

Distances and parallax bias in *Gaia* DR2

Ralph Schönrich¹,^{*} Paul McMillan² and Laurent Eyer³

¹Clarendon Laboratories, Rudolf Peierls Centre for Theoretical Physics, University of Oxford, Oxford OX1 3PU, UK

²Lund Observatory, University of Lund, Sölvegatan 27, SE-221 00 Lund, Sweden

³Department of Astronomy, University of Geneva, Chemin des Maillettes 51, CH-1290 Versoix, Switzerland

Accepted 2019 April 23. Received 2019 April 6; in original form 2019 February 6

ABSTRACT

We derive Bayesian distances for all stars in the radial velocity sample of *Gaia* DR2, and use the statistical method of Schönrich, Binney & Asplund to validate the distances and test the *Gaia* parallaxes. In contrast to other methods, which rely on special sources, our method directly tests the distances to all stars in our sample. We find clear evidence for a near-linear trend of distance bias f with distance s , proving a parallax offset δ_p . On average, we find $\delta_p = -0.054$ mas (parallaxes in *Gaia* DR2 need to be increased) when accounting for the parallax uncertainty underestimate in the *Gaia* set (compared to $\delta_p = -0.048$ mas on the raw parallax errors), with negligible formal error and a systematic uncertainty of about 0.006 mas. The value is in concordance with results from asteroseismic measurements, but differs from the much lower bias found on quasar samples. We further use our method to compile a comprehensive set of quality cuts in colour, apparent magnitude, and astrometric parameters. Lastly, we find that for this sample δ_p appears to strongly depend on σ_p (when including the additional 0.043 mas) with a statistical confidence far in excess of 10σ and a proportionality factor close to 1, though the dependence varies somewhat with σ_p . Correcting for the σ_p dependence also resolves otherwise unexplained correlations of the offset with the number of observation periods n_{vis} and ecliptic latitude. Every study using *Gaia* DR2 parallaxes/distances should investigate the sensitivity of its results on the parallax biases described here and – for fainter samples – in the DR2 astrometry paper.

Key words: astrometry – parallaxes – stars: distances, kinematics and dynamics – Galaxy: kinematics and dynamics – solar neighbourhood.

1 INTRODUCTION

This paper presents the first statistical evaluation of parallax biases for a general stellar sample and the first derivation of unbiased stellar distances for the subsample of *Gaia* DR2 with line-of-sight velocity measurements (Gaia Collaboration 2018a; Katz et al. 2019). Other than most other approaches, our method tests the distances of all stars in the selected sample and not just specific subgroups with different physical properties.

Since the first release of *Gaia* DR2 data, there has been clear evidence for offsets in the parallax measurements. In the data release papers, Lindegren et al. (2018) found a general offset of *Gaia* parallaxes by $\delta_p = -0.029$ mas (in the sense that the quoted values were too low) when studying known quasars in the *Gaia* catalogue. In addition, the median offset of the quasars is patchy on the sky (see also fig. 15 in Arenou et al. 2018), and thus we also have to expect an error term σ_p , which will in many aspects behave like a random

error (as long as these patches are not resolved). The calibration involved in the data analysis for the astrometric solution of *Gaia* depends on both the effective wavelength (colour) and magnitude (directly and through the ‘window class’), so both offsets/biases and random errors cannot be expected to be homogeneous along the entire magnitude and colour range of the sample. The quasar catalogue, which is the best direct benchmark for *Gaia* parallaxes, is very faint (typically at apparent magnitudes $G > 17$), and the colour distribution does not match the stellar colour distribution very well. In contrast their internal checks between different parts of the analysis show significant changes in behaviour in particular at magnitudes $G \sim 16, 13, 10$ (caused by their analysis windows; see fig. 16 in Lindegren et al. 2018) and changes in their astrometric measurement accuracy (likely near $G \sim 8$; see their fig. 9). In addition, sources in several remote objects, in particular the Large Magellanic Cloud, display a finely structured pattern in the sky position of mean offsets related to the scanning law of *Gaia*.

In the Gaia Collaboration (2018d) the same offset can be seen when comparing the globular cluster mean parallaxes to the more accurate (at small parallax) values from the Harris (1996) catalogue.

* E-mail: ralph.schoenrich@physics.ox.ac.uk

The Gaia Collaboration (2018d) also shows that the Large and Small Magellanic Clouds, and most of the satellite dwarf spheroidal galaxies it studies, have negative average parallaxes. Consistent with these general offsets, but differing in the magnitude of the effect, Stassun & Torres (2018) found an offset of (-0.082 ± 0.033) mas for a sample of 89 very bright ($G < 12$ mag) binary systems, however with a large scatter of the single measurements. Yet, given the patchiness of the offset, and the very different magnitude range probed, these values could well be in agreement. Just when we were about to submit this paper, the latest study of Graczyk et al. (2019) found a value of $\delta_p = -(0.054 \pm 0.024)$ mas, which by an amusing coincidence is exactly our result, albeit with a much larger uncertainty. We also note that binarity was not worked into the astrometric solution, so it is to be expected that binary systems carry a different bias from the main stellar sample. Similarly, comparison with asteroseismic values points to significant parallax underestimates (Zinn et al. 2018), though again for rather specific subsets of stars in magnitude and colour/stellar evolutionary stage. However, their result of $\delta_p = -0.05$ mas bears high confidence, and again differs from the quasar result. Sahlholdt & Silva Aguirre (2018) looked at a smaller sample of dwarfs with asteroseismic data, and the offset they found was closer to that of the quasar sample, but with more significant uncertainty (-0.035 ± 0.016 mas). No test so far has measured what we are really concerned about: the bias of the full stellar sample and how it depends on the other properties of the observation.

Unbiased distances are key to virtually every problem in modern astrophysics, and given the large sample sizes, we now need these distances at the 1 per cent level. For example the wave pattern discovered by Schönrich & Dehnen (2018), which was confirmed by Huang et al. (2018) and Kawata et al. (2018), and which is likely related to the later findings of Antoja et al. (2018), has a total amplitude of below 1 km s^{-1} . To compare this: Even at a perfect location towards the Galactic anticentre, the solar reflex motion will translate to an $\sim 0.1 \text{ km s}^{-1}$ bias for every 1 per cent in mean distance bias. Larger and more complex bias patterns will arise at other sky positions from cross-correlations between the velocity component measurements.

This work offers a solution to this problem. We will derive unbiased distances to the radial velocity (RV) subsample (about 7 million stars at magnitudes $G \lesssim 15$) using the method proposed in Schönrich & Aumer (2017). By measuring the selection function directly from the sample, we derive a data-informed, and thus nearly unbiased, prior, which suits a sample better than model-based priors (e.g. Astraatmadja & Bailer-Jones 2016; Bailer-Jones et al. 2018) designed to fit pre-existing models of the entire *Gaia* sample (e.g. Bailer-Jones et al. 2018). As discussed in Schönrich & Aumer (2017), mismatch or neglect of the selection function will result in systematic bias of a similar size to the measurement uncertainties for individual stars (i.e. of order 20 per cent for the common parallax quality cut of $p/\sigma_p > 5$). Our previous results have established that this method gives an unbiased translation from parallaxes to distances. Consequently, with these distances we can now directly measure and correct biases in the parallaxes using the statistical method of Schönrich, Binney & Asplund (2012).

Our paper is structured as follows. We start with a short description of the used data and coordinate system definitions in Section 2, followed by a description of our statistical distance estimator in Section 3. After this, we provide the formalism for deriving Bayesian distances in Section 4, including a derivation of the distance-dependent selection function $S(s)$. In Section 5 we quantify the different parallax biases in *Gaia* DR2, which is

followed by a comparison to previous distance derivations and a comment on the distance to the Pleiades. Section 8 provides a summary of quality cuts necessary in *Gaia* DR2, followed by the conclusions.

2 DATA AND DEFINITIONS

2.1 Coordinate frame and definitions

Throughout this paper, we will use the standard definitions for Galactic coordinates and the local standard of rest. We employ Galactic cylindrical coordinates (R, z, ϕ) , where R is the in-plane distance to the Galactic Centre, z is the altitude above or below the Galactic mid-plane, and ϕ is the Galactic azimuth, with the Sun placed at $\phi = 0$. The distance of a star to the Galactic Centre is termed $r = \sqrt{R^2 + z^2}$; for the solar galactocentric distance, we use the value $R_0 = 8.27 \text{ kpc}$ from Schönrich (2012), which is also in agreement with other determinations (Gillessen et al. 2009; McMillan 2017), and only slightly in tension with the latest estimates from measurements of stellar orbits around Sgr A* from the Gravity Collaboration (2018). The vertical displacement of the Sun from the mid-plane, $z_\odot = 0.02 \text{ kpc}$ is taken from Joshi (2007). We also tested that it does not have any significant impact on our results. The velocity vector in the heliocentric¹ Cartesian frame is defined as (U, V, W) with a right-handed set of components pointing radially inwards, in the direction of Galactic rotation, and upwards perpendicular to the plane. The velocity vector components in the Galactocentric cylindrical frame are analogously termed (U_g, V_g, W_g) . To translate these velocity components, we use the motion of the Sun against the local standard of rest as measured in Schönrich, Binney & Dehnen (2010) and Schönrich (2012): $(U_\odot, v_\odot, W_\odot) = (11.1, 250, 7.24) \text{ km s}^{-1}$, and if necessary, use the azimuthal velocity of the Sun against the local standard of rest ($V_\odot = v_\odot - V_c = 12.24 \text{ km s}^{-1}$). For simplicity's sake, we call p the parallax of a star and σ_p the effective uncertainty of the parallax measurement assumed in that instance, which, depending on the examined set of assumptions, may contain the additional $\delta\sigma_p = 0.043 \text{ mas}$ added in quadrature to the *Gaia* pipeline value $\sigma_{p,g}$.

2.2 Data

Here we use the *Gaia* RV sample (Cropper et al. 2018; Gaia Collaboration 2018c) from *Gaia* DR2 (Gaia Collaboration 2018a) with more than 7 million stars that have both astrometric and line-of-sight velocity measurements from the spectrograph (Sartoretti et al. 2018) onboard the *Gaia* spacecraft (Gaia Collaboration 2016). To ensure the quality of the data, we apply, if not stated otherwise for a specific task, a few quality cuts that were discussed in data release papers (e.g. Lindegren et al. 2018), namely the number of visibility periods $n_{\text{vis}} \geq 5$ to ensure a full astrometric solution, a parallax quality cut of $p/\sigma_p > 5$ and a lower limit $p > 0.1 \text{ mas}$, which translates to approximately demanding a distance $s \lesssim 10 \text{ kpc}$, an excess noise smaller than 1, and line-of-sight velocity limits of $|v_{\text{los}}| < 550 \text{ km s}^{-1}$ and $\sigma_{\text{los}} < 10 \text{ km s}^{-1}$. We usually remove the Galactic mid-plane from our sample, i.e. require $|b| > 10 \text{ deg}$.

¹We use the somewhat negligent Galactic dynamicists' term 'heliocentric', while in truth, *Gaia* is measuring quantities in the Solar system barycentric frame. With relative motions between the two frames below 0.1 km s^{-1} , this difference does not matter.

For our statistics the Galactic mid-plane carries no signal and we thus avoid problems with crowding and excessive reddening. We checked, though, that the measurement of the distance prior from low- $|b|$ data is similar to our higher latitude main sample, and provide distance estimates for these stars in the derived catalogue. In previous papers (see Schönrich & Aumer 2017) we uncovered major problems with v_{los} measurements, in particular with LAMOST. Here, we just note that our tests of v_{los} accuracy and precision looked very decent on the *Gaia* sample, and we will concentrate on the more pressing issue of distances and *Gaia* parallaxes. We follow the convention of *Gaia* papers to call their apparent magnitudes in the three broad colour bands (G_{BP} , G , G_{RP}). Problems with capitalization conventions do not arise, since we do not discuss absolute magnitudes through most of the paper.

3 STATISTICAL DISTANCE ESTIMATION

3.1 General thought

For the determination of distance bias, we rely on the method of Schönrich et al. (2012; hereafter **SBA**), which has been applied to various samples. The method relies on correlations between velocities, which depend on the position on the sky. The estimator is readily derived by writing down an estimate of stellar kinematics while allowing for a systematic distance bias $f = \langle s'/s \rangle$, where s' denotes the estimated distance and s the real distance to a star. This f affects both tangential velocity components simultaneously, correlating them. To explain this with a simple example, imagine approaching a mountain horizontally. Knowing your velocity, your mind automatically has a clear estimate of the distance to the summit. This is because any incorrect estimate would translate your horizontal motion to a vertical component of motion of the mountain; i.e., the top of the mountain would have to be growing or shrinking (if you had over- or underestimated the distance, respectively). Analogously, all parts of the mountain base below your level would appear to be moving downwards (upwards). And your brain knows this is not usually what mountains do. The **SBA** method allows us to extend this intuition and make it more robust against assumptions (i.e. we do not assume any mean motion or fixed velocity ellipsoid) to stellar samples.

The strength of the **SBA** method now lies in directly using the spatial dependence of this correlation of the heliocentric velocity components on the galactic longitude (l) and latitude (b). As long as we have a sufficient sky coverage, we do not depend on classic assumptions of other methods: Think of our mountain example. Observing many mountains around us, we gain a significant advantage in control of systematics over the use of just one single mountain. In a simplified picture, we just measure the pattern of apparent vertical velocities of all the mountain tops and mountain bases around us and try to find the distance correction that makes the angular dependence of this pattern disappear. Misjudging our own velocity, we would equally bias the motion of all summits and mountain roots, leaving the statistical distance evaluation unscathed. Similarly, we do not care about our horizontal velocity, since it can (i) be measured and (ii) would just affect our prediction for the strength of the effect around us; however, we just seek the distance factor at which the correlation of vertical velocity components with sky position disappears, making our own horizontal motion irrelevant. Translated to our real problem: Assumptions about the solar velocity do not matter for our method.

Similarly, let us assume that we are sitting in the middle of an orogeny (or reading Calvino's *Cosmicomics*) and both mountain

summits and roots are rapidly rising and sinking into the ground. Observing just one mountain in front of us, we would indeed infer that distances are overestimated, but behind us, the correlation term reverses sign; i.e., the apparent distance underestimate there cancels out the distance overestimate inferred from the opposite direction. We learn from this that typically modes of the disc cancel out in a sample with large sky area. Analogously, a wide halo stream passing through would cancel by the spatial terms. In short, galaxy physics can only affect our statistical measurements if they vary across the sky in a way that correlates with the angle terms of our method. In most cases (e.g. global breathing modes, streams) they will cancel out at first order.

Lastly, we note that other than our imaginary mountains, stars move horizontally, so in addition to the spatial correlations, we can benefit from two different horizontal velocity components with different dependences on sky position.

3.2 Formal argument and specific implementation

The formal derivation of our method (see **SBA** for a stringent treatment) is done by simply writing down what happens in the measurement. The vector of observed values is $(s' \mu_l, s' \mu_b, v_{\text{los}})$, where s' is the observationally inferred distance and μ_l and μ_b are the proper motions in Galactic longitude l and latitude b . This vector is translated into the measured velocity components (U, V, W) by a matrix \mathbf{M} depending on l and b . Since this is an orthogonal matrix, the inverse mapping (i.e. from the original velocity components) is done with the transverse \mathbf{M}' . If we now assume that distances are changed by some relative bias

$$f = (s' - s)/s, \quad (1)$$

where s is the real distance, we can relate:

$$\begin{pmatrix} U \\ V \\ W \end{pmatrix} = \mathbf{M}(\mathbf{I} + f\mathbf{P})\mathbf{M}' \begin{pmatrix} U_0 \\ V_0 \\ W_0 \end{pmatrix} \quad (2)$$

where the index 0 indicates the real values, and \mathbf{P} is $\text{diag}(1, 1, 0)$, which projects to the two proper motion components. Now, we see that the observed velocity components are correlated by f via the matrix $\mathbf{T} = \mathbf{M}\mathbf{P}\mathbf{M}'$. Since the equations are linear, the average f can thus be gained by a similar linear regression of any target velocity component v_i on to the other velocity components v_j multiplied with T_{ij} , the components of the matrix \mathbf{T} .

As discussed in **SBA**, using the in-plane velocity components (U, V) mixes the statistics with a Galactic rotation estimate, and given the very large sample size here, we make the choice to avoid this possible source of systematic bias. We thus limit this study to using the correlation of both U and V velocity components with the vertical motion W . The relevant part of \mathbf{T} is thus

$$\begin{pmatrix} T_{uw} \\ T_{vw} \\ T_{ww} \end{pmatrix} = \begin{pmatrix} \cos l \sin b \cos b \\ \sin l \sin b \cos b \\ 1 - \cos b^2 \end{pmatrix}. \quad (3)$$

Our method applies corrections for the following biases of this measurement:

- (i) v_{los} determination errors, σ_{los} , which would appear as distance underestimates (typically negligible due to the excellent precision and accuracy of the *Gaia* v_{los} estimates),
- (ii) proper motion determination errors, σ_{\pm} , which would appear as distance overestimates, but are again negligible by more than an order of magnitude,

(iii) the tilt of the velocity ellipsoid, which is of some importance for the statistics. This term is important, as the radially elongated velocity ellipsoid produces a locally changing correlation between the heliocentric velocity components, which can partially line up with the T_{uw} and T_{vw} angle combinations.

We have to add the systematic uncertainties from these terms to our error budget, assuming that the uncertainties in these terms are statistically independent of each other. As already done in previous studies (SBA) we assume a systematic uncertainty of 10 per cent of the calculated correction value for the first two terms, and an uncertainty of 30 per cent for the turn of the velocity ellipsoid correction.

As stated above, due to the unprecedented precision of *Gaia*, the exact values of proper motion errors do not matter here as long as the order of magnitude of the uncertainty estimates in the *Gaia* pipelines is correct. Similarly, the error correlations are mostly inconsequential: Two team members did independent tests on independently calculated mock samples, where we folded the mock measurements with the full error matrix between parallaxes and proper motions as given for each star in the sample. In these tests, the effect of error correlations on our statistics is more than one order of magnitude less than our systematic and formal error budget for whole-sky measurements. On pencil beams, like in tests of high β stars, it contributes of order one-tenth to the residual bias (see analysis below).

For the velocity ellipsoid correction, we assume that the velocity ellipsoid points to the Galactic Centre at every position. Other than in previous applications of the SBA method, the *Gaia* sample spans a large volume throughout the disc, and when we select by distance, stars in each sample will cover regions with vastly different values of the velocity dispersion. To optimize the estimate for this correction term, we directly measure the velocity dispersions in the Galactocentric spherical coordinate frame weighted by their impact on the distance estimator.

As described in Section 3.1 the method does not assume any velocity ellipsoid, and does not even require knowledge of the correlations between the velocity components (e.g. U and W). The only important requirement is that Galactic structure does not infer a correlation between this velocity correlation and Galactic position. Realistic structure (e.g. a stream passing through the survey, or disc breathing modes) might produce a local velocity correlation (and so distance statistics on small patches of the sky are uncertain; see the β issue below), but cancels out to first order with large sky coverage. We have already tested and confirmed this on realistic simulations in the appendix of Schönrich & Aumer (2017).

4 DISTANCE DETERMINATIONS AND SELECTION FUNCTION FOR GAIA DR2

Before we can test distances/parallaxes with the SBA method, we have to derive distance expectation values for all stars, using the method of Schönrich & Aumer (2017). From parallax measurements we calculate the probability distribution in distance $P(s)$ for each star by

$$P(s) = N^{-1} s^2 G(p, p_0, \sigma_p) \rho(s(p), l, b) S(s(p)), \quad (4)$$

where

$$N = \int ds s^2 G(p, p_0, \sigma_p) \rho(s(p), l, b) S(s(p)) \quad (5)$$

is the normalization, s is the distance from the Sun, p denotes a parallax, $G(p, p_0, \sigma_p)$ is the (Gaussian) observational likelihood

distribution in parallax, given the measurement p_0 and effective uncertainty σ_p , and $\rho(s(p), l, b)$ is the assumed density model. $S(s)$ denotes the selection function, i.e. number of stars detected in the sample divided by the number of stars actually there. As in previous papers, we use the simple density model from Schönrich & Bergemann (2014), which contains a thin-disc, thick-disc, and halo component (we can neglect the bulge because we only select stars with $|b| > 10$ deg). All calculations are done with a self-adaptive trapezoid integration, where we start from the suspected maximum of the probability distribution function (PDF) and integrate to both sides in distance simultaneously, adapting the step length upwards when the relative contribution of each segment to the integrated value falls below a threshold. We tested that our results are the same when lowering the initial step length and threshold by a factor 10, or when using a simple step-wise integrator.

We tested by cutting the sample in Galactic latitude and longitude that this model is sufficiently close to the underlying distribution to ensure a good distance measurement. The most important quantity in the above equation, however, is the distance-dependent selection function $S(s(p))$, which arises from the magnitude-dependent cuts of the sample. Predicting $S(s)$ directly would require a full chemodynamical galaxy model including a three-dimensional reddening map. With this choice, we would be vulnerable to systematic uncertainties of the model choice, stellar evolution models, and the reddening map.

Here, we instead measure $S(s)$ from the data themselves. Never the less, it is useful to formulate a model-based expectation, which acts as a sanity check to our results. The top panel of Fig. 1 shows the selection function $S'(s)$ calculated from simple population synthesis using the machinery of Schönrich & McMillan (2017) using B.A.S.T.I. stellar models (Pietrinferni et al. 2004, 2009) and a simple Salpeter (1955) initial mass function (IMF). The shown $S'(s)$ expresses the number of stars per solar mass (at birth) of a stellar population at a given distance s . Since the normalization factor here is irrelevant for our purposes, we multiplied $S'(s)$ with an arbitrary factor 80 to facilitate direct comparison with the measured $S(s)$ in the bottom panel. For this figure, we suppose a simple magnitude-dependent selection, setting the selection probability constant below Johnson V -band $m_V < 12.8$ mag and then going linearly to zero for $m_V \geq 13.5$ mag.

$S(s)$ has three main regions: (i) a quite steep roughly exponential decrease of $S'(s)$ with distance s in the near field, stemming from the magnitude limit moving up through the main sequence and turn-off region with increasing distance modulus. The exponential behaviour is expected, since the magnitude scale is logarithmic, the luminosity of stars is roughly the fourth power of the mass, and the IMF is close to a simple power law (as we can neglect the impact of stars below $\sim 0.5 M_\odot$ where this is no longer true). This decrease is followed by (ii) a weakly inclined plateau associated with the subgiant branches and red clump, followed by a downward knee and somewhat shallower decline at the largest distances due to (iii) the red/asymptotic giant branch stars. There are some differences between different ages and metallicities: Metal-poor stars are somewhat more luminous on average, shifting the features to the right; older populations have their subgiant branch at fainter magnitudes and less stars overall on the giant branch (the lengthening of stellar main-sequence lifetimes outweighs the increase in the IMF). However, the common features let us predict the functional shape for the selection function rather well, when averaged over all stellar populations. We choose

$$S(s) = aA(s)B(s)C(s) \quad (6)$$




Figure 1. Top panel: selection function $S(s)_{\text{model}}$ in distance from a simple population synthesis for populations with fixed metallicity and age and a simplified magnitude-dependent selection function resembling the *Gaia* DR2 subset with v_{los} measurements. The normalization is irrelevant, so we multiplied with an arbitrary factor to match the bottom panel. Bottom panel: selection function measured from the data after several iterations with different quality cuts on the relative parallax quality. The least tight cut ($p/\sigma_p > 2.5$) should not be used in stellar sample selections as it will comprise a large number of catastrophic distance errors. It can, however, serve to educate us on the real shape of the distance-dependent selection function $S(s)$. Note that the normalization is irrelevant and here comprises amongst other effects a geometric factor 2π .

where a is a normalization constant, and the three multipliers are:

$$A(s) = \exp(-bs) + c \exp(-ds) \frac{1.0}{1.0 + \exp(-h(s-j))}$$

$$B(s) = \frac{1.0}{0.5\pi + k_2} (\tan^{-1}(l(l_2 - s)) + k_2)$$

$$C(s) = 1 - \exp(-zs),$$

The rationale behind this is to capture in $A(s)$ the general shape with two exponentials of scale length b^{-1} and d^{-1} for the short-range (out to ~ 1 kpc) and long-range behaviour (beyond 3 kpc) and model the step in $S(s)$ with \tan^{-1} in $B(s)$, as well as the (unimportant) drop-out of luminous or otherwise too close stars (potential loss due to proper motions) with $C(s)$.

We proceed as in Schönrich & Aumer (2017), iterating the distance calculation with the adapted prior. However, due to the good nominal quality of *Gaia* parallaxes, the general appearance of

Table 1. Parameters of the selection function (equation 6).

Parameter	Value	Unit
a	97.9594	–
b	4.10228	kpc $^{-1}$
c	0.0159642	–
d	0.15	kpc $^{-1}$
l	2.22279	kpc $^{-1}$
l_2	2.97547	kpc
k_2	2.0704	–
j	0.956364	kpc
h	4.91193	–
z	0.024	kpc $^{-1}$

the measured $S(s)$ is already present from the first iteration using a flat prior. A fit of the function to data corrected for a mean parallax offset of $\delta_p = -0.048$ mas (see following below) after several iterations of the prior fitting is presented in Fig. 1 (the values are provided in Table 1). The shape matches the prediction from the population synthesis. We also checked that the selection function does not vary strongly with galactic latitude, which signals that the spatial prior is sufficient and population differences throughout the galaxy will not strongly bias our distance estimates.

Apart from d , which sets the scale length of the long exponential component, all parameters are well constrained in the fit. The latter suffers from the fact that we cannot fit beyond the point where the parallax quality cut affects the sample. This cut must not enter $S(s)$, since the stars are culled after they have passed the magnitude limits. Inspection of Fig. 1 shows that given the relative drop-out rates of stars at different quality cuts, the fit must be close to the real shape. Extensive further tests of the far-distance end favoured an additional flattening of $S(s)$ by multiplying with $\exp(-(s'/\text{kpc} - 4)/0.07)$, where $s' = \min(s, 10 \text{ kpc})$. The difference can be inspected between the top panel of Fig. 2, which does not contain this factor, and the bottom panel, which does include it. It mostly serves to improve a slight kink in the distance statistics. The latter factor only has a minor effect on very remote stars beyond $s > 3$ kpc, and we will apply it to all data shown from Fig. 3 onwards. Both the model expectations and the derived distance statistics for very remote stars, which react strongly to the prior, support this choice.

5 DISTANCE BIAS VERSUS DISTANCE: EVIDENCE FOR THE PARALLAX OFFSET

5.1 Scanning versus distance: detection of the parallax offset

Fig. 2 measures the relative distance bias $1 + f$ when using the SBA method for the stellar sample when ordered and binned in the measured distance s . The top panel shows this scan of $1 + f$ versus s for two different quality cuts in the relative parallax error, allowing a maximum σ_p/p as given by *Gaia* DR2 of 10 per cent or 20 per cent, respectively. It is apparent that $1 + f$ increases nearly linearly with s , reaching a relative distance bias in excess of 5 per cent at $s \sim 1$ kpc and errors in excess of 25 per cent near $s \sim 5$ kpc. Such a bias precludes a precision measurement of galactic kinematics even in the relative near field. The comparison of different σ_p/p cuts proves that this cannot be an issue with the distance priors or other parts of our measurement method. Otherwise, a tighter quality cut would drastically reduce the bias $1 + f$ at a given distance. This leaves only one culprit: a bias in the *Gaia* parallaxes.

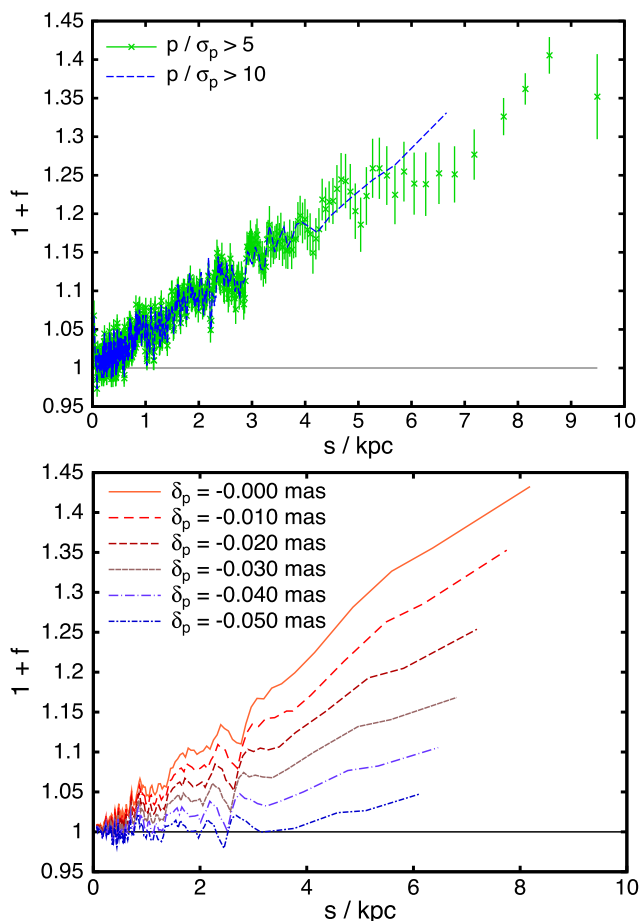


Figure 2. Relative distance error $1 + f$ in the *Gaia* DR2 RV sample when binning the sample in distance. In each panel, we sort the sample in distance s and then let a mask of 15 000 stars width slide by steps of 5000; i.e., every third data point is independent. The $1 + f$ shown on the y-axis, measured with the SBA method, denotes the factor by which the average distance in each subsample is wrong. E.g. $1 + f = 1.1$ means that stellar distances are on average 10 per cent too large. The top panel shows that the result does not depend on the value of the quality cut in the relative parallax error p/σ_p . The bottom panel (using a sample size of 90 000 in steps of 30 000) uses $p/\sigma_p > 5$ and varies a fixed parallax offset, i.e. incrementing all measured parallaxes by δ_p , which reduces the estimated distances.

In fact, as we can see from the bottom panel in Fig. 2, this linear trend is a signature imprint of such a parallax bias. When we apply the correction δ_p to all parallax measurements, we can minimize the trend for $\delta_p = -0.048$ mas with an uncertainty of about 0.006 mas. The most distant bins beyond $s > 3$ kpc are in line with this estimate within the uncertainties. Small differences at these distances should be ascribed to uncertainties in $S(s)$ as discussed above. The offset is about double the amount found by Lindegren et al. (2018), but exactly in line with the asteroseismic evaluation of *Gaia* DR2 (Zinn et al. 2018). We note again that we do not believe the evaluations in Arenou et al. (2018) and Lindegren et al. (2018) on quasars to be applicable to our case, since those are in different apparent (magnitude) window classes with separate astrometric calibrations, and have all different kinematics/zero intrinsic parallax.

We further note that there is a significant positive bias $f \sim 5$ per cent for the distances of the nearest stars or, equivalently, brightest parts of the sample (seen in the left-most green data point in




Figure 3. A quantification of the parallax bias δ_p in the sample. Here, we evaluate $f(s)$ for different values as in Fig. 2. These values we fit in the interval $0.1 < s/\text{kpc} < 3$ both with a constant value and with a linear regression $f(x) = a + \frac{df}{ds}s$ and show the results with the red and blue lines (including their formal 1σ error intervals depicted with short-dashed lines). For comparison, we show the same evaluation, but now increasing the parallax error $\sigma_p = \sqrt{\sigma_{p,g}^2 + (\delta\sigma_p)^2}$ by adding $\delta\sigma_p = 0.043$ mas in quadrature (long-dashed without error bars).

the top panel Fig. 2). This bias is near impossible to explain with bad v_{los} measurements, which would feign a negative f . This bias is for these nearby stars orders of magnitudes larger than the previously discussed parallax offset, which for these stars is negligible. Some of this may be traced back to the magnitude-dependent deviations (see below, Fig. 9), some to misidentifications in the near field.

So far, we simply applied a parallax offset. However, from Lindegren et al. (2018) a position-dependent variation in the parallax bias is expected. There has been discussion in the *Gaia* collaboration whether such a bias should be added to the error budget or not. Now, in our case, we are interested in the uncertainty for each single star. A priori, we do expect a random but spatially correlated fluctuation of the parallax offset to enter single-star parallax uncertainties just like an additional term that has to be added in quadrature to the formal uncertainty given by the pipeline, setting $\sigma_p = \sqrt{\sigma_{p,g}^2 + (\delta\sigma_p)^2}$.

We thus quantify the correction δ_p on two versions of the sample, once with and once without systematic $\delta\sigma_p$, which we take to be 0.043 mas following the quasar analysis of Lindegren et al. (2018, table 4). One can argue for two ways of measuring the parallax bias: On one hand, we could demand that the average f in the safe region $0.1 < s/\text{kpc} < 3$ should be exactly 0. While we have high confidence in the accuracy of the distance statistics in this area, one might still want to get rid of the need for a correct zero-point. Indeed, since a parallax offset gives an almost linear $f(s)$ dependence, we can alternatively demand that the estimated slope df/ds should be zero. However, within our systematic uncertainties, both methods, shown in Fig. 3, yield the same values for δ_p . Drawing a mean estimate from the shown results and further tests of varying quality cuts, we conclude that $\delta_p = -0.048$ mas with a negligible formal uncertainty and a systematic uncertainty of ~ 0.006 mas. The systematic uncertainty is a cautious estimate. The systematic uncertainties on the velocity ellipsoid and measurement uncertainties are already priced into the formal errors, but we priced in a correlated systematic error between evaluation bins, and performed a variation of fit parameters (region in s to fit on), changes in quality cuts (see below), and comparison of the different




Figure 4. Relative distance error in the *Gaia* DR2 RV sample when binning the sample in parallax quality p/σ_p . We use the same binning scheme as in the previous figures. In addition we remove all stars with $s, p^{-1} > 10$ kpc from the sample. Again the y-axis shows the distance bias factor measured with the SBA method. Both panels reveal a strong increase. The top panel shows the two possibilities of adding an additional parallax measurement uncertainty of 0.04 mas in quadrature versus not adding it; the bottom panel displays the same statistics for different cuts in the maximum value of p^{-1} .

methods. These would advise a mildly smaller number, and we added a budget for effects that we may have missed. Separate from these effects, accounting for $\delta\sigma_p = 0.043$ mas further increases the $|\delta_p|$ estimate by about 0.006 mas to $\delta_p = -0.054$ mas. The small difference between the f and the df/ds estimation for δ_p can be ascribed to bad luck as it is within $\sim 2\sigma$. We think, however, that it has contributions by one or several of the secondary problems identified below.

5.2 Quantifying parallax offset versus uncertainty

Could we make different assumptions for the parallax error visible with our distance method? In principle, yes – the background is that the expectation value of a single star’s distance shifts systematically with the assumed error. Consequently, a misjudgement of uncertainties will typically show as a distance bias. However, as we saw in Fig. 3, this effect is subordinate to the other problems in the sample. Never the less, assessing the right value for the p/σ_p quality cut requires a sample scan in this quantity. This is done in the top panel of Fig. 4. We remind the reader that apart from stars at very small distances (corresponding to large p/σ_p) in this plot, the distance scan (see Fig. 2) shows no significant deviations after

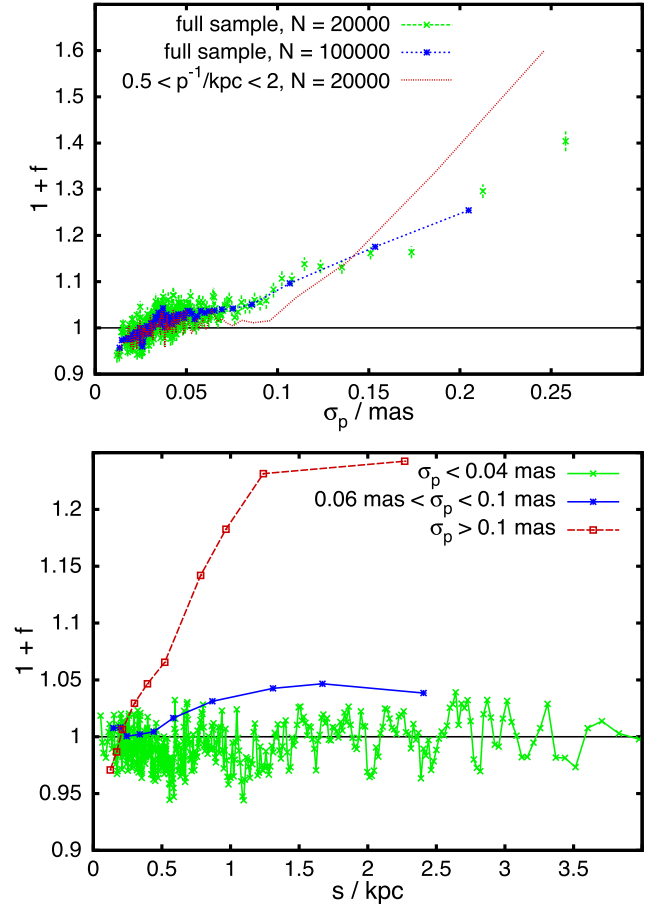


Figure 5. Top panel: distance bias versus the the parallax error σ_p as provided in the *Gaia* DR2 data set. The bottom panel displays the distance bias versus distance to uncover the nature of the trends/biases observed in the top panel. The subsample with large parallax uncertainty, $\sigma_p > 0.1$, shows the typical signature of a larger parallax offset; i.e., the distance bias rises almost linearly with distance. The other subsamples have a weak indication of the same trend.

correcting for δ_p . Yet, the red line in the top panel of Fig. 4 reveals an impressive increase in $1+f$ towards lower p/σ_p . Note, however, that we have extended the sample to stars beyond our usual quality limit, admitting all stars with $p/\sigma_p > 1$; below $p/\sigma_p < 4$ there will definitely be a sizeable fraction of stars with catastrophic distance misestimates. Also, the large parallax uncertainty allows for major parts of the probability distribution $P(s)$ to cover highly uncertain regions of $S(s)$. However, extensive experimentation with $S(s)$ could not rectify the abnormalities in Fig. 4. One could now be tempted to argue that the suggested inclusion of $\delta\sigma_p = 0.043$ mas removes the problem, but this is a deception caused by the left shift of the graph, since we increased σ_p . The $S(s)$ uncertainty can be resolved by limiting the p^{-1} range to contain the PDFs of even uncertain parallax measurements within the safe region $s < 3.5$ kpc. The result of this is shown in the bottom panel of Fig. 4: There always remains a spike in $1+f$ on the left-hand side. Something else is going on here.

We now restore the parallax quality limit and plot $1+f$ against the pipeline parallax error σ_p itself in the top panel of Fig. 5. For $\sigma_p > 0.12$ mas distance bias sky-rockets, even for the rather conservative cut $p^{-1} < 1.5$ kpc applied here. This suddenly clarifies the tension in the plots of $1+f$ against p/σ_p , because we were just moving




Figure 6. Studying the hypothesis that the parallax error predicts the parallax offset. The top panel shows a scan of distance bias versus parallax error when assuming that the parallax offset equals the parallax error, i.e. $\delta_p = -\sigma_p$, and attempting to correct the error by adding σ_p to the parallax before evaluating the distance. The two versions of doing this are shown with green errorbars (adding $\delta\sigma_p = 0.043$ mas in quadrature to σ_p) and the blue line (not adding a parallax error). It is striking how well the data set is constructed when assuming this, instead of the constant parallax offset discussed previously (and shown for comparison with the purple line). In the bottom panel, we show a scan in distance of the sample, again showing that the method of first adding $\delta\sigma_p = 0.043$ mas in quadrature produces a near-perfect outcome.

the region of large parallax errors via the distance/parallax cut. A further unsettling observation is the uptrend of the distance bias with parallax error between $\sigma_p = 0.045$ mas and $\sigma_p = 0.06$ mas, turning from a negative to a positive distance bias of order 2 per cent.

The bottom panel of Fig. 5 attempts to qualify the nature of this failure. One could argue that stars with large parallax uncertainties noted by the pipeline should be binaries, affecting their proper motions, and thus our distance statistics. However, we see a clear increase of the distance bias with distance, putting the blame at a parallax offset exceeding by far the ~ 0.05 mas of the entire sample, for which we already correct in the shown data. Given the displayed results, the only viable explanation is that δ_p is at least to some extent proportional to σ_p .

With such an unusual finding, it is natural to point the search for an honest error at our own code. For example, there could be a typo in our distance integral creating the dependence on σ_p . Apart from controlling and testing our code, Fig. 6 investigates this by checking different assumptions for the distance error. The purple line displays




Figure 7. Top panel: distance bias versus additional parallax error $\delta\sigma_p$. Like in Fig. 3 we bin the sample in distance and measure both the average distance bias (red lines) and the trend of the distance bias with distance for distances smaller than 3 kpc. To ensure a clean sample, we removed all stars with $\beta < -55$ deg, which, however, has only quite a minor impact on our statistics. Bottom panel: the same statistics as the top panel, now plotted against the parameter q when we assume that $\delta_p = -q\sigma_p$.

the original trend as found with a constant $\delta_p = -0.048$ mas. The blue line corrects this to setting $\delta_p = \sigma_p$ and the green line in addition adds $\delta\sigma_p = 0.043$ mas in quadrature. If we had made a σ_p -dependent error ourselves, the green line and the blue line should deviate in the same way (as the distance estimates depend very weakly on the assumed σ_p). In contrast, we see that only with the full correction and assuming that $\delta_p = -\sigma_p$ can we rectify the trend in the sample.

While we do not want to adhere to the notion that the parallax offset is perfectly equal to the parallax error, we do in Fig. 7 attempt to show a quantification of this dependence. In both panels we use our usual quality cuts in colour, g -band magnitude, and $p/\sigma_p > 4$. In addition, we removed the parts of the sample with $\beta < -55$ deg, which, however, has a very minor effect. The top panel demonstrates that if we assume that $\delta_p = -\sigma_p$, both the average distance error and the trend of distance error with distance are within the systematic uncertainties, in line with assuming the usual additional parallax error $\delta\sigma_p \sim 0.043$ mas added in quadrature. The bottom panel tests different values of the proportionality constant q , when we set $\delta_p = -q\sigma_p$ after adding the additional term to σ_p . Both statistics are in line with a value very close to 1.




Figure 8. Distance bias $1 + f$ versus parallax error σ_p . Using the usual quality cuts, we probe different assumptions for changing δp and σ_p . The plot uses samples of 90 000 stars, moving the mask in steps of 30 000.

While the notion of $\delta p = -\sigma_p$ is simple at first hand, it is rather unreasonable to believe that there should be such a perfect equality. To this end we tested a third option: While leaving q free again, we added the offset found for the quasars to the parallax, i.e. set $-\delta p = 0.029 \text{ mas} + q\sigma_p$. Since about half the offset in parallax needed is now again captured by a constant term, we can expect that a good solution will be found for $q \sim 0.5$ mas. This result is shown in Fig. 8, where we plot the distance bias against σ_p for both $q = 1$ and $q = 0.5$ plus the constant term. However, various experiments show that we cannot get rid of the trend of f at small σ_p if q is not close to 1 in this region. The equality is definitely not perfect, since we would require a q closer to 0.5 in all possible assumptions near $\sigma_p \sim 0.09$ mas and a larger $q > 1$ for $\sigma_p > 0.012$ mas.

To summarize this: The best choice for the sample is to add $\delta\sigma_p \sim 0.043$ mas in quadrature, and to add ~ 1 times the parallax error σ_p to the parallax. When concerned about precision, we further advise to remove all stars with $\sigma_p > 0.08$ mas, and recommend to separately test stars with $\sigma_p \lesssim 0.047$ mas for anomalies.

5.3 Bias versus colour and magnitude

After these more complicated considerations, the derivation of safe limits in colour and magnitude for the sample is quite straightforward. We simply order our sample by G - or G_{RP} -band magnitude (for an overview of the photometric instrumentation and calibration in *Gaia*, see Evans et al. 2018; Riello et al. 2018), as well as by the given $G_{\text{BP}} - G_{\text{RP}}$ colour, and then iteratively tighten the quality cuts. The results after a first round of clean-ups is shown in Fig. 9. The top panel shows the distance bias versus G magnitude after removal of all stars with $G > 14.5$ mag and $G_{\text{BP}} - G_{\text{RP}} > 1.5$ mag, while the bottom panel shows a scan in $G_{\text{BP}} - G_{\text{RP}}$ colour after removal of all stars with $G > 14.5$ mag. If we had not censored the faintest stars in the top panel, they would be shown with $1 + f \sim 0.8$, i.e. a very strong distance underestimate. Since we recall line-of-sight velocity measurement errors (even if unbiased) look like distance underestimates, by far the most likely explanation for the decline in $1 + f$ is not a failure of *Gaia* parallaxes, but a much larger than indicated error from the v_{los} measurements. It would remain to be explored if the mild decline around $G > 13$ mag may also be related to a change of magnitude window in the *Gaia* astrometry. Of course, the change of colour range implies a change in the

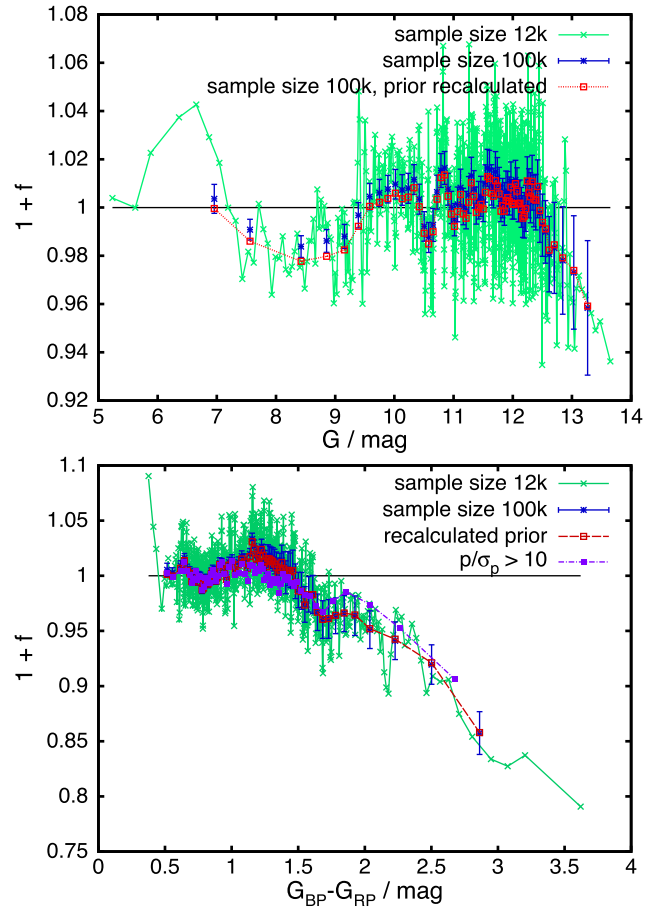


Figure 9. Top panel: distance bias versus G magnitude of the stars. Here we assume the $G_{\text{BP}} - G_{\text{RP}} < 1.5$ mag quality cut. The figure displays different sampling sizes (100 000 moving in steps of 50 000 versus 12 000 moving in steps of 4000), and gives one comparison where we recalculated the prior. Bottom panel: scanning the sample in a similar way in $G_{\text{BP}} - G_{\text{RP}}$ colour. Here we assumed a $G < 14.5$ mag quality cut and used the assumption ($\delta p = -0.048$ mas, $\delta\sigma_p = 0$).

selection function. We thus devised an automated measurement of the selection function $S(s)$, where we use about 40 base points for $s < 4$ kpc, on which we remeasure $S(s)$ and then calculate a grid of relative factors between which we interpolate linearly on $\log s / \text{kpc}$. As the reader can easily see from the difference between the red and blue points in the top panel of Fig. 9, this more appropriate but far more costly calculation does not significantly change the results. The largest change is around $G \sim 8$ mag, where the recalibration shortens all distances a little, since $S(s)$ drops more steeply with s , and thus exacerbates a little the negative bias in this area. We also note that the relatively sharp and borderline significant feature just below $G \sim 11$ mag resembles a lot what was shown in the recalibration comparisons in Lindegren et al. (2018). However, the sample size and small amplitude of the effect preclude further investigation.

5.4 Bias versus astrometric parameters

After discussing the main problems with the sample, it is worth looking at the sample quality versus the different astrometric parameters. Fig. 10 shows the distance bias against the number of visibility periods n_{vis} (top panel) and the ecliptic latitude β (bottom




Figure 10. Examining the dependence of the residual distance bias on measurement geometry. In the top panel, we order the sample by the number of visibility periods, n_{vis} , using a sliding mask of width 200 000 moved in steps of 100 000, finding a significant downtrend when using a constant parallax offset (red line). This trend disappears (blue line) when we assume $-\delta_p = \sigma_p$ and limit $\sigma_p < 0.12$ mas. The last data point contains only 20 000 stars and is likely an outlier. Similarly, the bottom panel shows the systematic distance bias against ecliptic latitude, β , using a sliding mask of width 100 000, moved in steps of 50 000. The general improvement is as in the top panel. The ecliptic poles imply pencil beams and are thus not reliably measured.

panel). In both the top and bottom panels we compare the results obtained with the simpler assumption of a constant parallax offset $\delta_p = -0.048$ mas (red lines) versus setting $\delta_p = -\sigma_p$ and limiting $\sigma_p < 0.12$ mas (blue lines). The σ_p limit does not significantly alter the results. We were first rather astounded by the strong dependence of the distance bias on n_{vis} in the naive evaluation (red lines). However, most of this effect is readily explained by the dependence of δ_p on σ_p . First, we note that due to the scanning law of *Gaia*, n_{vis} very strongly correlates with β , so both trends have a common explanation. The culprit is easily found when we correct for the apparent dependence of δ_p on σ_p (blue lines), which diminishes the trend. Consistently with the top panel, most of the dependence of $1+f$ on β in the bottom panel disappears when we apply this correction, with a mild failure of order $f \sim 5$ per cent remaining around the ecliptic south pole. This is far less concerning than it looks since the ecliptic poles imply a pencil beam, where our method loses most of its advantages, and furthermore, the ecliptic poles are in a near-worst-case location, since they are close to the Galactic plane almost exactly in the azimuthal direction: This implies weak statistics, and only the $V-W$ correlation term, which carries some

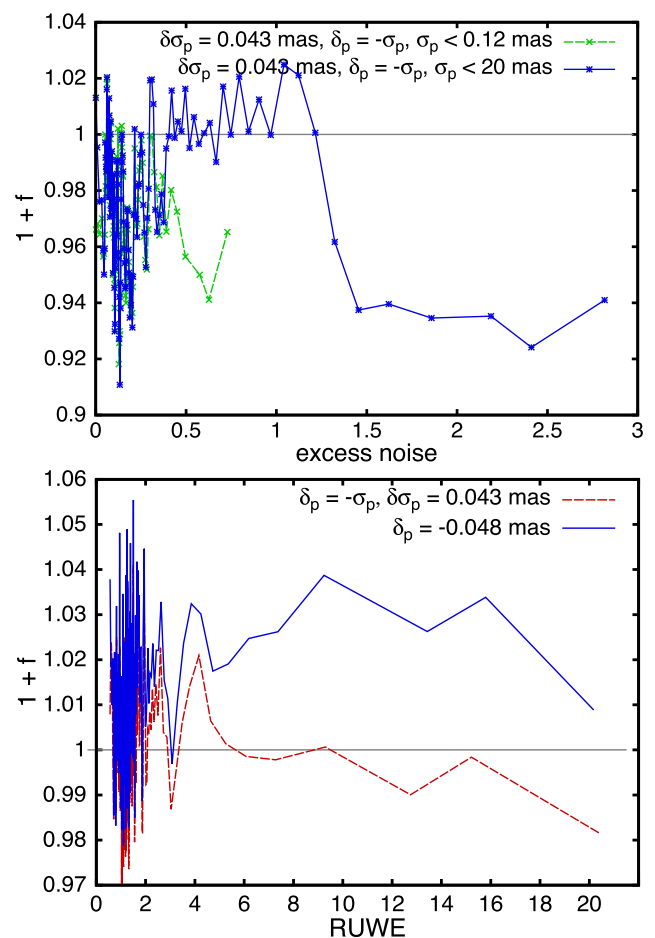


Figure 11. Top panel: distance bias versus excess noise. Few stars have a positive excess noise, so we scan in subsamples of 9000 stars, moving the mask on the ordered sample by 3000. Excess noise correlates strongly with the *Gaia* pipeline’s parallax error, so in this plot we use the assumption $\delta_p = \sigma_p$ with $\delta\sigma_p = 0.043$ mas. The blue line relaxes the cut on parallax quality, which in turn allows us to show the tail of large-excess-noise values. Bottom panel: distance bias versus RUWE. No significant trend can be detected when we use the $\delta_p = \sigma_p$ assumption; a mild trend exists when $\delta_p = -0.048$ mas, since the RUWE is correlated strongly with σ_p and thus high values of RUWE imply a large σ_p .

mild signal from the galactic warp, and no information from the $U-W$ term. We thus ran a few checks. The distance method did not reveal anything else than near-perfect distance bias when scanning along Galactic l and b (which would have argued against a failure here), but in turn, we did not see the apparent dip in the bottom panel of Fig. 10 when limiting the sample to the more robust stars with low V_g . We further compared the locus of the lower main sequence in absolute G -band magnitude for stars with $\beta \lesssim -55$ deg with equivalent fields in Galactic latitude and longitude, finding no significant difference. To summarize this paragraph: Assuming the dependence of $-\delta_p \sim \sigma_p$ makes all problems in n_{vis} and β go away, with some possibility for a problem at $\beta \lesssim -55$ deg, so in sensitive studies, one might want to touch these stars with a long pole.

Fig. 11 shows the distance bias $1+f$ versus the astrometric excess noise (top panel) and versus the RUWE (a quantity based on the chi-squared of the astrometric fit and stellar colour). To create these plots we used the usual limits in colour and magnitude, and waived the excess noise limit in the top panel. We note that




Figure 12. Distance bias $1 + f$ versus the bp-rp excess noise factor. We use the assumption $\delta_p = \sigma_p$ with $\delta\sigma_p = 0.043$ mas and scan the ordered sample with a mask of width 12 000 in steps of 4000 stars each.

only a very minor fraction of our stars have a non-zero excess noise value. These stars were censored from the sample. One could expect that at least some stars with larger excess noise should be binaries. In the discussion of Schönrich & Aumer (2017), the *Gaia* DR1 sample had a far larger temporal baseline for the astrometry compared to the v_{los} measurements, which raises the expectation of an apparent distance underestimate in the statistical distance estimator, since the v_{los} measurements would carry the additional velocity dispersion from the binary. Here, the case is less clear-cut. However, we can still note that stars with a positive excess noise have a slight tendency to distance underestimates in our method (understandable, since the astrometric effect is not clear-cut and we still have some binary dispersion affecting the estimator). If we admit stars with very large σ_p (which correlates with the excess noise), we find that the sample shows clear signs of a break-down for excess noise values larger than ~ 1 . Thus, we recommend applying a cut at this value. Further we note, the excess noise measurement is limited to stars with very good signal to noise in *Gaia* and thus leads to a concentrated sample in distance; neglecting this effect should cause slight distance overestimates in our selection, so the true distance bias on the large-excess-noise stars is likely slightly underestimated.

The bottom panel of Fig. 11 provides our two suggestions for distance evaluation versus the RUWE as defined in the additional release notes from *Gaia*. The usual argument is that stars with a large RUWE value are bad and should not be used. However, as we see from the figure, the only trend of $1 + f$ that we can detect is in the evaluation where we hold the parallax offset fixed at $\delta_p = -0.048$ mas. However, large values for RUWE (which is an expression for the quality of the astrometric fit) duly correlate with a larger σ_p given by the *Gaia* astrometric pipeline. If we correct for the trend in the *Gaia* parallax offset (green line), this slight bias versus RUWE vanishes. In short, we cannot see any reason for applying a quality cut in RUWE – of course, stars with large RUWE have worse measurements, but the *Gaia* pipeline appears to be perfectly fine in not producing any bias versus RUWE and mirrors the larger uncertainties in larger σ_p values.

Fig. 12 finally shows $1 + f$ versus the BP/RP flux excess factor E_{bprp} . The quantity is often cited as an important quality measure for *Gaia* data. It compares the flux in the BP+RP bands to the total flux in the *Gaia* *G* band and relies on their very similar coverage. Due




Figure 13. Distance bias $1 + f$ versus distance for the distance sets of Bailer-Jones et al. (2018) and McMillan et al. (2018). We use a sample size of 60 000 stars each, stepping by 20 000.

to the *Gaia* passband definitions, E_{bprp} increases on average for red stars, but primarily expresses contamination by neighbouring stars and background, or misidentifications. We note that here we only measure the average distance error; i.e., we are not concerned with single outliers. The result is quite clear-cut: Both very low values of $E_{\text{bprp}} < 1.172$ and large values of $E_{\text{bprp}} > 1.3$ signal compromised stars. We also note, however, that we needed to scan the sample with a very fine sample size mask of 12 000 stars moved in steps of 4000 stars each, since the number of compromised stars is so low in the sample with $b > 10$ deg. The error distribution within the bulk of the sample shows the usual number of outliers. Scanning with a larger sample size (100 000) on the centre of the distribution shows that all subsamples there have $|f| < 2$ per cent, with a slightly suspicious region almost exactly at $E_{\text{bprp}} \sim 1.2$.

6 COMPARISON TO PREVIOUS DISTANCE DERIVATIONS

It is appropriate to ask how our distances differ from the two previous distance derivations by Bailer-Jones et al. (2018) and McMillan et al. (2018), so Fig. 13 shows a distance scan of $1 + f$ versus their distance s . In both cases the distances were estimated under the assumption that the *Gaia* DR2 parallax zero-point is -0.029 mas, based on the quasar results from Lindegren et al. (2018), which was the best available estimate at the time they were computed. We have shown that this is an underestimate for these stars. It is therefore no surprise that Fig. 13 shows that both studies overestimated stellar distances substantially, with the typical increase of $1 + f$ with distance. In short, due to the parallax offset these distances are compromised and should not be used. The reason why the bias in the Bailer-Jones et al. (2018) distances is somewhat smaller than in the McMillan et al. (2018) distances is predominantly the fact that we do not have proper expectation values for Bailer-Jones et al. (2018). Instead, we have only the mode (maximum) of their posterior distance probability distribution, which, due to the skew probability distribution, is significantly smaller, for usual distance PDFs, than the expectation value. While this compensates for some of their intrinsic distance overestimate, the mode infers a difficult-to-predict and variable bias relative to the expectation value of a distribution; also, one should not rely on two

opposite biases to partly cancel. A longer discussion of this issue can be found in Schönrich & Aumer (2017).

7 THE DISTANCE TO THE PLEIADES

Since there has been so much discussion about the distance to the Pleiades, let us quickly analyse the effect that the offset has on their distance estimate. Historically, there has for a long time been a tension between a low-astrometric-distance estimate by *Hipparcos* (van Leeuwen 2009), placing the distance of the Pleiades at $s = 120$ pc with a formal uncertainty of 1.5 per cent, and results from isochrone fits to their photometry (Meynet, Mermilliod & Maeder 1993; Stello & Nissen 2001), as well as eclipsing binaries (Zwahlen et al. 2004; Southworth, Maxted & Smalley 2005), which placed their distance in the range $\sim(130 - 137)$ pc with similarly small uncertainties. The *Gaia* DR2 release (Gaia Collaboration 2018b) estimates the Pleiades' distance at 135.8 ± 0.1 pc. Applying our parallax offset to this estimate brings the distance moderately down to 134.8 ± 0.2 pc, and back towards the average of the stellar-physics-based estimates.

8 SUMMARY OF SUGGESTED QUALITY CUTS

Since the previous discussion was rather lengthy with many details, we provide here our suggestions for quality cuts on the *Gaia* sample, to ensure a minimization of the detected kinematic biases:

- (i) A colour cut $G_{BP} - G_{RP} < 1.5$ mag. To be entirely safe, we suggest $0.5 < G_{BP} - G_{RP} < 1.4$.
- (ii) A magnitude cut for $G < 14.5$ mag, and $G_{BP}, G, G_{RP} > 0$. A safer limit is $G < 12.5$ mag and $G_{RP} < 13.7$ mag.
- (iii) $p/\sigma_p > 4$; safer is $p/\sigma_p > 10$.
- (iv) $\sigma_p < 0.1$ mas with σ_p as given by the *Gaia* pipeline; safer is $\sigma_p < 0.07$ mas.
- (v) $n_{\text{vis}} > 5$ as pointed out in Lindegren et al. (2018) and excess noise < 1 .
- (vi) For the BP–RP excess flux factor, use $1.172 < E_{\text{bprp}} < 1.3$. Tighter cuts might apply if the number of outliers is important.
- (vii) $s > 80$ pc for studies that need assurance of distance systematics < 4 per cent.

Some notes: The σ_p dependence of the parallax offset is comparably well controlled as long as we choose a safe limit on σ_p . If one needs to use a looser limit on σ_p , we strongly advise to control for the dependence of δ_p on σ_p . If needed, the upper limit on the excess flux factor E_{bprp} can be relaxed a bit with proper caution (see Fig. 12). As detailed in Section 5.4, we find no reason to test or cut in RUWE, when applying our set of quality cuts on this sample. Note further that the faint magnitude limit and likely the colour limit are an imprint of the v_{los} measurement quality in the *Gaia* RV subset; i.e., these should be validated separately for different samples.

9 CONCLUSIONS

We have used the methods of SBA and Schönrich & Aumer (2017) to derive bias-free expectation values for stellar distances in the RV subset of *Gaia* DR2.² While *Gaia* parallaxes have been extensively

tested in Lindegren et al. (2018), their quasar sample has no overlap with our study in magnitude; thus the quasar sample is inadequate for testing distances in the important *Gaia* RV subset for two reasons: Quasars have all $p = 0$ – i.e. differences in the astrometric pipeline for $p > 0$ cannot be tested – and basically none of their quasars shares the same evaluation cohort as the stellar sample.

We derived Bayesian distances for all stars in the subset and validated the distance expectation values for stars with $p/\sigma_p > 4$, the minimum justifiable quality requirement for relative parallax error. All distances and derived kinematics will be made available with this work.

Our study provides clear proof for an average parallax offset $\delta_p = -0.054$ mas (*Gaia* parallaxes are too small) with negligible formal uncertainty and a systematic uncertainty of ~ 0.006 mas. The parallax offset is clearly identifiable as such as it results in an almost perfectly linear uptrend of distance bias with distance s that reaches values in excess of $f > 30$ per cent for $s > 3$ kpc. This offset is comparable to the findings of Zinn et al. (2018) using asteroseismic data. It is significantly larger than the value of $-\delta_p = 0.029$ mas found by Lindegren et al. (2018) for quasars. This also advocates a reanalysis of cluster distances. Even the very nearby Pleiades are pulled back by this offset by ~ 1 pc. Similarly, even a comparably benign bias of 10 per cent creates larger deviations of mean velocities, than found e.g. for the warp/wave pattern in the local Galactic disc. Every study using *Gaia* DR2 parallaxes/distances should investigate the sensitivity of its results on the parallax biases described here and – for fainter samples – in the DR2 astrometry paper.

We evaluated different assumptions for the parallax error in the *Gaia* pipeline and found that our estimate for δ_p is nearly unaffected by changing the parallax error. Not adding the additional error $\delta\sigma_p = 0.043$ to the astrometric pipeline value σ_p in quadrature decreases our estimate for $-\delta_p$ by about 0.006 mas.

As we used a self-informed prior for the distance-dependent selection function $S(s)$, the method provides a good approximation for $S(s)$, which we provide in equation (6). Assuringly, $S(s)$ displays the behaviour expected from population synthesis and the magnitude limits of the RV sub-sample of *Gaia*: an almost perfectly exponential decrease for $s < 1$ kpc related to the main sequence, a knee at intermediate distances, when the magnitude cut passes the level of the subgiant branch, and a slower decrease of $S(s)$ towards large s .

After correction for a constant parallax offset we still found a highly significant correlation of the distance bias f with σ_p . While this could point to a problem with the assumptions in our Bayesian distances, this explanation is unlikely since $S(s)$ can be measured to high confidence from the data. To the contrary, our results suggest that $-\delta_p$ is roughly proportional to σ_p (best-fitting value $q = 1.05$) after adding the 0.043 mas additional error to the *Gaia* parallaxes, and further show that stars with $\sigma_p \gtrsim 0.1$ mas should be discarded from analysis. It is unreasonable to think that the parallax uncertainty is added to the parallax value in a simple way, and so it is no surprise that the required factor is not constant in parallax: q depends significantly on p and around $p \sim 0.09$ mas, it is closer to $q = 0.5$.

Resolving this dependence of p on σ_p also removes a highly significant trend of our measured distance bias with the number of visibility periods n_{vis} , and consequently with ecliptic latitude β . We note that our method would still flag a distance problem at the ecliptic south pole. However, when used on a very narrow area on the sky, we lose most of our statistical corrections and do not trust the evaluation. In fact consistent with this expectation, when

²Please find the data sets with the derived distances and simple estimates of stellar velocities and positions in Galactic coordinates either in the MNRAS online materials or at <https://zenodo.org/record/2557803>.

limiting the sample to the more robust stars with low azimuthal velocity, this dependence was not confirmed, consistent with an evaluation of the derived HR diagram.

We further used the method to evaluate safe limits to be imposed on both apparent magnitude and stellar colour, finding that red stars at $G_{BP} - G_{RP} > 1.5$ are compromised as well as that stars with $G \gtrsim 14$ mag are flagged for distance underestimates. The most likely explanation is a decline in quality of the otherwise very well determined v_{los} .

We further tested for astrometric parameters, finding no biases related to RUWE (after removal of the aforementioned problems), and no strong correlation of f with astrometric excess noise values smaller than 1. At least for getting distance expectation values in the *Gaia* RV sample, this strongly argues for not using RUWE as a quality indicator. A mild decrease in distance estimates could point to stellar binaries.

A summary of all quality cuts is provided in Section 8.

ACKNOWLEDGEMENTS

We thank our referee, U. Bastian, for very thorough and insightful comments to the paper. It is a pleasure to thank Lennart Lindegren, J. Magorrian, J. Binney, F. van Leeuwen, and A. Mora for helpful comments. RS is supported by a Royal Society University Research Fellowship. This work was performed using the Cambridge Service for Data Driven Discovery (CSD3), part of which is operated by the University of Cambridge Research Computing on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The DiRAC component of CSD3 was funded by BEIS capital funding via STFC capital grants ST/P002307/1 and ST/R002452/1 and STFC operations grant ST/R00689X/1. DiRAC is part of the National e-Infrastructure. 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.

REFERENCES

- Antoja T. et al., 2018, *Nature*, 561, 360
 Arenou F. et al., 2018, *A&A*, 616, A17
 Astraatmadja T. L., Bailer-Jones C. A. L., 2016, *ApJ*, 832, 137
 Bailer-Jones C. A. L., Rybizki J., Foesneau M., Mantelet G., Andrae R., 2018, *AJ*, 156, 58
 Cropper M. et al., 2018, *A&A*, 616, A5
 Evans D. W. et al., 2018, *A&A*, 616, A4
 Gaia Collaboration, 2016, *A&A*, 595, A1
 Gaia Collaboration, 2018a, *A&A*, 616, A1

- Gaia Collaboration, 2018b, *A&A*, 616, A10
 Gaia Collaboration, 2018c, *A&A*, 616, A11
 Gaia Collaboration, 2018d, *A&A*, 616, A12
 Gillessen S., Eisenhauer F., Trippe S., Alexander T., Genzel R., Martins F., Ott T., 2009, *ApJ*, 692, 1075
 Graczyk D. et al., 2019, *ApJ*, 872, 85
 Gravity Collaboration, 2018, *A&A*, 615, L15
 Harris W. E., 1996, *AJ*, 112, 1487
 Huang Y. et al., 2018, *ApJ*, 864, 129
 Joshi Y. C., 2007, *MNRAS*, 378, 768
 Katz D. et al., 2019, *A&A*, 622, A205
 Kawata D., Baba J., Ciucă I., Cropper M., Grand R. J. J., Hunt J. A. S., Seabroke G., 2018, *MNRAS*, 479, L108
 Lindegren L. et al., 2018, *A&A*, 616, A2
 McMillan P. J., 2017, *MNRAS*, 465, 76
 McMillan P. J. et al., 2018, *MNRAS*, 477, 5279
 Meynet G., Mermilliod J.-C., Maeder A., 1993, *A&AS*, 98, 477
 Pietrinferni A., Cassisi S., Salaris M., Castelli F., 2004, *ApJ*, 612, 168
 Pietrinferni A., Cassisi S., Salaris M., Percival S., Ferguson J. W., 2009, *ApJ*, 697, 275
 Riello M. et al., 2018, *A&A*, 616, A3
 Sahlholdt C. L., Silva Aguirre V., 2018, *MNRAS*, 481, L125
 Salpeter E. E., 1955, *ApJ*, 121, 161
 Sartoretti P. et al., 2018, *A&A*, 616, A6
 Schönrich R., 2012, *MNRAS*, 427, 274
 Schönrich R., Aumer M., 2017, *MNRAS*, 472, 3979
 Schönrich R., Bergemann M., 2014, *MNRAS*, 443, 698
 Schönrich R., Dehnen W., 2018, *MNRAS*, 478, 3809
 Schönrich R., McMillan P. J., 2017, *MNRAS*, 467, 1154
 Schönrich R., Binney J., Dehnen W., 2010, *MNRAS*, 403, 1829
 Schönrich R., Binney J., Asplund M., 2012, *MNRAS*, 420, 1281(SBA)
 Southworth J., Maxted P. F. L., Smalley B., 2005, *A&A*, 429, 645
 Stassun K. G., Torres G., 2018, *ApJ*, 862, 61
 Stello D., Nissen P. E., 2001, *A&A*, 374, 105
 van Leeuwen F., 2009, *A&A*, 497, 209
 Zinn J. C., Pinsonneault M. H., Huber D., Stello D., 2018, *ApJ*, preprint (arXiv:1805.02650)
 Zwahlen N., North P., Debernardi Y., Eyer L., Galland F., Groenewegen M. A. T., Hummel C. A., 2004, *A&A*, 425, L45

SUPPORTING INFORMATION

Supplementary data are available at Zenodo: <https://zenodo.org/record/2557803>.

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

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