writing of this manuscript. KE was supported by an NSF graduate research fellowship and a Hellman fellowship from UC Berkeley. T.M.H. acknowledges support from the National Science Foundation under Grant No. AST-1908119.

This project was developed in part during the 2020 virtual eDR3 Unboxing *Gaia* Sprint. This work has made use of data from the European Space Agency (ESA) mission *Gaia* (<https://www.cosmos.esa.int/gaia>), processed by the *Gaia* Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the DPAC has been provided by national institutions, in particular the institutions participating in the *Gaia* Multilateral Agreement. Guoshoujing Telescope (the Large Sky Area Multi-Object Fiber Spectroscopic Telescope LAMOST) is a National Major Scientific Project built by the Chinese Academy of Sciences. Funding for the project has been provided by the National Development and Reform Commission. LAMOST is operated and managed by the National Astronomical Observatories, Chinese Academy of Sciences. This research made use of Astropy,⁵ a community-developed core Python package for Astronomy (Astropy Collaboration et al. 2013, 2018). This research made use of the cross-match service provided by CDS, Strasbourg.

DATA AVAILABILITY

All the data used in this paper is publicly available. The *Gaia* data can be retrieved through the *Gaia* archive (<https://gea.esac.esa.int/archive>), and the LAMOST data are available at <http://dr6.lamost.org>. The binary catalog and code to produce it can be found at <https://zenodo.org/record/4435257>.

REFERENCES

- Abrahams E. S., Bloom J. S., Mowlavi N., Szkody P., Rix H.-W., Ventura J.-P., Brink T. G., Filippenko A. V., 2020, arXiv e-prints, p. [arXiv:2011.12253](https://arxiv.org/abs/2011.12253)
- Andrews J. J., Agüeros M. A., Gianninas A., Kilic M., Dhital S., Anderson S. F., 2015, *ApJ*, 815, 63
- Andrews J. J., Chanamé J., Agüeros M. A., 2017, *MNRAS*, 472, 675
- Arenou F., et al., 2018, *A&A*, 616, A17
- Astropy Collaboration et al., 2013, *A&A*, 558, A33
- Astropy Collaboration et al., 2018, *aj*, 156, 123

⁵ <http://www.astropy.org>

- Bédard A., Bergeron P., Brassard P., Fontaine G., 2020, *ApJ*, **901**, 93
- Belokurov V., et al., 2020, *MNRAS*, **496**, 1922
- Bergeron P., et al., 2011, *ApJ*, **737**, 28
- Bergeron P., Dufour P., Fontaine G., Coutu S., Blouin S., Genest-Beaulieu C., Bédard A., Rolland B., 2019, *ApJ*, **876**, 67
- Binney J., Tremaine S., 2008, Galactic Dynamics: Second Edition
- Brandeker A., Cataldi G., 2019, *A&A*, **621**, A86
- Brandt T. D., 2018, *ApJS*, **239**, 31
- Buder S., et al., 2020, arXiv e-prints, p. arXiv:2011.02505
- Catalán S., Isern J., García-Berro E., Ribas I., Allende Prieto C., Bonanos A. Z., 2008, *A&A*, **477**, 213
- Chanamé J., Gould A., 2004, *ApJ*, **601**, 289
- Chanamé J., Ramírez I., 2012, *ApJ*, **746**, 102
- Cheng S., Cummings J. D., Ménard B., Toonen S., 2020, *ApJ*, **891**, 160
- Clarke C. J., 2020, *MNRAS*, **491**, L72
- Close L. M., Richer H. B., Crabtree D. R., 1990, *AJ*, **100**, 1968
- Coronado J., Rix H.-W., Trick W. H., El-Badry K., Rybizki J., Xiang M., 2020, *MNRAS*, **495**, 4098
- Cui X.-Q., et al., 2012, *Research in Astronomy and Astrophysics*, **12**, 1197
- 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
- Duchêne G., Kraus A., 2013, *ARA&A*, **51**, 269
- El-Badry K., 2019, *MNRAS*, **482**, 5018
- El-Badry K., Rix H.-W., 2018, *MNRAS*, **480**, 4884
- El-Badry K., Rix H.-W., 2019, *MNRAS*, **482**, L139
- El-Badry K., Rix H.-W., Ting Y.-S., Weisz D. R., Bergemann M., Cargile P., Conroy C., Eilers A.-C., 2018, *MNRAS*, **473**, 5043
- El-Badry K., Rix H.-W., Tian H., Duchêne G., Moe M., 2019, *MNRAS*, **489**, 5822
- Evans D. W., et al., 2018, *A&A*, **616**, A4
- Fabricius C., et al., 2020, arXiv e-prints, p. arXiv:2012.06242
- Farihi J., 2016, *New Astron. Rev.*, **71**, 9
- Fouesneau M., Rix H.-W., von Hippel T., Hogg D. W., Tian H., 2019, *ApJ*, **870**, 9
- Gaia Collaboration et al., 2016, *A&A*, **595**, A1
- Gaia Collaboration et al., 2018, *A&A*, **616**, A1
- Gaia Collaboration Brown A. G. A., Vallenari A., Prusti T., de Bruijne J. H. J., Babusiaux C., Biermann M., 2020a, arXiv e-prints, p. arXiv:2012.01533
- Gaia Collaboration et al., 2020b, arXiv e-prints, p. arXiv:2012.02061
- Hartman Z. D., Lépine S., 2020, *ApJS*, **247**, 66
- Hawkins K., et al., 2020, *MNRAS*, **492**, 1164
- Herschel W., 1782, Philosophical Transactions of the Royal Society of London Series I, **72**, 112
- Herschel W., 1803, Philosophical Transactions of the Royal Society of London Series I, **93**, 339
- Holberg J. B., Bergeron P., 2006, *AJ*, **132**, 1221
- Hwang H.-C., Ting Y.-S., Schlaufman K. C., Zakamska N. L., Wyse R. F. G., 2020, *MNRAS*,
- Jiang Y.-F., Tremaine S., 2010, *MNRAS*, **401**, 977
- Jurić M., et al., 2008, *ApJ*, **673**, 864
- Justesen A. B., Albrecht S., 2020, *A&A*, **642**, A212
- Katz D., et al., 2019, *A&A*, **622**, A205
- Kervella P., Arenou F., Mignard F., Thévenin F., 2019, *A&A*, **623**, A72
- Koester D., 1987, *ApJ*, **322**, 852
- Koester D., Gänsicke B. T., Farihi J., 2014, *A&A*, **566**, A34
- Kowalski P. M., Saumon D., 2006, *ApJ*, **651**, L137
- Lépine S., Bongiorno B., 2007, *AJ*, **133**, 889
- Lindgren L., et al., 2018, *A&A*, **616**, A2
- Lindgren L., et al., 2020a, arXiv e-prints, p. arXiv:2012.01742
- Lindgren L., et al., 2020b, arXiv e-prints, p. arXiv:2012.03380
- Lutz T. E., Kelker D. H., 1973, *PASP*, **85**, 573
- Luyten W. J., 1971, *Ap&SS*, **11**, 49
- Luyten W. J., 1979, Proper Motion Survey, University of Minnesota, 52
- Maíz Apellániz J., Pantaleoni González M., Barbá R. H., 2021, arXiv e-prints, p. arXiv:2101.10206
- Moe M., Di Stefano R., 2017, *ApJS*, **230**, 15
- Oh S., Price-Whelan A. M., Hogg D. W., Morton T. D., Spergel D. N., 2017, *AJ*, **153**, 257
- Perpinyà-Vallès M., Rebassa-Mansergas A., Gänsicke B. T., Toonen S., Hermes J. J., Gentile Fusillo N. P., Tremblay P. E., 2019, *MNRAS*, **483**, 901
- Pittordis C., Sutherland W., 2019, *MNRAS*, **488**, 4740
- Qiu D., Tian H.-J., Wang X.-D., Nie J.-L., von Hippel T., Liu G.-C., Fouesneau M., Rix H.-W., 2020, arXiv e-prints, p. arXiv:2012.04890
- Raghavan D., et al., 2010, *ApJS*, **190**, 1
- Riello M., et al., 2020, arXiv e-prints, p. arXiv:2012.01916
- Robin A. C., Reylé C., Derrière S., Picaud S., 2003, *A&A*, **409**, 523
- Rowell N., et al., 2020, arXiv e-prints, p. arXiv:2012.02069
- Rybizki J., et al., 2020, *PASP*, **132**, 074501
- Salim S., Gould A., 2003, *ApJ*, **582**, 1011
- Sartoretti P., et al., 2018, *A&A*, **616**, A6
- Sharma S., Bland-Hawthorn J., Johnston K. V., Binney J., 2011, *ApJ*, **730**, 3
- Soderblom D. R., 2010, *ARA&A*, **48**, 581
- Tian H.-J., El-Badry K., Rix H.-W., Gould A., 2020, *ApJS*, **246**, 4
- Tokovinin A., 2014, *AJ*, **147**, 87
- Tokovinin A., Smekhov M. G., 2002, *A&A*, **382**, 118
- Tokovinin A., Hartung M., Hayward T. L., 2010, *AJ*, **140**, 510
- Tremblay P.-E., Bergeron P., Gianninas A., 2011, *ApJ*, **730**, 128
- Weinberg M. D., Shapiro S. L., Wasserman I., 1987, *ApJ*, **312**, 367
- Winters J. G., Charbonneau D., Henry T. J., Irwin J. M., Jao W.-C., Riedel A. R., Slatten K., 2020, arXiv e-prints, p. arXiv:2011.09409
- Yoo J., Chanamé J., Gould A., 2004, *ApJ*, **601**, 311
- Zinn J. C., 2021, arXiv e-prints, p. arXiv:2101.07252

APPENDIX A: CHANCE ALIGNMENT PROBABILITIES

We estimate the local density in parameter space of binary candidates and known chance alignments from the shifted catalog using a Gaussian kernel density estimate (KDE). The parameters (“features”) we use are listed in Table A1. Most of them are described in Section 2.

We also add a measure of the local sky density, Σ_{18} . This represents the number of sources per square degree that (a) pass the cuts of our initial query (Section 2) and (b) are brighter than $G = 18$. We calculate the value of Σ_{18} around every binary candidate, counting the number of sources within 1 degree of the primary and dividing by π . Values of Σ_{18} range from 280 toward the Galactic poles to 8700 toward the Galactic center. A significant fraction of the sources toward the Galactic center are likely background stars that are not actually within the 1 kpc search volume.

We rescale the features so that they have similar dynamic range. This is accomplished by applying a few simple functions, which are listed in the “scaled parameter” column of Table A1. The rescaled parameters all have a dynamic range of about 4. We then calculate a 7-dimensional Gaussian KDE using a bandwidth $\sigma = 0.2$.

The distance metric in this parameter space is somewhat ill-defined due to the different units and distributions of the features. We nevertheless proceed boldly, making no claim that the set of features, rescalings, or the choice of kernel are optimal. Our choices are designed to make the kernel (a) narrow enough that it does not smooth over the sharpest features in the data, and (b) wide enough to prevent density peaks around individual, discrete binaries (overfitting).

When calculating the KDE for the binary candidates, we use a leave-10% out method wherein the density at the positions of 10% of the binary candidates is evaluated using a KDE constructed from the other 90%. To minimize discreteness noise in the chance-alignment KDE, we produce 30 different realizations of the shifted chance alignment catalog, shifting the declination of each star by a random variable $\mathcal{U}(-0.5, 0.5)$ degrees for each realization. We combine the

realizations when calculating the KDE, and then divide the calculated density by 30 to reflect the number of pairs in a single realization.

Figure A1 compares the distributions of shifted chance alignments, all binary candidates, and candidates with $\mathcal{R} < 0.1$, in the space of features used for the KDE. It is clear that there are two modes in the binary candidate distribution, only one of which has a corresponding population in the chance alignment catalog. The clearest divisions between chance-alignments and binaries are in the dimensions of angular separation and proper motion difference (see also Figure 12), but their distributions also differ in other features. For example, a binary candidate is more likely to be genuine if the parallax errors are small (low $\sigma_{\Delta\varpi}$), or if it is found in a region of low stellar density (low Σ_{18}).

Figure A2 shows the ratio of the number of chance alignments (from the shifted catalog) that have a given \mathcal{R} value to the number of binary candidates with similar \mathcal{R} . If \mathcal{R} is interpreted as the probability that a candidate is a chance alignment, one would expect this ratio to follow the one-to-one line. It does indeed fall close to the one-to-one line (dashed) but with some deviations, likely due to over-smoothing of the KDE. This figure suggests, for example, that about 6% of binary candidates with $\mathcal{R} \sim 0.1$ are chance alignments, implying that chance alignment probabilities inferred when \mathcal{R} is interpreted as a probability are conservative.

A small fraction of binary candidates in the catalog have extremely small \mathcal{R} values; i.e., 0.5% have $\log(\mathcal{R}) < -10$, and 0.09% have $\log(\mathcal{R}) < -20$. These are primarily at close separations, where the chance alignment probability is indeed very low, but in this regime the \mathcal{R} values should not be interpreted as probabilities due to the finite size of the shifted chance alignment catalog.

A1 Sources with spurious astrometry

A non-negligible fraction of sources in *Gaia* eDR3 have spurious astrometric solutions, meaning that they have large reported parallaxes and small reported uncertainties, but the parallaxes are significantly in error. The types of problems that can cause spurious solutions – typically crowding and marginally resolved sources – are generally equally likely to produce positive and negative parallaxes. The impact of spurious sources on our sample can thus be assessed by considering sources with significant negative parallaxes.

To this end, we repeat our initial ADQL query (Section 2) but require `parallax < -1` and `parallax_over_error < -5`. This yields 2,877,625 sources, implying that about 4.5% of the sources returned by the initial query have spurious solutions. We add the sources to our initial sample, treating their parallaxes as if they were positive. We then repeat the neighbor-counting procedure described in Section 2.1 for these sources, again removing objects with more than 30 neighbors. Of the 2,877,625 known spurious sources, only 380,379 (13%) survive this cut. That is, spurious sources are overwhelmingly found in regions of high source density, and a majority of them are removed by the first pass of cleaning.

We then carried out the full catalog construction procedure, now operating on an input sample that includes the initially selected sources as well as the known spurious sources with negative parallaxes, where the sign of the parallax is inverted for the known spurious sources. This yielded 15,852 candidate pairs in which one component is from the spurious sample. As expected, these pairs are concentrated at large separations; only 187 (593) have projected separations $s < 10,000$ au ($s < 30,000$ au). Finally, we repeat the calculation of \mathcal{R} on the candidates in which at least one component is known to be spurious, yielding 133 pairs with $\mathcal{R} < 0.1$. This implies

that about 1 in 10,000 binary candidates with $\mathcal{R} < 0.1$ contains a source with a spurious parallax as defined here.

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

Parameter	units	scaled parameter	(1,99)% range	scaled (1,99)% range	Description
θ	arcsec	$\log \theta$	(0.7, 607)	(-0.16, 2.78)	angular separation
ϖ_1	mas	$4/\varpi_1$	(1.02, 10.7)	(0.37, 3.93)	parallax (primary)
$\sigma_{\Delta\varpi}$	mas	$4\sigma_{\Delta\varpi}$	(0.02, 0.74)	(0.08, 2.96)	parallax difference error
Σ_{18}	deg^{-2}	$4 \log(\Sigma_{18})$	(345, 2790)	(10.15, 13.78)	$G < 18$ local source density
$v_{\perp,1}$	km s^{-1}	$v_{\perp,1}/50$	(3.4, 121)	(0.07, 2.43)	tangential velocity (primary)
$\Delta\varpi/\sigma_{\Delta\varpi}$	–	$ \Delta\varpi /\sigma_{\Delta\varpi}$	(-3.24, 3.30)	(0.02, 3.95)	normalized parallax difference
$(\Delta\mu - \Delta\mu_{\text{orbit}})/\sigma_{\Delta\mu}$	–	$2\text{erf}[(\Delta\mu - \Delta\mu_{\text{orbit}})/\sigma_{\Delta\mu}]$	(-88, 1.97)	(-2, 1.99)	scaled proper motion difference

Table A1. Features used by our Gaussian KDE in computing the local density of binaries and chance alignments. The feature vector is the “scaled parameter” column, in which all variables have been rescaled to have comparable dynamic range. We list the middle-98% ranges of both the raw and scaled parameters.

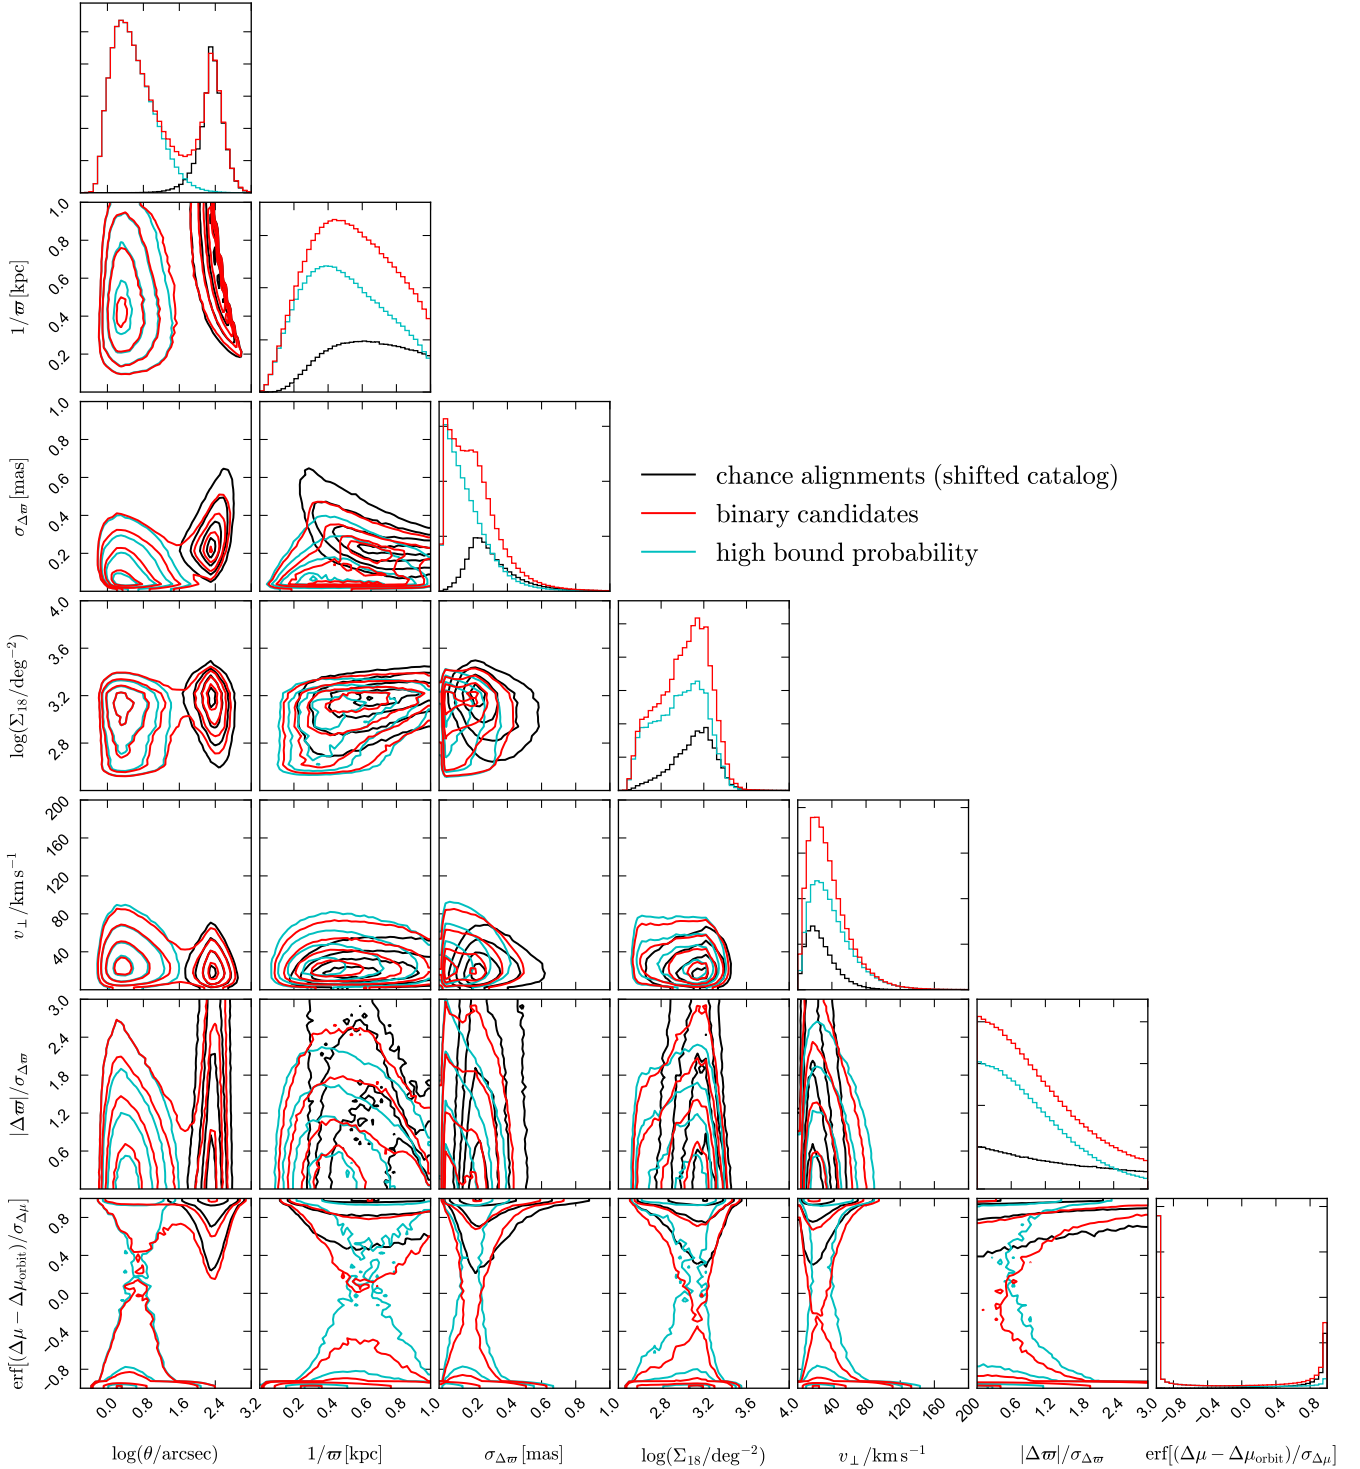


Figure A1. Parameter distribution of chance alignments from the shifted catalog (black), binary candidates (red), and high-confidence binary candidates with $\mathcal{R} < 0.1$ (section 3.2; cyan). Compared to chance alignments, the high-probability binaries have smaller angular separations, closer distances, larger parallax uncertainties, higher local source densities, larger tangential velocities, and more consistent parallaxes and proper motions.




Figure A2. \mathcal{R} is the ratio of the local “density” of chance-alignments from the shifted catalog to that of binary candidates (Equation 8). We compute \mathcal{R} for binary candidates and for chance alignments from a realization of the shifted catalog. We then plot the ratio of the number of pairs in the chance alignment catalog that have a given \mathcal{R} value to pairs in the binary candidate catalog with the same \mathcal{R} value. Dashed line shows a one-to-one relation for comparison.