Gaia Data Release 3: Spectroscopic binary-star orbital solutions

E. Gosset, Y. Damerdji, T. Morel, L. Delchambre, J.-L. Halbwachs, G. Sadowski, D. Pourbaix, A. Sozzetti, P. Panuzzo, F. Arenou

A&A, 693, A124 (2025)
https://doi.org/10.1051/0004-6361/202450600

Astronomy & Astrophysics

Gaia Data Release 3: Spectroscopic binary-star orbital solutions

The SB1 processing chain

E. Gosset[1,⋆,⋆⋆], Y. Damerdji[2,1], T. Morel[1], L. Delchambre[1], J.-L. Halbwachs[3], G. Sadowski[4], D. Pourbaix[4,†], A. Sozzetti[5], P. Panuzzo[6], and F. Arenou[6]

[1] Institut d'Astrophysique et de Géophysique, Université de Liège, 19c, Allée du 6 Août, 4000 Liège, Belgium
[2] CRAAG-Centre de Recherche en Astronomie, Astrophysique et Géophysique, Route de l'Observatoire, Bp 63 Bouzareah, 16340 Algiers, Algeria
[3] Université de Strasbourg, CNRS, Observatoire Astronomique de Strasbourg, UMR 7550, 11 rue de l'Université, 67000 Strasbourg, France
[4] Institut d'Astronomie et d'Astrophysique, Université Libre de Bruxelles CP 226, Boulevard du Triomphe, 1050 Brussels, Belgium
[5] INAF – Osservatorio Astrofisico di Torino, via Osservatorio 20, 10025 Pino Torinese (TO), Italy
[6] GEPI, Observatoire de Paris, Université PSL, CNRS, 5 Place Jules Janssen, 92190 Meudon, France

Received 3 May 2024 / Accepted 8 October 2024

ABSTRACT

Context. The *Gaia* satellite constitutes one of ESA's cornerstone missions. Being primarily an astrometric space experiment measuring positions, proper motions, and parallaxes for a huge number of stars, it also performs photometric and spectrophotometric observations. *Gaia* operates a medium-dispersion spectrometer, known as Radial Velocity Spectrometer (RVS), which provides spectra and radial velocity (RV) time series.
Aims. The paper is focussed on the analysis of the RV time series. We fit orbital and trend models, restricting our study to objects of spectral types F-G-K that are brighter than a magnitude of 12, presenting only one single spectrum (SB1).
Methods. Suitable time series were processed and analysed on an object-per-object basis, providing orbital or trend solutions. The results of the various fits were further filtered internally on the basis of several quality measures to discard spurious solutions. The objects with solely a spectroscopic solution were classified in one of the three classes: `SB1` (eccentric model), `SB1C` (circular model), or `TrendSB1` (mere trend model).
Results. We detail the methods used in this work and describe the derived parameters and results. After a description of the models considered and the related quality tests of the fit, we detail the internal filtering process aimed at rejecting bad solutions. We also present a full validation of the pipeline. A description of the current content of the catalogue is also provided.
Conclusions. We present the `SB1`, `SB1C`, and `TrendSB1` spectroscopic solutions contained in the SB subcatalogue, part of the DR3 catalogue. We deliver some 181 327 orbital solutions in class `SB1`, 202 in class `SB1C`, and 56 808 in the associated class `TrendSB1`. This is a first release and the delivered SB subcatalogue could be further tuned and refined. However, the majority of the entries are correct. Thus, this data set constitutes by far the largest set of spectroscopic orbital solutions to be computed.

Key words. methods: data analysis – techniques: radial velocities – catalogs – binaries: spectroscopic

1. Introduction

The *Gaia* satellite constitutes one of the cornerstone missions of the European Space Agency (ESA). This astrometric space experiment is intended to represent a new breakthrough, compared to the previous pioneering one, HIPPARCOS (Perryman 2009). The *Gaia* satellite has two telescopes, observing at 106.°5 of each other, and three associated scientific instruments that are continuously scanning the celestial sphere thanks to its six-hour rotation (around an axis perpendicular to the two telescope's lines of sight), as well as the precession of the spinning axis around the direction of a fictitious, nominal Sun (compensating for the orbit of the satellite around the Sun-Earth Lagrangian point L2), with a 63-day period, and thanks to the one-year orbit of the satellite around the Sun, while it accompanies the Earth in its revolution (Gaia Collaboration 2016b). Multiple visits are necessary to reach a convenient signal-to-noise (S/N) ratio, but also to gather information on binarity (or higher order multiplicity) and variability.

The *Gaia* satellite is measuring the positions and transverse proper motions for some 1.8 billion stars, as well as precise parallaxes and thus distances for a very large number of objects. The detectors are working in a Time Delay Integration mode (TDI). The satellite was launched on 19 December 2013, and has been scanning the sky since the end of the commissioning period in July 2014. The detailed description of the project and of the payload can be found in Gaia Collaboration (2016b), published alongside the first Data Release (DR1). The description of DR1 and of the related data products is provided in Gaia Collaboration (2016a) whereas DR2 was detailed in Gaia Collaboration (2018a). The present paper is to be associated with the DR3 (Gaia Collaboration 2023d). In addition to astrometry performed with the Astrometric Field (AF) instrument (Lindegren et al. 2018, 2021), the satellite also produces

⋆ Corresponding author; eric.gosset@uliege.be
⋆⋆ Research Director, F.R.S.-FNRS (Belgium).
† Deceased 2021.

photometry from the same instrument as well as spectrophotometry provided by another channel: the BP/RP spectrophotometer purveyor of very low-resolution spectra as well as of photometric colours (Riello et al. 2018, 2021; Evans et al. 2018). The two above-mentioned channels are supposed to observe stellar objects down to a magnitude of about 20 in the G band.

In addition to astrometry and spectrophotometry, the satellite also incorporates a third channel: a medium-resolution spectrometer to measure the radial velocities (RVs) of stellar objects, known as the Radial Velocity Spectrometer (RVS). This channel and the related instrument are described in great detail in Cropper et al. (2018), which also includes a comprehensive review of the genesis of the RVS instrument. The RVS brings the third component of the space velocity (to join with the transverse proper motion on the sky). It also provides, for the brightest stars, information on the nature of the object ($T_{\rm eff}$, $\log g$, chemical abundances, chromospheric activity, etc.). Stars brighter than a magnitude of $G^{\rm int}_{\rm RVS} = 2.76$ produce saturated pixels and, thus, are not considered here (for the definition of the $G_{\rm RVS}$ magnitudes, see Sartoretti et al. 2023).

The scanning law for RVS is intended to typically produce 40 transits over the nominal five-year mission. The focal plane of the RVS is covered by 12 CCDs (four rows of three in a row). The RVS instrument acquires, in TDI mode, spectra in the range 845–872 nm, including the well-known Ca II triplet (in late-type stars). With a resolving power of about $R \sim 11\,500$ (see Sects. 8.2 and 11 of Cropper et al. 2018), it obtains three spectra per transit, each corresponding to a maximum exposure time of 4.42 s (they are analysed individually and concomitantly). These spectra are spread over about 1100 pixels (Cropper et al. 2018), but the part used for RV determination is restricted to the range 846–870 nm (Sartoretti et al. 2023). Very bright objects benefit from 2D spectra whose across scan component is 10 pixels.

DR2 was the first release concerning the spectrometer channel and some 280 million spectra were treated and resulted in the publication of median (over transits) RVs for over 7 million objects (brighter than mag. 12, being single-line spectra and corresponding to assumed constant RVs) in the range of effective temperatures [3550, 6900] K. This catalogue of single-line, non-exotic stars has already been widely used to study the structure of the Galaxy (Gaia Collaboration 2018b; Antoja et al. 2018; Kawata et al. 2018; Ramos et al. 2018; Monari et al. 2019; Koppelman et al. 2019; Marchetti et al. 2019; Trick et al. 2019; López-Corredoira & Sylos Labini 2019; López-Corredoira et al. 2020; Khoperskov et al. 2020; Nitschai et al. 2020).

The data coming out of the RVS are treated by ground-based spectroscopic processing. The spectroscopic processing is ensured by the Coordination Unit 6 (CU6) of the Data Processing and Analysis Consortium (DPAC; Gaia Collaboration 2016b). The first description of this processing was given in Sartoretti et al. (2018) for DR2 and the properties and validation of the RV set were described in Katz et al. (2019). The stars exhibiting a double-line spectrum (also known as the composite spectrum) were rejected from the process (as described in Sartoretti et al. 2018); (Damerdji et al., in prep.). The spectroscopic processing pipeline is made of a series of modules, each taking into account a series of tasks. The first modules are working per transit and per trending epochs. The flowchart is illustrated in Sect. 2.3 of Sartoretti et al. (2018). The single transit analysis (STA) refers to the workflow that actually performs the measurements of the RV per transit. Three submodules are measuring the RVs of all the selected convenient stars (single spectrum ones). They are based on the idea to derive the RV by cross-correlating the observed spectrum with a synthetic spectrum built in `generate_template` (Sartoretti et al. 2018). The incentive in making use of several methods is explained in the early investigation of the problem, with conclusions detailed in David et al. (2014) (see also Sect. 7 of Sartoretti et al. 2018). The STA measures the RVs per object and per transit. Thereafter, the multiple transit analysis (MTA) workflow performs some statistical analysis tests based on the RVs from all objects at all transits. It computes (per star) a probability of constancy that is used to spot out variable objects. It also computes the median of the RVs over all transits for all the bright objects (Sartoretti et al. 2018). The constancy probability is not used to define the list of assumed constant objects that are considered to compute their median RV. The updated version of the spectroscopic chain for DR3 is described in Sartoretti et al. (2022); Katz et al. (2023). In DR3, MTA further computes mean spectra for assumed constant objects (Seabroke et al., in preparation).

In the context of DR3, the RV values corresponding to variable objects are transmitted to the CU4 (object processing) and in particular to the non-single star (NSS) part of the object processing that analyses these data in order to derive possible orbital solutions. Three channels running in parallel deal with the astrometric data (Halbwachs et al. 2023; Holl et al. 2023b), photometric data (Siopis et al., in prep.; Pourbaix et al. 2022, see Sect. 7.6 by Siopis & Sadowski), and the spectroscopic data (the present paper; Damerdji et al., in prep.). The detailed description of the CU4 spectroscopic processing chain for SB1 (spectroscopic binaries with only one spectrum visible; hereafter labelled NSS-SB1) orbital solutions or trend ones is the subject of the present paper whereas Damerdji et al. (in prep.) describe the SB2 channel. These DR3 products are preliminary results; this is the first time that orbital solutions issued from *Gaia* have been delivered. The results of the three channels can be further processed by combining their independent outputs. This is the task of the combiner (Gaia Collaboration 2023a; Pourbaix et al. 2022, see Sect. 7.7 by Gavras & Arenou). The scientific conception of the SB1 chain described here was elaborated on by Yassine Damerdji, Eric Gosset and Thierry Morel; most of the codes were benchmarked in MATLAB and finally written in Java by YD with some help from TM; the validation tests (including the critical building of the ground-based data catalogues used for comparison) were mainly performed by TM. The whole work was carried out in the highly favourable and boosting context of the NSS team.

Section 2 describes the nature of the data that are entering the process, whereas Sect. 3 describes the simple mathematical formalism used and lists the models that are considered. Section 4 explicits the importance of the period search and the related difficulties. Section 5 briefly summarises the main trials we did to arrive at the method described in Sect. 6, where we present the chain as adopted for the DR3 (NSS-SB1 main processing). In Sect. 7, we introduce the notion of post-processing. Section 8 illustrates the results. Section 9 concerns the validation of the results of the chain. Additional considerations in Sect. 10 explain the details of the post-filtering and provide additional information in the light of the conclusions from the previous section. The catalogue description, the caveats and the precautions necessary in using the results of the NSS-SB1 processing are given in Sect. 11. Section 12 presents our conclusions. The contents of the present paper are more detailed and contain material that supplants the information given in the *Gaia* archives (see Sects. 7.4 and 7.5 in Pourbaix et al. 2022)[1] at the time of the release. The main core of the paper is accompanied by

[1] https://gea.esac.esa.int/archive/documentation/GDR3/

appendices. Appendices D to J are only available on Zenodo (see https://zenodo.org/records/13990211).

2. Description of the input data

The DR3 data correspond to the first 34 months of the nominal mission (25 July 2014–28 May 2017, almost 1038 d, 2.84 y). The instrumental spectra obtained at each RVS transit are converted to proper physical spectra (reduced) by the pipeline developed by the CU6 (spectroscopic processing), see Sartoretti et al. (2022); Katz et al. (2023). Per transit, one spectrum is acquired by each of the three CCDs in a row of the RVS (Cropper et al. 2018). The RVs are measured individually on the three spectra. Three methods are used to derive the radial velocities (Sartoretti et al. 2018), each of them providing three velocities (one per CCD) and a combined RV value (over the three CCDs). The three combined RVs issued from the three methods are transmitted to the *Integrator* that provides a single combined value (per transit). Only this single combined value is considered here but the other intermediate RVs are used as a sanity check to validate the final combined value (see Sects. 6.4.8.2 and 6.4.8.5 of Sartoretti et al. 2022). The RVs are measured by cross-correlating the observed reduced and calibrated spectra with a theoretical/synthetic spectrum computed on the basis of various stellar atmosphere models. These so-called templates are identified through a set of three atmospheric parameters (effective temperature, surface gravity, metallicity). The selection of the template used to measure the RVs is based on these three parameters only; the alpha element content is not considered as an independent parameter. The assignation of the set of parameters to the object treated is performed on the basis of three databases. They are ordered here in the order of preference. The first possible origin is an internal compilation of ground-based catalogues. The second one is from an early run of the *Gaia General Stellar Parametrizer from Spectroscopy* (GSP-Spec, Recio-Blanco et al. 2023). The third one is from an early run of the *Gaia General Stellar Parametrizer from Photometry* (GSP-Phot, Andrae et al. 2023). Finally, if none of these sources provides the necessary parameters, the latter were estimated as previously done for DR2 (*DetermineAP*: Sect. 6.5 of Sartoretti et al. 2018). This process *DetermineAP* cross-correlates the observed spectrum with each template from a limited subset of 28 selected ones, and this for every transit. The template giving per transit the highest correlation is retained, whatever is the velocity. The final adopted template is the one that is the major contributor over all the transits. More details can be found in Sect. 3.6.1 of Katz et al. (2023). The repartition among the various origins for the data set treated here is the following: ground-based catalogues (14.4%), GSP-Spec (33.4%), GSP-Phot (45.6%) and *DetermineAP* (6.6%). The RV measurements from transit to transit are performed with the same template. Consequently, a template mismatch is expected to essentially introduce a systematic RV shift (for single-line objects).

The measurements of the RVs were performed under the supervision of the STA Development Unit. The list of good RVs corresponding to a particular star are then analysed by the MTA to detect and separate variable RV sets from constant RV sets. For each star, the median of the RVs is computed and appears in the main catalogue. The stars with variable RVs were further considered by the global pipeline and forwarded to the CU4 (Object processing: NSS). For DR3, the adopted cutting threshold corresponds to a probability of variability of 0.99 (`rv_chisq_pvalue` ≤ 0.01). The composite spectra and SB2 objects were already treated separately at the level of the CU6 processing (Damerdji et al., in prep.) for the determination of both RVs and are then dispatched to CU4/NSS in order to search for the corresponding orbital solutions (Damerdji et al., in prep.).

The non-single star (NSS) processing is intended to perform the reduction of the corresponding data from the point of view of multiple star studies. In the present paper, we consider the sole spectroscopic channel for the SB1 objects.

The input data were drawn from a list of RVs (and their 1σ-uncertainties) that were corrected to make reference to the barycentre of the solar system. It should be pointed out that for DR3 the rms error on the barycentric correction ($\sigma \lesssim 0.05\,\mathrm{km\,s^{-1}}$) was not propagated to the uncertainty on the individual RVs.

The RVs are accompanied by a time of observations expressed in barycentric Julian days (BJD). The number of points of the time series corresponds to the number of transits for which the RV has been properly measured with success. The data could present several weaknesses such that outliers and wrong RVs due, for instance, to template mismatch errors (for the latter case, most likely the same systematic for all values within a time series). These problems and additional selection rules are further documented in Katz et al. (2023) and will not be further detailed here. It should however be noticed that some objects have been rejected at the level of the input data filtering in CU6 because they correspond to a poorly performing *Astrometric Global Iterative Solution (AGIS)* solution or to a large error on the along-scan field angle η (see its definition in Sartoretti et al. 2018). The threshold has been chosen at $\sigma_\eta = 200\,\mathrm{mas}$, equivalent to about $29\,\mathrm{km\,s^{-1}}$. Other filtering applied to the input CU6 data is described in Sect. 4 of Katz et al. (2023). In the course of the CU6 STA/MTA validation, a bias on output RVs was noticed for stars fainter than $G_{\mathrm{RVS}}^{\mathrm{int}} = 11$ (being $0\,\mathrm{km\,s^{-1}}$ at magnitude 11 and reaching $75\,\mathrm{m\,s^{-1}}$ at magnitude 12, Katz et al. 2023). The correction was not applied for the catalogue, nor for the transmission of the RVs to the NSS processing. The effect is quite small and should only induce a global shift of all the RVs per object along with, as a consequence, a small bias on the systemic velocities (see Sect. 3).

The current astrometric solution for DR3 has been released as part of the EDR3 and contains, among others, positions for a large number of objects. The astrometric solution (AGIS) is based on a treatment assuming the objects are single, unresolved stars and thus having an image characterised by the instrumental line-spread function (for the 1D observations in the astrometric channel) or point-spread function (for the 2D observations), as described in Lindegren et al. (2021). Some objects do not behave as expected according to this assumption. The pathological behaviour could be due to anomalies, but a large contribution to this problem originates in unresolved or poorly resolved double (or multiple) objects. The separation threshold is considered to be around $0\farcs4$–$0\farcs5$ (see Fabricius et al. 2021; Holl et al. 2023a). The double objects not following the behaviour for single stars are associated with a bad adjustment to the global astrometric solution and are thus accompanied by a large renormalised unit weighted error (`ruwe`). Lindegren et al. (2021) also associated an excess noise (impacting the position) to these objects, which is delivered in the catalogue under the name `astrometric_excess_noise` and which corresponds to the additional dispersion necessary to accept the concerned object in the global solution. Objects with a `ruwe` larger than 1.4 are certainly problematic, but smaller values already have an impact on the RVS calibration. Since the RVS is a slitless spectrometer, the calibration in wavelength of the spectra is requesting an epoch position of the object that is accurate enough. Therefore,

for the objects that are deviating from the single star behaviour, we could expect that the wavelength calibration is not correct, and thus the measured transit RVs are somewhat biased and this creates artefacts in the time series. No correction of this effect is offered for DR3. As a first approach, we preferred not to discard these objects on the basis of the `ruwe`. In order to take into account this problem, it was decided to inflate the uncertainty on the RV value by introducing (quadratically) an additional component linked to the astrometry. This term is derived as being $\sigma_{\rm add}$ (km s^{-1}) = 0.146 × `astrometric_excess_noise` (mas) where the factor 0.146 is related to the dispersion of the spectrometer (transformation of mas into km s^{-1}). Typically, the values of `astrometric_excess_noise` are below 10 mas and, thus, $\sigma_{\rm add}$ is below 1.46 km s^{-1}.

Some selections were also made on the objects at the moment of their entry in the NSS-SB1 processing. In the case of DR3, these kinds of stars are restricted to the range of effective temperatures 3875 K to 8125 K (mainly ((M))-K-G-F-(A) stars). No emission-line object was selected and, in principle, only stars without spectral peculiarities were permitted. The objects with a `rv_renormalised_gof` value greater than 4 were selected (Katz et al. 2023) to enter the NSS-SB1 processing because they were considered as clearly non-constant. The measurements of the individual (per transit) RVs are restricted to the range –1000 to +1000 km s^{-1}. A global view of the selections applied beforehand is available in Table H.4. This measurement of the individual RVs on individual spectra is only possible for bright objects. The distribution in magnitude of the treated objects is provided in Fig. 15 in the validation Sect. 9.2.1. A threshold on $G_{\rm RVS}^{\rm int}$ was selected at 12 mag. Fainter objects are considered to produce individual spectra for which the S/N ratio does not enable the derivation of a sound radial velocity. This strongly reduces the amount of selected objects. The stellar RVs are not processed by the NSS spectroscopic channel if the number of transits (data points) is less than 10 (see details in Sect. 6.3.1).

The nature of the measurement itself (cross-correlation with templates) implies that it is not possible to discriminate between a change in RV due to a global shift of the line (as expected for SB1 objects) and a shift due to a line-profile variation (as expected e.g. for some intrinsic variables). Although possibilities to get rid of this problem exist, they are not implemented for DR3 and this certainly constitutes a limitation of the present orbital-solution pipeline. In particular, fake SB1 could persist in the sample corresponding to never-deblending SB2.

When the SB2 chain (Damerdji et al., in prep.) analyses an SB2 time series and is not able to derive the corresponding SB2 orbital solution, the object can be redirected to the SB1 chain. An SB1 time series is built with the RVs characterised by the smallest uncertainties. The object is then considered as an SB1 and is forwarded to the present chain. This is a second possibility of entering the SB1 channel. These objects are flagged (for more details, see Damerdji et al., in prep. and also Appendix I).

3. The underlying orbital models

Four kinds of orbital solutions are considered here: they are labelled SB1, SB1C, TrendSB1, and StochasticSB1. They are described here below. The single-line RVs treated here and consequently the deduced motion is intended to refer to the sole object that is visible in the spectra. The model will be fitted through a classical least-square procedure and the selected objective function is the χ^2-function except when otherwise stated. The uncertainties on the parameters are calculated in the classical way on the basis of the diagonal elements of the inverse of the curvature matrix (i.e. the variance-covariance matrix; see Bevington & Robinson 2003). Despite the fact that we use the Greek symbol σ for all (i.e. concerning input data as well as output parameters) uncertainties, it should be pointed out that all of them are of course estimators. They are always expressed as 1σ values.

3.1. Orbital models of type SB1

The general eccentric Keplerian model is, for the spectroscopic channel, expressed in terms of the Campbell coefficients (see the appendices in Halbwachs et al. 2023; Binnendijk 1960, as the seminal reference). The model RV is given by

$$RV(t) = \gamma + K\,[\cos(v(t)+\omega) + e\,\cos\omega\,], \tag{1}$$

where $v(t)$ is the true anomaly, which is deduced from the eccentric anomaly $E(t)$ by

$$\cos v(t) = \frac{\cos E(t) - e}{1 - e\,\cos E(t)}, \tag{2}$$

$$\sin v(t) = \frac{\sqrt{1-e^2}\,\sin E(t)}{1 - e\,\cos E(t)}, \tag{3}$$

which in turn is expressed as a function of the mean anomaly $M(t)$

$$M(t) = E(t) - e\,\sin E(t) = \frac{2\pi}{P}(t - T_0). \tag{4}$$

Here, T_0 is the time of passage at periastron, while ω is the argument of periastron measured in the sense of the orbit with its origin at the ascending node. In this model, the eccentricity e enters as a strongly non-linear parameter. This is also the case for T_0. The parameters to be determined are P, γ, K, e, ω, T_0 (respectively, the period, the centre-of-mass velocity, the semi-amplitude, the eccentricity, the argument of periastron and the time of passage at periastron). The first step of the computation is devoted to the determination of the period. Some significance tests are made on the very existence of the period. If the period is not significant, the model is fitting noise and is considered as being invalid. This is a statistical decision. Owing to the insufficient span of time covered for the DR3, no apsidal motion is considered; ω is considered constant in time.

3.2. Orbital models of type SB1C

Another particular kind of orbit is the circular one, i.e. $e = 0$. In this case, the model is much more simple and is expressed by the following equation

$$RV(t) = \gamma + K\left[\cos\left(\frac{2\pi}{P}\,(t - T_0)\right)\right], \tag{5}$$

where here T_0 is the time of maximum velocity and the ω is fixed at zero corresponding to the maximum velocity. The free parameters to determine are restricted to P, γ, K, T_0. The model is linear except, of course, for the determination of the period. Concerning the T_0, the model can be linearised by using the following equivalent equation

$$RV(t) = \gamma + A\left[\cos\left(\frac{2\pi}{P}\,t\right)\right] + B\left[\sin\left(\frac{2\pi}{P}\,t\right)\right]. \tag{6}$$

We then computed

$$K = \sqrt{A^2 + B^2} \tag{7}$$

and

$$T_0 = \frac{P}{2\pi} \times \text{atan2}(-B, A), \tag{8}$$

where this T_0 refers to a cyclic variable and is assigned the value closest to $t = 0$. The significance level of the period is here also determined.

3.3. Models of type TrendSB1

It can happen that the RV evolution as a function of time just exhibits a trend. This could be a transient behaviour (e.g. outburst) just observed by chance or correspond to a small piece of slowly accelerated motion (e.g. of a stellar atmosphere). It could also correspond to an orbital motion with a period well longer than the span of observing time covered by the satellite. This span of time is 34 months for the DR3. However, it could be slightly shorter for some objects because it combines with the gaps in the scanning law. The present modelisation is restricted to polynomials of one to four degrees and presenting only one extremum on the interval. The degree of the polynomial is chosen on the basis of an F-test for nested models. In practice, it turned out that only the one- and two-degree polynomials were necessary. No higher degree fit was required. The formalism used for the fit is the same as for Hipparcos (ESA 1997). The TrendSB1 solution is compared to the orbital solutions using a Fisher Snedecor F-test on the comparison of the respective χ^2s. It should be made clear that this comparison could be an approximation.

3.4. Models of type StochasticSB1

When the fit of a TrendSB1 solution is not possible nor the fit of an SB1 or SB1C orbital solution corresponding to a significant period, the solution is classified as StochasticSB1. The median RV is considered as constant, but with an extra dispersion compared to the adopted RV uncertainties. This extra dispersion is estimated. Properly speaking, this simple model is not an orbital one. Some badly sampled orbits and high-order multiple systems could fall into this category, perhaps momentarily. Some of them will certainly benefit from an increase of the number of transits in future releases. Short-period intrinsic variables (e.g. δ Scuti stars) may also receive such solutions.

3.5. The downstream combination of solutions

The outputs of the NSS spectroscopic pipeline are delivered to the combiner that operates downstream. The task of the combiner is to inspect the output of the astrometric, photometric, and spectroscopic NSS channels in order to detect possible solutions corresponding to the same object and that could be combined. This will be described in Sect. 7.2.

4. Determination of the periodicity

The first step in the analysis is the search for a periodic behaviour of the RVs. The analysis is based on a periodogram (actually a frequencygram), which has the nature of the Fourier power spectrum. Since the time sampling of the *Gaia* satellite is particularly complex, the classical Fourier formalism is not applicable. To aptly understand the limitation of the method, it is necessary to have some knowledge of the impact of the sampling and of a few characteristic timescales.

4.1. Effect of the gaia sampling

To explain the effect of the *Gaia* sampling, we have to start from a single, idealistic, case and progressively add certain degrees of complexity. We start by considering the spinning of the satellite on itself that makes that a star is measured every 0.25 d (this is our first basic idealistic case since this simple view is never occurring in practice). The existence of the sample periodicity concerns both fields of view individually. Therefore, 0.25 d is the typical, smallest regular and periodic timescale in the sampling. The 0.25 d step should generate in the time series some aliasing with a typical 2 d^{-1} Nyquist frequency. This means that a progenitor frequency ν will be duplicated at $4 - \nu$, $4 + \nu$, $8 - \nu$, $8 + \nu$ and so on (in d^{-1} units). The duplication is also extended to the negative frequencies. This duplication represents a pure aliasing, since the sampling is in this hypothetical case regular. Therefore, the aliasing is pure and the Nyquist frequency exists. In such a case, intermediate and large periods could have aliases towards periods of 0.25 to 0.5 d (and vice versa).

However, the existence of two fields of view separated by 106.°5 on the sky renders the situation more complex by adding a regularity that is not properly speaking periodic. Thus, the aliasing is no more pure, we will call it pseudo-aliasing. The semi-amplitude spectral window is a tool that facilitates the analysis of such phenomenon. It has been introduced by Deeming (1975, 1976). In the case of an equidistant sampling, the spectral window is the well-known Dirac comb (or Shah function; see Bloomfield 1976). The general spectral window represents an extension of this Dirac comb to the general case of odd sampling. The property that the generalised Fourier Spectrum of the sampled data is the result of the convolution of the underlying Fourier Spectrum of the signal by the spectral window remains true for almost any kind of sampling (Deeming 1975, 1976). If the satellite was limited to a rotation on itself, the semi-amplitude spectral window corresponding to the combined two fields of view would resemble the plot in the upper panel of Fig. 1. Contrary to the classical case of pure aliasing, the aliasing is here not pure since all the peaks have different heights. Besides the main peak at $\nu = 0$, the next first peak is at $\nu = 4\,d^{-1}$, but it is mildly high. The peak having a height comparable to the $\nu = 0$ peak is the one at $\nu = 12\,d^{-1}$ which could represent some almost pure aliasing. This figure represents an idealistic case. First, it is assumed that the data points in each field of view are perfectly regularly distributed. This is not the case because several phenomena induce a small perturbation of the time of the observations. The variations of the Basic Angle as well as micro-clanks and meteoritic-hits are perturbing the regularity (Gaia Collaboration 2016b). We do not further consider these effects in our explanation. The satellite is not only spinning on itself, but also follows a more complicated motion. First, the rotation axis is precessing along an axis directed towards a fictitious nominal position of the Sun (compensating for the orbit of the satellite around the Sun-Earth Lagrangian point L2). In addition, the satellite located around the Lagrangian point L2 is orbiting around the Sun with the Earth annual motion (change in ecliptic longitude of the Sun, Gaia Collaboration 2016b). The major consequence is that, for a given object, the majority of the turns of the spinning satellite will not correspond to a transit of the object and thus no observation will be acquired for this object.

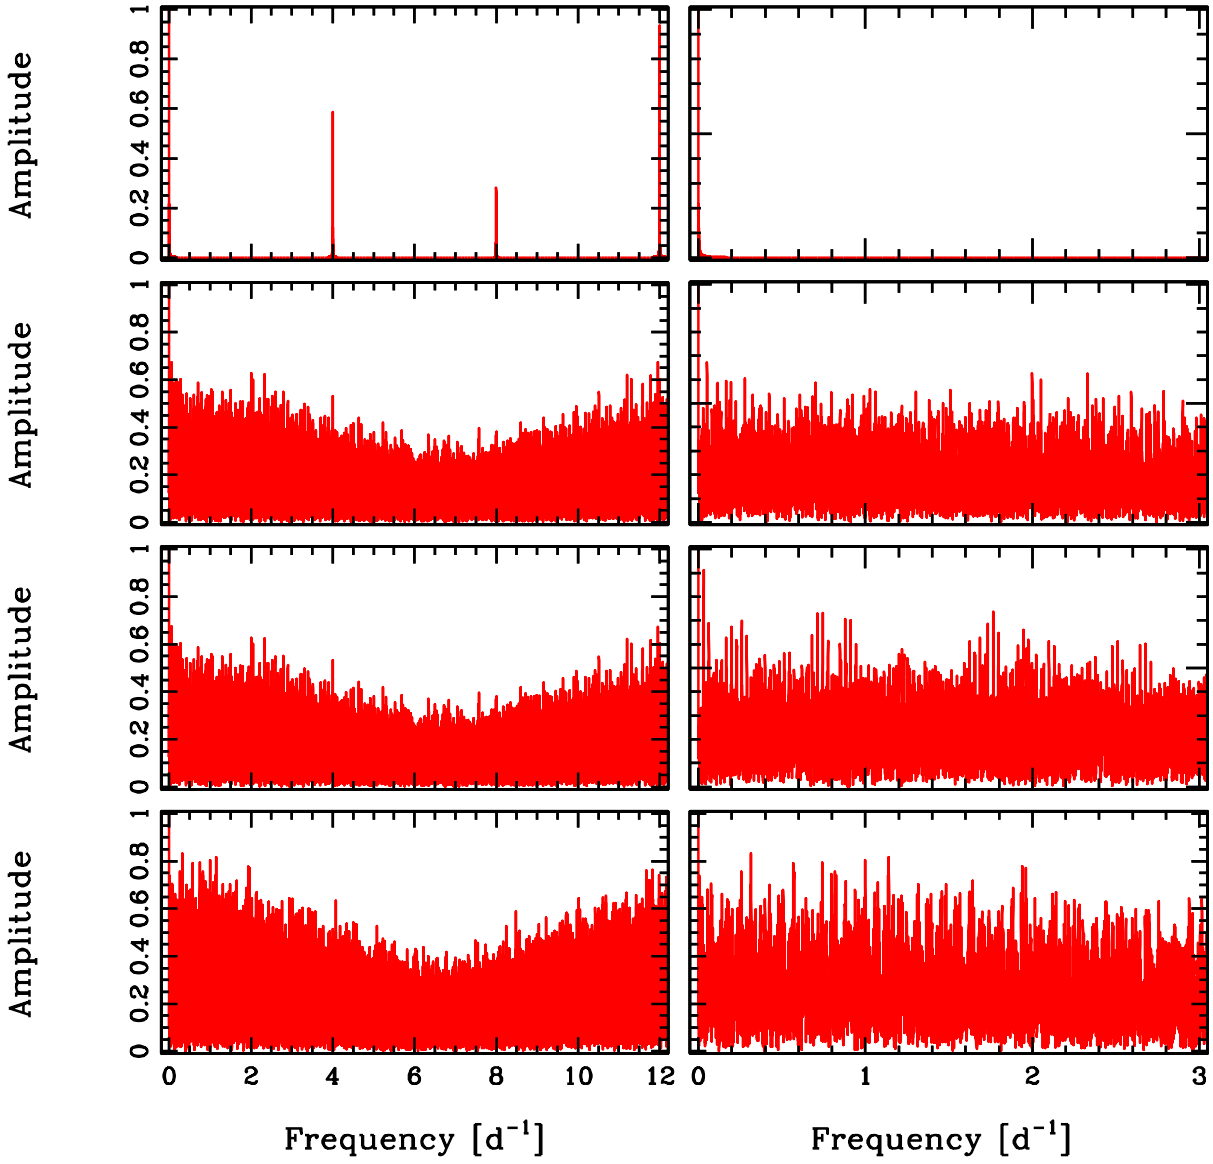

Fig. 1. Amplitude spectral window corresponding to various sampling. By definition, the peak at $\nu = 0\,\mathrm{d}^{-1}$ is always going up to the value 1. The upper panel corresponds to some 8000 data points and illustrates the sampling that *Gaia* would generate if the sole effect would correspond to the spinning of the satellite on itself. The three other panels represent typical spectral windows corresponding to objects at various positions in the sky. The number of data points are, from the second panel to the bottom one, respectively, 26, 19 and 14. The right panels are zooming on the low frequency domain. All the individual peaks have by definition the same shape and the same width. They are not resolved in this figure.

Most of the observations are not done and medium and large gaps occur during which the field of view is not directed towards the star. This effect is dependent on the position of the object on the sky (Gaia Collaboration 2016b). Therefore, the sky sampling is markedly non-uniform (see Sect. 5.2 of Gaia Collaboration 2016b).

Practically, this generates in the spectral window a power leakage from the main peak at $\nu = 0\,\mathrm{d}^{-1}$ towards other frequencies. This effect is illustrated in the other panels of Fig. 1. We produce typical time series of *Gaia* RVS with the GOST tool[2]. Fig. 1 second panel represents a typical spectral window of a time series containing 26 observations, the third panel 19 and the bottom one 14 observations (against the ~8000 observations for the first panel). This small number of observations induces an aliasing but the aliasing typical of the *Gaia* sampling is not leaking power towards a very few other frequencies but towards numerous frequencies. Besides the $\nu = 0$ main peak, there is a leakage towards a forest of frequencies. The contrast between the parasitic peaks and the main one is decreasing with the number of data points (following a $1/\sqrt{N}$ for the semi-amplitude). Typically, pseudo-aliasing can occur at any $\delta\nu$. If existing, the value of $\delta\nu$ is most probably located below let say $3\,\mathrm{d}^{-1}$. The structure due to the motion considered in the top panel of Fig. 1 almost totally disappeared. In the power spectrum of the signal, the combination of noise and of aliasing is going to generate parasitic peaks at frequencies that mainly depend on the position of the object on the sky. Consequently, the time sampling is complex, and the aliasing is dependent on the star being observed.

Looking at the third panel of Fig. 1, we notice a marked aliasing at $\delta\nu = 0.0317\,\mathrm{d}^{-1}$ and at $\delta\nu = 0.0634\,\mathrm{d}^{-1}$. Therefore, in a few cases, noise could conspire to shift the largest peak position

[2] https://gaia.esac.esa.int/gost/

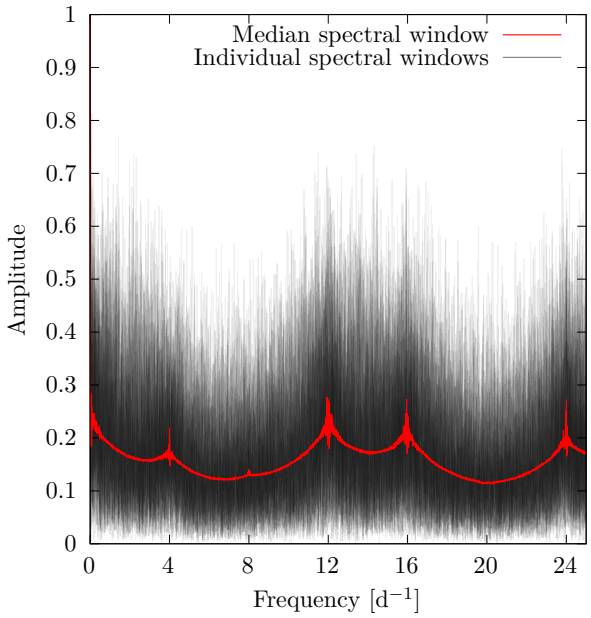

Fig. 2. Median of the simulated amplitude spectral windows (in red). In black are shown some 32 of the individual realisations in order to illustrate the marked dispersion among the population of simulated spectral windows. By definition, all the peaks at $\nu = 0\,\mathrm{d}^{-1}$ are here also going up to 1.

from the progenitor peak by this amount, thus producing pseudo-aliases of the progenitor that could dominate. However, in the present case, predominant aliasing seems to be restricted to small $\delta\nu$, i.e. to small shifts of the progenitor peak. Other values of $\delta\nu$ will be associated (or not) to other regions of the sky. Large shift aliasing is clearly possible but rather for values $\delta\nu = 12\,\mathrm{d}^{-1}$. This should have no impact in the framework of spectroscopic binaries.

Here above, we have shown three examples illustrating the strong leakage characterising the *Gaia* sampling (for RVS). An alternative view would be to compute some kind of typical spectral window. We obtain such a result using simulations. We generated 65 536 positions randomly and uniformly distributed over all the sky. For each position, we extracted from the GOST software the series of observational epochs typical of this particular position and corresponding to the nominal scanning law (NSL) for RVS. On purpose, we discarded the first month of scanning which was characterised by the ecliptic pole scanning law (EPSL). For the latter, the precession of the satellite was stopped and the scanning law clearly becomes more relevant to the scanning due to the spinning of the satellite, except for a slow rotation due to the orbital motion around the Sun. We thus neglected the EPSL to instead concentrate on the effect of the NSL. For each simulation, we derived the corresponding amplitude spectral window (for the RVS instrument). Figure 2 illustrates (in red) the median of the amplitude spectral windows. Despite the fact that we restricted our simulations to the NSL, the aliasing due to the spinning of the satellite is clearly visible. However, even the peaks at $\nu = 12\,\mathrm{d}^{-1}$ (and at $\nu = 16\,\mathrm{d}^{-1}$) are rather low, compared to the peak at $\nu = 0\,\mathrm{d}^{-1}$, at most 30% (in amplitude). We also plot in Fig. 2 (in black) 32 randomly selected amplitude spectral windows. Clearly, the dispersion of the various realisations of the amplitude spectral window is rather large. Also, near the possible alias frequencies, the amplitude spectral window could present values nearing zero. This confirms the fact that in most cases, the spinning pseudo-aliasing will have little, if any, impact on the global aliasing. Although a few positions on the sky could still exhibit the effect of the spinning aliasing. Of course, if, over the 34 months, the EPSL scanning is dominating the definition of the times in the analysed time series, the effect will increase exhibiting a stronger spinning aliasing. We must insist here on the fact that no signal is analysed in the present section; the dispersion of the amplitude spectral windows is originating from the randomness of the selection of the sky positions and from the variety of nature of the aliasing as a function of the latter. For any time series, the observing time could be used to pre-calculate the particular spectral window independently of the signal.

4.2. Search for the periodicity

When a periodicity is present in a time series, the periodogram exhibits a peak at the corresponding frequency (the progenitor), but other peaks are also present at the location of the pseudo-aliases of the progenitor. Contrary to the case of even sampling, the actual sampling generates pseudo-aliases that are in principle lower than the progenitor peak. However, the combination of the aliasing with background noise or other phenomena (blending) induces the fact that the choice of the true progenitor among the aliased peaks is never fully secure.

The total span of time of the DR3 time series is at most of 34 months and thus the largest periodicity accessible from the point of view of the Fourier corresponds roughly to a frequency of 1/1000 d^{-1}. This value could actually be larger taking into account the gaps generated by the actual sampling law. The constraints should thus be expressed as a function of the actual length of the time series and thus as a function of ΔT. It has been decided that it would be adequate to limit the range of investigation for periodicities to $1.5\,\Delta T$ since beyond the possible orbit is certainly not well determined. However, from the practical point of view, we decided to limit the range for period search to $2.0\,\Delta T$ although no orbit are computed for periods located in this extension (see Sect. 6.3 for more details). The adopted frequency step is actually $0.01/\Delta T$ where ΔT is the actual total time-span of the individual time series. This represents an oversampling of a factor of 100, which gives direct access to the amplitude of the signal as the height of the peak. Based on the discussion detailed in the previous section, we decided that the adopted largest frequency to be explored is 4 d^{-1}. This choice is rather atypical since it was chosen to correspond to the spin period.

If the star indeed varies in a periodic way, the largest peak in the periodogram is an estimate of the possible period, but does not stand as definitive proof of its existence. We should consider that a periodicity is proven if a minimum of three identical cycles are observed. Although the long periods are indicative, they are associated to large errors due to the small number of cycles included in the computation.

The method used here is that of Heck-Manfroid-Mersch (hereafter HMM: Heck et al. 1985; Gosset et al. 2001), which advantageously replaces the classical Fourier method (see e.g. Bloomfield 1976) in the case of odd sampling. In principle, the highest peak is assumed to be the progenitor one and to correspond to the orbital period. This last point is particularly true for the case of RV orbital motion, since most of the power is gathered in the fundamental frequency; the power of the peaks corresponding to high harmonics is gradually decreasing in the case of the Keplerian motion. This is due to the fact that RV curves always present one simple cycle during the orbital period. On the contrary, this is usually not the case for light curves of

eclipsing binaries and ellipsoidal variables where most of the power is gathered in twice the orbital frequency, corresponding to the average depth of the minima, whereas the power at the orbital frequency is dominated by the difference between the primary and secondary minima and thus is much weaker. In their paper, HMM introduced the θ_F statistic which is a power related statistic (see Eq. (A.11) in Gosset et al. 2001). At a predetermined frequency and under the null-hypothesis of Gaussian white-noise, it has the nature of the ratio of two independent χ^2 statistics and it behaves as a Snedecor F statistic with 2 and $N-3$ degrees of freedom. It is reminiscent of a ratio of two quadratic forms, similarly to the use of nested models (here a pure cosine model versus a constant term). Since the global power is proportional to the variance of the white-noise process, it is mandatory to normalise the power by a function of this variance. This is precisely the purpose of the θ_F statistic.

The periodogram is explored and the frequencies corresponding to the various peaks that are above a certain threshold are further retained. This threshold is defined as the limit above which the highest peak can be considered as significant at a certain probability level. This is further described in the next section. We keep in a list up to 100 candidate frequencies. This is more than the amount of independent ones in the classical Fourier. If too many peaks are listed, only the 100 highest ones are kept. For each of these selected candidate frequencies, a fit is made of an eccentric Keplerian orbital model. At this stage, there is a possibility to add one additional candidate frequency coming for example from CU7 (e.g. on the basis of a photometric curve). In DR3 in particular, this possibility is only used for the subset of solutions of the type `EclipsingSpectro` combined with SB1 solutions (see Sect. 7.6 by Siopis & Sadowski in Pourbaix et al. 2022; Siopis et al., in prep.).

4.3. Significance level of the period

Thus, to search for the period, a Fourier method is used. In power spectra, the peaks in the periodogram are quite numerous. At a fixed frequency and under the null-hypothesis of white-noise with known variance, the height of the peaks is distributed as a decreasing exponential (χ^2 with a degree of freedom of two). A test must be made to decide if a candidate frequency is significant by computing the significance level which is the probability under the null-hypothesis of a white-noise process of unknown variance to observe a peak at least as high at a predetermined frequency: this is the above-mentioned F distribution. Since the algorithm is considering the highest peak over all frequencies, the adopted significance level is the probability under the null-hypothesis of white-noise to observe the largest peak at least as high. The actual statistic is thus

$$\Theta_F = \sup_{\text{all } \nu} \theta_F(\nu). \tag{9}$$

If the significance level is very close to zero, the peak could not be considered as generated by the noise; thus, there is considerable suspicion that the outstanding peak is due to a deterministic periodic signal. In the case of even sampling of the time series, such a significance could be easily computed. However, in the case of irregular sampling and particularly in the case of the *Gaia* sampling, the situation is much more complex. Unfortunately, no rigorous analytical formula actually exists for these precise circumstances. Consequently, to test for the significance of the period, the procedure relies on the use of Monte-Carlo simulations performed in the conditions of the *Gaia* RVS sampling. In principle, this should be performed individually for each object characterised by their own sampling, but this is not affordable for such a large number of sources.

The first dependency of the statistical distribution is on the number N of data points in the time series. The secondary dependency concerns the actual distribution in time of the data points. The dispersion of the distribution decreases with N increasing. Therefore, we first attempted to perform such simulations for each of the possibilities on this number N. Keeping this number constant, it is easy to notice through the simulations that the statistical distributions of the highest peak at fixed N remain quite similar whatever is the actual distribution in time of the data points (i.e. over the various simulations). This is particularly true for the interesting regime where the significance level (SL) is lower than 0.2, but also when N is getting large (see Sect. 4.4). Therefore, various Monte-Carlo simulations were performed and the value of the height of the highest peak associated to various fixed significance levels has been computed. From these simulations, a look-up table is built that is used by the algorithm.

The probabilities (`conf_spectro_period`) given in the data model (DM) are $1 - SL$. Despite the given name, this is thus the probability of rejecting the null-hypothesis of white-noise and is by no means the probability to have a periodicity. Independently, it should be made clear that the frequently encountered notion of significance in the database has nothing to do with the presently defined significance levels. The significance rather corresponds to the value of the parameter over its uncertainty. Such a significance is also used in the spectroscopic channel in the particular case of the semi-amplitude K where it is defined as K/σ_K.

4.4. Build-up of the look-up table

The computation of the empirical SL function in our software can either be performed on-line (on an individual source basis) or read from a look-up table. The on-line computation, being necessarily more accurate but requiring an important amount of computational time (more than 10 minutes per source), was no longer considered by the processing centre after it threw time-out exceptions during pre-operation tests. Therefore, for DR3, we adopted the method of the look-up table which is approximate since it is based on typical average properties of the actual samplings (with N data points) and not on the actual individual sampling.

To build the look-up tables off-line, we gathered the observation dates from final RV curves of RVS sources used as standards, kindly provided by CU6, and we grouped them on the basis of their number of valid RV measurements. We preferred to rely on concrete observation dates rather than on the predicted NSL dates provided by the GOST tool in this exercise. Indeed, there are significant discrepancies between the NSL prediction and the final actual observation dates. The reason is that many spectroscopic observations fail to provide a valid RV measurement at the end of the CU6 spectroscopic processing. This implies a sparser sampling.

For each source, we computed 10 000 HMM periodograms of Gaussian white-noise realisations on the very same date sampling as that of the considered source, by adopting the same frequency step and limits as in the operation pipeline. For each periodogram, we retained the largest value of the θ_F, the so-called Θ_F statistic. Then we merged all the values and ranked them by increasing order and build up an empirical cumulative distribution function (ECDF) with the 10 000 values of the Θ_F statistic. At N fixed, each realisation of the sampling provides

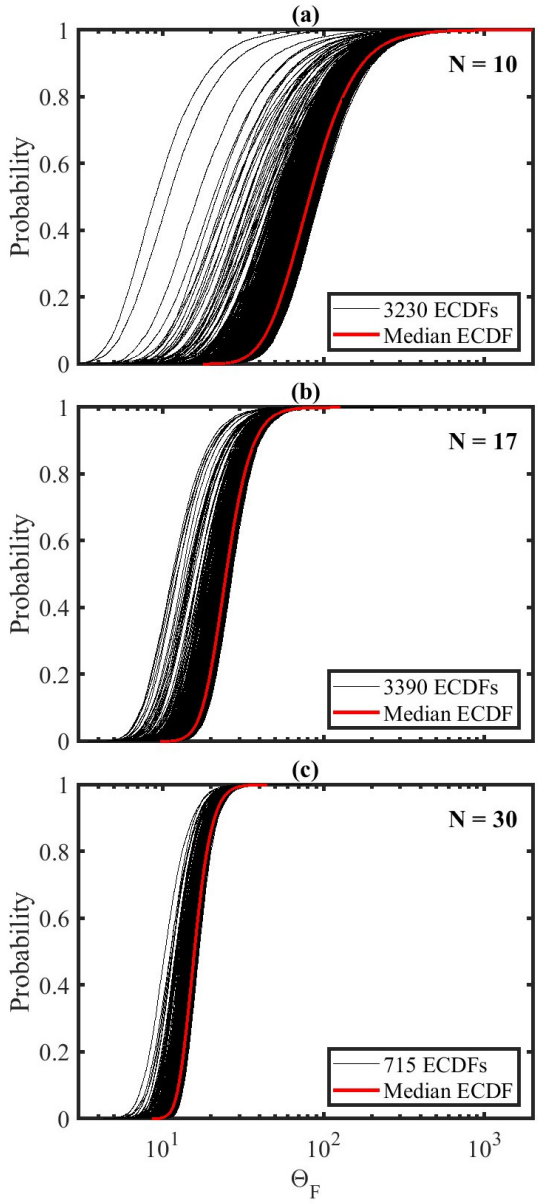

Fig. 3. Examples of simulated ECDFs (in black) for different valid numbers of transits N and their median (in red). The ordinate gives the probability to reject the null-hypothesis of white-noise as a function of the height of the peak (Θ_F statistic of Eq. (9)). Panels a, b and c correspond to N = 10, 17 and 30 respectively. The number of simulated ECDFs in each panel mirrors the proportion of valid transits in the actual data. The sampling time statistically tends to be similar for large N values, leading to a decrease in ECDF dispersion when N increases.

a particular ECDF function. The choice of 10 000 in the framework of the ECDF allows us to assign false alarm probabilities with a precision better than 0.5% at 95% confidence.

The look-up table contains the set of median ECDFs for every given amount N of valid transits. The real distribution of the number of observations per source is reflected in the number of ECDFs accounted for in each median with this scheme (i.e. sources with 17 valid observations, which are the most common in DR3, have a median ECDF covering the greatest number of individual ECDFs).

Figure 3 shows the simulated ECDFs for N = 10, 17 and 30 and their corresponding median. As time sampling approaches the similarity of the samplings for large N values, the ECDF dispersion decreases and their median approaches the regular time sampling CDF. The values in ECDFs also decrease when N increases. Indeed, fitting a cosine function to a large sample of pure noise will not significantly improve its variance, unlike in small samples where perfect fits may occur, especially when high frequencies are explored. In addition, the signal-to-noise contrast increases with N.

When calculating the periodogram of an input time series with the length N, its corresponding adopted ECDF (on the basis of the N value) is read from the look-up table. The table has a comprehensive coverage of the number of transits, without gaps from N = 10 to 72, but there are some entries for large N values that are not present. In this case, the algorithm opts for the closest lowest N input.

5. Calculating the orbital solutions

Our main task is to derive simple orbital motion based on Keplerian models. The general model (see Eq. (1)) is highly non-linear and usually necessitates an a priori approximate knowledge of the true values of the parameters to fit. This is not possible in the framework of a pipeline treating so many objects. To define the proper method to use, we investigated numerous possibilities; they are evoked in Appendix A. Despite the many trials performed, it turned out that approaches based on classical Levenberg-Marquardt iterations were much more robust and precise, if appropriate precautions are taken. As explained in Appendix A, the alternative idea to perform simulations were rejected for excessive processing time reasons.

The Keplerian fitting, even at fixed period, is time-consuming if it is to be performed over a vast set of periods: as for example in a genuine Keplerian periodogram (see e.g. Cumming 2004; Zechmeister & Kürster 2009, their Sects. 2.3 and 5, respectively).

Therefore, we decided to strew subtleties over our procedure in order to render it much more robust. The first step is to search for periodicities using a periodogram (HMM method) as described in Sect. 4. This corresponds to a fit of sine/cosine, which is a linear process (at fixed frequency). This is particularly convenient: even if the RV curve is highly eccentric, and thus the power of the signal is spread over several harmonics, the power related to the period of the orbit is still dominating the noiseless periodogram. If the number of cycles is sufficient, the periodogram will always present increased power at the correct period. However, the detection of a periodicity does not imply that the motion is Keplerian. Therefore, after the establishment of a list of candidate periods, the Keplerian model needs to be fit for each of them to deduce the orbital solution and corresponding parameters that are a better fit to the data. We first performed a Keplerian fit that is linearised at the maximum (but not entirely) and represents a preliminary fit, not definitive but robust. The actual least-square fit was performed with the true anomaly as the independent variable.

The Fourier periodograms are known to be inefficient at addressing the case of eccentric RV curves (for $e \gtrsim 0.5$, see e.g. Endl et al. 2002; Cumming 2004). We tend to think that this is due to the fact that for eccentric orbits by opposition to the circular case, not all the data points contain the same amount of information. In fact, for very eccentric orbits, the data points related to observations around the maximum separation are more determining, in particular for the semi-amplitude. If such data points are missing, have not been acquired, this will have major impact at least on the determination of the K value. Consequently, this will have a huge importance on the detectability of

the signal: more abruptly, we cannot detect something that is not observed. Actually, this is not a problem of ability of the Fourier periodogram to address eccentric orbits, this is a problem of fair sampling of the RV curve. The correct way to interpret things is that if a period is detected, we are authorised to have some confidence in its existence. Therefore, the nature of the Fourier periodogram improves the reliability and robustness of the algorithm at the expense of a difficulty to have a proper estimation of the completeness. This is a choice we made.

We finally decided to adopt the method fully described here below, which is based on fully analytical and classical numerical developments. After several tests, the adopted method appeared to be faster, more stable, trustworthy and robust.

6. Main NSS-SB1 processing chain

6.1. Global structure of the chain

Five Java packages operate in sequential order in our pipeline. These are *InputChecker*, *PeriodFinder*, *SB1Refiner*, *SB1CRefiner* and *TrendSolution*. The package *InputChecker* reads the input time series in the NSS data model and copies them into arrays of primitives that are shared with the entire pipeline. This reduces the number of modifications required after every data model update. It also performs sanity checks, data cleaning and some other modifications. Finally, it passes the rearranged data to the four remaining processing packages.

The second package (*PeriodFinder*) is working in two steps. Its first step is to establish a list of candidate periods that are significant. Its second task is to test each of the significant candidate periods by fitting the corresponding Keplerian model to the data, and to select the best period on the basis of the comparison of the various models. The best of these models represents a first approximation of the orbit. The third package (*SB1Refiner*) has the charge to start from the selected approximate solution and to refine it. As part of this computation, it derives the variance-covariance matrix of the fit. The fourth package (*SB1CRefiner*) addresses the case of circular orbits when the eccentric solution generates a badly configured variance-covariance matrix, and tries to adopt the best compromise between a circular solution and a singular eccentric one. After these computations, a decision is to be taken to adopt the most adapted periodic orbital solution. The fifth package (*TrendSolution*) deals with the RV time series exhibiting a trend as a function of time (see Sect. 3.3). This means time series that can be fitted by polynomials. These are acceleration models. After the execution of all these packages, a choice is made to select the solution that presents the best fit: SB1, SB1C, TrendSB1, or StochasticSB1.

The five packages are fully described in the respective five Sects. 6.3.1, 6.3.2, 6.3.3, 6.3.4, and 6.3.5. The flow chart of the full NSS-SB1 processing is described in Fig. 4, along with the relations between the five packages mentioned above. It is interesting to notice that after the upper violet lozenge that contains a test on $N_{\rm good}$ (defined in Sect. 6.3.1), there is a bifurcation towards two sub-chains that receive the data. One sub-chain is treating the objects exhibiting a periodicity that could conduct to an orbital solution. The other sub-chain treats the cases where no periodicity is involved (only trends). From the purely logical point of view, they are working in parallel although from the practical point of view, they are executed in sequence. During the processing of the full NSS-SB1 chain, some flags are raised to retain more information about the path followed by the time series, as well as to point out a few characteristics that could be useful for further treatment. In Sects. 6.3.1, 6.3.2, 6.3.3, 6.3.4, and 6.3.5, the flags are identified by their number and a full description of their nature is given in Sect. 6.5.

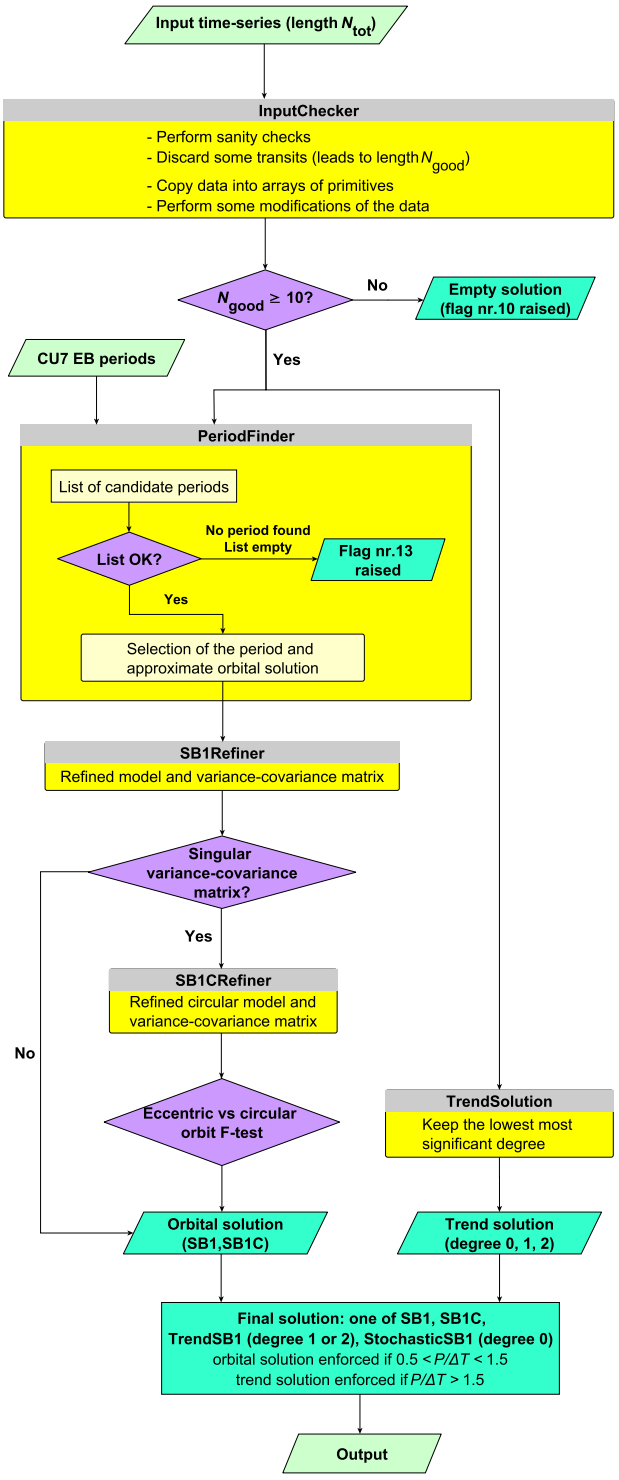

Fig. 4. Flowchart of the single-line channel, as described in Sect. 6.1. The details are fully explicited in the description of the individual processing steps (see Sect. 6.3). ΔT is the actual time span of the observations.

6.2. Generalities about the fits

After the *InputChecker* intervention, the four other NSS-SB1 processing packages are performing a set of linear or

non-linear fits, involving linear algebra operations on moderate rank matrices (rank of matrices never exceeds 6). A home-made linear algebra toolbox, based on the Eigenvalue and Cholesky decompositions, has been developed for that purpose.

Although the former is known to be time-consuming compared to QR or Cholesky algorithms, it gives the best estimates of the condition number of involved algebraic systems, and thus ensures a better robustness. This is why it is applied all over the pipeline except for trend analysis, for which we adopted the Cholesky decomposition instead.

Moreover, thanks to the eigenvectors, singular algebraic systems can still be solved by the Moore-Penrose pseudo-inverse technique (Press et al. 2007, Sect. 2.6.2). This is extremely useful because it solves the discontinuity problem between circular and eccentric models. Indeed, when eccentricity approaches zero, ω and T_0 parameters become progressively redundant (Sterne 1941), and thus the eccentric model becomes overstated, too powerful and thus prone to overfitting. The boundary between the two regimes, circular and eccentric, cannot be set by a fixed value of the eccentricity since it is governed by the signal-to-noise ratio of the input time series, and thus depends on many parameters such as the sampling, the amplitude of the signal and the involved noise (Lucy & Sweeney 1971).

Eigenvalue decompositions of matrices with rank smaller than 5 are performed analytically to improve robustness and reduce the processing time. Cardan-Tartaglia and Ferrari methods are applied to compute eigenvalues of rank 3 and 4 matrices, respectively.

The Cholesky decomposition is very suitable for the trend analysis in which we successively increment the polynomial degree. Since the k^{th} polynomial order is nested into the $(k+1)^{\text{th}}$ one, we just need to update the k^{th} order curvature matrix decomposition to the new dimension instead of completely recomputing it.

Our home-made linear algebra toolbox defines the same singularity threshold for both decompositions. It is chosen equal to $10 \times 2.22 \times 10^{-16} \times$ the square of the input matrix size. The value 2.22×10^{-16} is the lowest non-zero double precision number in current processors. The threshold corresponds respectively either to the ratio between the extreme eigenvalues (in the eponymous decomposition method), or to the square of the ratio of the extreme diagonal elements of the Cholesky matrix.

6.3. Processing steps

6.3.1. Ingestion of the data

All stars that present RV variations according to the following criteria `rv_chisq_pvalue`≤0.01 and `rv_renormalised_gof`>4 are processed by the pipeline starting with the *InputChecker* described here. Other details about these input data were already provided in Sect. 2.

In the ingestion process, the NSS input data are requested from the Main Data Base (MDB) and copied into `StarObject` Java objects during data ingestion. A `StarObject` contains the single star astrometric solution (the reference epoch, the mean position, the proper motions, the parallax, `ruwe`, and `astrometric_excess_noise`), the atmospheric parameters as measured by CU8 (Coordination Unit Astrophysical Parameters, Andrae et al. 2023; Recio-Blanco et al. 2023), as well as the light curve classification and the orbital period if the corresponding source is suspected to be an eclipsing binary by CU7 (Coordination Unit Variability Analysis, Mowlavi et al. 2023). In addition to these mean data, a `StarObject` contains an array of N_{tot} `Transit` Java objects, where each `Transit` contains the relevant epoch data such as the source position, the integrated flux in the G, G_{BP} and G_{RP} bands, the RV (along with its uncertainty) as well as some related quality flags. Observation dates at the centre of the astrometric field are given in BJD. With the satellite spin, an observation at the centre of the RVS CCDs is performed roughly 75 s after the observation at the centre of the astrometric field: this delay has been neglected. The observation date, the RV and its uncertainty are extracted from each `Transit` and copied into arrays of primitives for limiting code modifications required after every data model update. Those that are not finite numbers are filtered out. Observation dates and error bars must be positive numbers, while RVs must fall within the range of $[-1000, +1000]\,\text{km s}^{-1}$. The length of the time series N_{tot} is expected to be larger than or equal to 10 and if it is not the case a flag is raised (flag nr.8; see Sect. 6.5) and the star is rejected from further processing.

After these sanity checks, we start by cleaning up non-valuable entries. A `Transit` is discarded when its RV is detected by CU6/MTA as an outlier, either on its own or together with the rotational velocity or the integrated flux in the G_{RVS} band or both. In addition, we also discard the `Transit` if it exceptionally corresponds to a composite spectrum (see Damerdji et al., in prep.). The resulting time series after cleaning contains N_{good} data points distributed over a genuine individual time span ΔT (the difference between the two dates of observations at the extreme of the interval).

Finally, we perform a few modifications. We correct uncertainties for sources with large `ruwe`, as described in Sect 2. The reference epoch J2016.0 (BJD 2 457 389.0) of the astrometric solutions, is subtracted from the observation dates for reducing the correlations between the derived solution parameters. A cleaned time series of N_{good} RVs, still containing the required number of 10 observations, is shared with the whole software for the rest of processing, whilst, in the opposite case, it is rejected from the pipeline, and the corresponding flag is raised (flag nr.10; see Sect. 6.5 for flag descriptions). The rejection from the pipeline is done by outputting an empty solution containing the flag and no parameter values. Such an empty solution is for internal use and statistical analysis. It is not further considered for the final output of the chain.

6.3.2. Inventory of the periods and approximate model

The package *PeriodFinder* is first establishing a list of candidate frequencies. The time series of length N_{good} is analysed thanks to an HMM periodogram, as explained in Sect. 4.2. The package reads the ECDF corresponding to N_{good} (see Sect. 4.4) and establishes a threshold above which the peak is considered significant against the null-hypothesis of white-noise. All the peaks in the periodogram that are above the threshold value have their position in frequency entering the list. This is clearly an approximation since the test is constructed on the rigorous concept that we are looking at the highest peak only (see Sect. 4.3). No fully sound counter-indication to this assumption has been identified yet. Up to 100 frequencies/periods are accepted in the list. No additional smaller peak is considered above this amount. If the object under analysis is an Eclipsing Binary, we also add to the list the value of the period issued from the photometry. We test the significance of any available external period (i.e. in DR3 from CU7) before adding it to the set of candidate periods. For DR3 the periods entering the list were selected on the basis of an adopted threshold probability value `conf_spectro_period` larger than 0.95 (*SL* smaller than 0.05). If the list of significant

candidate periods is empty, no full orbital solution is possible. The relevant flag is raised (flag nr.13).

Since the aim of the pipeline is to compute orbital solutions, it is worth trying to limit the possibilities to Keplerian motion. For each of the frequencies ν_k in the list, a Keplerian model is fitted. The pipeline is not performing the fit in phase (defined by the candidate frequency under test) but in true anomaly. In this domain, the RV curve to fit is a cosine. This of course necessitates the knowledge of T_0 and e. Since the eccentricity is a strongly non-linear parameter, the pipeline proceeds as follows. Starting from a two-parameter grid of $e = (0, 0.5)$ and $T_0 = (0, 0.25, 0.5, 0.75)/\nu_k$, we minimise, using a Levenberg-Marquardt procedure, the quadratic sum of the differences between the observed RVs and the model ones (χ^2). The vector of model RVs (coming from an expansion of Eq. (1)) is given by

$$RV(t_i) = \begin{bmatrix} 1, & \cos v(t_i) + e, & \sin v(t_i) \end{bmatrix} \begin{bmatrix} \gamma \\ \alpha \\ \beta \end{bmatrix}, \tag{10}$$

where $\alpha = K \cos\omega$ and $\beta = K \sin\omega$. From these α and β, the Campbell parameters K and ω are easily extracted. In Eq. (10), γ, α, and β appear in a linear way. They can be optimised through an ordinary linear least-square regression and this for every set of trial $v(t_i)$. This system of three linear and three non-linear parameters, can be efficiently solved with the separable variable technique (Golub & Pereyra 2003). Indeed, we optimise three non-linear parameters only: e, T_0 and ν_k ($\equiv 1/P_k$). Consequently, the dimensional space of variables is reduced, leading to more robust and rapid convergences. The pseudo-inverse implemented in the linear algebra toolbox smooths singularities near zero eccentricities. The rank of the involved linear systems is smaller than or equal to three, allowing them to be solved using more precise analytical computations. We give in Appendix C a summary of the involved computations. The minimisation is limited to 1000 iterations for each of the 8 considered starting points, and this for every candidate frequency ν_k. The solution is taken into account for further comparison, even when the maximum number of iterations is reached. The relative convergence of the 8 different descending tracks are analysed to estimate the robustness of the convergence (the arrival point should be the same in a perfect world). Ultimately, the best solution is selected for the considered candidate frequency. This algorithm is quite robust and provides adequate fits for orbital curves up to eccentricities of 0.995, well beyond what is necessary for the majority of the objects. The procedure is executed for every candidate frequency and gives a χ^2 of the fit that is translated into an F_2 statistic. The smallest value identifies the best fit and the best set of parameters, including the period. The algorithm is efficient at finding the best minimum, but the fit is performed in two steps, which is clearly an approximation concerning the cross-covariance terms between families of parameters (linear versus non-linear). Therefore, we need to refine the approximate model.

The use of a fit in the true anomaly space ensures that no analytical solution could present a bad behaviour in the phase space. Zones in the phase diagram where no data are present, the phase gaps, could exist in some pathological cases. This problem is tackled in Sect. 7.1.

6.3.3. Final fit of the eccentric orbital solution

After having identified the best approximate orbital solution, the global solution refiner comes into play (*SB1Refiner*). It is a Levenberg-Marquardt minimisation based on first order derivatives as given by Schlesinger (1910). It returns the eccentric model as well as its variance-covariance matrix. The algorithm thus performs, at the best minimum (the approximate model), a classical least-square fit over all the parameters with all the non-diagonal terms being taken into account. This has the advantage to provide sound (but classical) estimations of the uncertainties on the parameters. The variance-covariance matrix may be singular due to an over-determination of the model. In that case, a flag (flag nr.15) is written before activating the following package. If the matrix is not singular, the SB1 solution is forwarded to the final decision module. If the iterative process does not converge after 1000 iterations, a flag is raised (flag nr.14). If the numerical variance-covariance matrix is not semi-definite positive, a flag is raised (flag nr.18).

The eccentricity is a parameter that suffers from a bias related to the noise propagation. The RV noise could enforce a fitted value for e that is much larger than the true one. This effect is sometimes compensated by enforcing $e = 0$ for the solution when the eccentricity is not significant. This is common practice in the literature to apply a test and to decide to go circular (i.e. to decrease the number of free parameters) for low-amplitudes and noisy cases. This has the pitfall to enforce the circular solution for objects whose eccentric solution is only marginally good (non-significant eccentricity). This produces a deficit of non-zero but small eccentricity solutions. The present algorithm does not do that. It never decides that an orbit could be forced circular. It is allowed to possibly (but not necessarily) go circular when the eccentric model cannot lead to a well-behaved non-singular variance-covariance matrix.

6.3.4. Fit of the circular orbital solution

The module described here is the *SB1CRefiner*. The best approached period obtained in the *PeriodFinder* package is refined and fitted along with the linear parameters γ, K and T_0 to fit a circular model to data (see the equations in Sect. 3.2). When a circular orbit is selected, T_0 corresponds to the maximum of velocity by definition. A flag is written when the derived variance-covariance matrix of the circular solution is singular; this is flag nr.16. If the iterative process does not converge after 1000 iterations, a flag is raised (flag nr.21). If the numerical variance-covariance matrix is not semi-definite positive, a flag is raised (flag nr.19).

In some sense, the less complex circular model is nested into the eccentric model through the eccentricity. Indeed, if we put $e = 0$ into the eccentric model, we obtain the equations of the circular model (except for the trivial change of the definition of T_0). It is thus tempting to opt for a classical statistical test for nested models to choose between the two. This is actually what is done. However, this is not perfectly rigorous since, although the circular model is linear, the eccentric one is not. Therefore, the number of degrees of freedom of the χ^2 is not rigorously defined. In any case, as a first approach, we decided to use the classical nested model statistic F (also described in Lucy & Sweeney 1971, i.e. Lucy's test). We thus test here the eccentric orbit model (as computed in *SB1Refiner* but corresponding to a possibly unreliable solution having necessarily at least one of the three flags nr.14, nr.15 or nr.18 raised) against the circular orbit model taken as null-hypothesis. The eccentric model is thus retained when the Snedecor-F test significance level is lower than (or equal to) 0.05. This decision is taken regardless of the quality of the variance-covariance matrix of the circular model. The *SB1CRefiner* thus

produces SB1 and SB1C type candidate solutions. When Lucy's test has been applied, flag nr.22 is raised.

6.3.5. Fit of trends

Not all the observed RV curves correspond to a true periodicity nor to a pseudo-one. One widespread alternative possibility is the existence of a mere trend: the data as a function of time can be represented by a polynomial function of t (see Sect. 3.3). The last computing package *TrendSolution* applies polynomials to data to search for acceleration models (trends) rather than periodic orbits. We start by applying a zero degree polynomial (corresponding to a noisy constancy), then we successively increase the order at each step. The lower degree polynomial is thus nested in the new one. The choice between the lower and the higher degree polynomial is done through a classical Snedecor-F test on the ratio of the improvement of the χ^2 to the value of the χ^2 for the higher level model, since the polynomials given above are indeed nested models. A Snedecor-F test for nested models is thus applied. The higher degree polynomial is accepted by the Snedecor-F test for a confidence larger than 0.99865 (3-σ probability). We stop iterating at the 4$^{\text{th}}$ order, or when the higher-degree model exhibits more than one extremum over the time baseline. This last criterion has been adopted to avoid any possible overlaps with the orbital models with periods between ΔT and $1.5\,\Delta T$ (see below). Most of the objects entering the trend category are probably long-term periodic curves but with a timescale much longer than the present time span of the mission (ΔT). The code is built to accept polynomials of degrees 1, 2, 3, and 4 but in DR3, only degrees 1 and 2 were encountered (at the selected probability threshold). A flag is written when the derived variance-covariance matrix of the trend solution is singular; this is flag nr.17. If the numerical variance-covariance matrix is not semi-definite positive, a flag is raised (flag nr.20). A regrettable mistake occurred in the code and, for the last operational run, the flags raised were not the correct nr.17 and nr.20 but instead flags nr.15 and nr.18, respectively. This problem has been corrected for DR4.

The coefficients of the polynomial are included in the database but the work is performed in an Hipparcos-like formalism (see Sect. 2.3.3, ESA, 1997). The coefficients saved in the database are related to the derivatives of the velocity curve and are expressed with respect to the DR3 reference epoch (J2016.0, equivalent to BJD 2 457 389.0).

We thus decided to express the trends as a function of the various derivatives. The simple mathematical model is

$$V(t) = V_0 + \frac{\partial V}{\partial t}\, t + \frac{1}{2}\,\frac{\partial^2 V}{\partial t^2} t^2, \tag{11}$$

where t is the time (in days) relative to the reference epoch. We thus fit to the time series the polynomial,

$$a + bt + ct^2 \tag{12}$$

with the c value potentially fixed to zero. In the case of the first degree polynomial ($c = 0$), we adopt the following rigorous formulae,

$$V_0 = a, \qquad \frac{\partial V}{\partial t} = b. \tag{13}$$

The a and b parameters are not correlated if $\sum t_i = 0$. This is approximated by subtracting to the BJD dates the reference epoch (J2016.0) as a proxy of the middle of the current mission. Practically, not all actual time series are centred on this value since some transits might be missing (e.g. more numerous before the mid-epoch or alternatively after it). This is in any case still an approximation since it is implicitly assumed that the distribution of the sample times are homogeneously apportioned. This is certainly not true, particularly if we include the EPSL data.

In the case of the second degree polynomial (c fitted), we note that the fitted polynomial can be expressed as

$$a' + bt + c\left(t^2 - \frac{(\Delta T_{\rm m})^2}{12}\right), \tag{14}$$

where $\Delta T_{\rm m}$ is the total duration of the mission treated in the framework of DR3. We thus adopt the approximate formulae,

$$V_0 = a' - c\,\frac{(\Delta T_{\rm m})^2}{12} \qquad \frac{\partial V}{\partial t} = b \qquad \frac{\partial^2 V}{\partial t^2} = 2c. \tag{15}$$

The term shifting t^2 by $\frac{(\Delta T_{\rm m})^2}{12}$ has been introduced such that the contribution by the second derivative is tending to be neutral over the interval $\Delta T_{\rm m}$. This of course induces a shift on V_0. The first derivative is considered as averaged over the full interval. This property is here also true for an homogeneous distribution of the observing times. A strong disparity of the observing times could have introduced some biases.

Initially, the threshold was chosen at a probability level of 0.95 (following the Snedecor-F test of nested models) and polynomials of first, second, and third order (with only one extremum) were found. Following validation, we decided to be more conservative and we instead adopted a new threshold (0.99865); this induced a downwards cascade among the different degree categories and restricted the palette of solutions to first and second order polynomials. But the decision occurred too late and no more operational run took place. Therefore, we kept formally adopting the threshold of 0.99865 and decided to redistribute the various solutions and to blacklist the objects that, in case of a rerun, would have been discarded as trend solutions. This certainly induced the loss of a small number of objects in other categories.

The acceleration identification as performed in DR3 can be affected by false detection, revealing the subtle effect of an unfortunate sampling on the polynomial fit. In the upcoming data releases, changes to the algorithm will be made to ensure that trend models are first detected in the lowest frequency region of the periodogram, as introduced in Boulkaboul et al. (2022).

6.3.6. Selection of the best solution

Therefore, the *SB1Refiner* and the *SB1CRefiner* packages provide an SB1 or SB1C candidate solution. They could in some cases provide no solution if no period is found. The *TrendSolution* package delivers in parallel polynomial fits and thus candidate trend solutions (of degree 0, 1 and 2). It is now necessary to decide which model is performing the best interpretation of the data. The adopted orbital, circular or eccentric, candidate model is ultimately compared to the less complex best trend candidate (i.e. the null hypothesis is that the model is a trend). Acceptance of the orbital model is granted when the Snedecor-F test confidence is higher than 0.95. If this is the case, the final delivered solution is the selected SB1 or SB1C. If this is not the case, the trend solution is adopted. If the degree of the fit is 1 or 2, the object is entering the delivered output as a TrendSB1 solution. If the degree is zero, we have to deal with a constant RV with excess noise compared to the expected uncertainty. In

this last case, the solution is output as a StochasticSB1. Thus, when no significant periodic variability nor trend are possible models, the algorithm associates the star to a stochastic solution. These are objects that have very little chance of being associated to a mere orbital or trend motion. As explained elsewhere, some badly sampled orbits and high degree multiple systems could fall into this category, perhaps momentarily. Some of them will certainly benefit from an increase in the number of transits. These StochasticSB1 solutions are not provided in the DR3 catalogue. The correct choice between an orbital solution and the trend one is particularly difficult because the independent variable is time for the trends and let say phase for the well covered orbits. The time entering in a non-linear way in the Keplerian equations and the first selection of the period being based on a HMM periodogram, the statistical Snedecor-F test is not fully correct.

Actually, the procedure described above has been applied when orbital solutions presented a period lower than $0.5\,\Delta T$; it concerns the regime where the orbital motion is expected, when selected, to be convincingly robust. For longer periods, the folding with the period does not permit us to test the perfect overlap homogeneously at all phases. We found it preferable to adopt an alternative procedure. If the orbital period is comprised between $0.5\,\Delta T$ and $1.5\,\Delta T$, we decided to enforce the choice of the orbital solution. For periods larger than $1.5\,\Delta T$, i.e. outside the zone of possible convincing orbital motion, we decided to enforce the choice of the TrendSB1 solution. This last point explains why we decided to explore the periodogram in the range between $1.5\,\Delta T$ and $2.0\,\Delta T$.

6.4. Quality of the solution

For each object processed, along with the various resulting model-parameters, the code also computes some statistics and various entities that might help in estimating the quality of the solution. The first statistic is the F_2 one given by

$$F_2 = \sqrt{\frac{9n}{2}}\left[\left(\frac{\chi^2}{n}\right)^{1/3} + \frac{2}{9n} - 1\right], \tag{16}$$

where χ^2 is the classical statistic of the least-square fit and where n is the related number of degrees of freedom. This statistic was first introduced by Wilson & Hilferty (1931). Since the eccentric fit is non-linear in nature, we performed various simulations to test the estimation of the number of degrees of freedom to use. Quite surprisingly, it turned out that, for the Keplerian motion, the number of degrees of freedom n is very close to the number of data points minus the number of parameters (similarly to the linear case).

Besides the F_2, the code also computes the $\delta\phi$ which is the largest gap in the phase diagram where no data points are present. It also computes the solution efficiency whose definition can be found in Eichhorn & Xu (1990), and is corresponding to a measure of the correlation between the computed parameters. A total absence of correlation, and thus a diagonal variance-covariance matrix, will give an efficiency of 1. Values of the efficiency tending to 0 indicate strong correlation. Practically, an efficiency identically equal to 0 is not possible since this is associated to a perfectly singular matrix to invert.

In addition, we also compute the solution significance defined for spectroscopic data as K/σ_K. Finally, we should add the significance level of the periodicity as described in Sect. 4.3 and expressed as a confidence $(1 - SL)$.

6.5. Flagging

Beyond the determination of the various parameters, the code also produces various flags that allow the user to have a good understanding of the results and to select some solutions that could be better or more secure than others. The flags generated by our NSS-SB1 processing are incorporated into a more global set of flags issued from the global NSS processing. We will restrict the discussion here to the set of flags directly related to the SB1 processing chain, with a few information linked to flags issued from the SB2 processing that is fully described in Damerdji et al. (in prep.). All the NSS flags are included in a unique long variable, whose every bit acts as an individual logical flag. All the flags involved in the NSS spectroscopic processing are enumerated in Appendix I. As an example, the *SB1Refiner* processes the time series of an object and produces a singular matrix. The flag nr.15 is thus raised, giving to the global flag an additive value of 32 768. The object is then processed by the *SB1CRefiner* and, let say, it concluded to a singular matrix associated to the circular orbit. The flag nr.16 is thus raised (associated to a value of 65 536). Consequently, the object is associated to both flags nr.15 and nr.16 raised, giving together a value of $32\,768 + 65\,536 = 98\,304$. The object could appear in the catalogue as a `TrendSB1` with the two above-mentioned flags raised. Although the majority of the flags are important for the decision tree of the algorithm, some of them are purely indicative. This is the case of the flag nr.24 that draws the attention to the fact that the attributed eccentric orbit appears, in the eccentricity versus log$P$ diagram, outside the zone where the circularisation of the orbit could have partly taken place. The borderline used is inspired from the one defined by Halbwachs et al. (2005) but modified by us[3]. Flag nr.25 just indicates that the selected adopted period was compatible with one period suggested by the photometry. Flag nr.22 informs that the Lucy test has been applied. Flags nr.11 and 12 concern SB2 objects that turned out to be impossible to treat as SB2 either because the amount of RVs was too small for the secondary (but not for the primary) or because the period found by the SB1 chain was not coherent with the one of the SB2 chain (see Damerdji et al., in prep.). Flag nr.26 is raised by the SB2 processing to indicate that the pair of RVs of the SB2 measurements were not anticorrelated (more details are given in Damerdji et al., in prep.).

7. Post-processing

The post-processing described here further cleans and distributes the results of the NSS-SB1 main processing chain. The detailed flowchart is provided in Fig. 5 and additional explanations are given in this section.

7.1. Internal filtering

From summer 2020 to spring 2021, the NSS pipeline processed the *Gaia* data in the framework of several validation runs. At least seven of them were concerning the software linked to the present paper. This processing helped in finding remaining failures of the algorithms and in correcting them. This also helped in fine-tuning the decision tree of the pipeline by adjusting some

[3] Using the formalism of Halbwachs et al. (2005), we estimated a value of $P_{\rm cut-off}$ on the basis of a fairly crude analysis of the full SB9 catalogue (Pourbaix et al. 2004). This estimate has been deduced without proper treatment of the dispersion of e due to noise and is considered as a conservative border. The purpose is to flag extremely pathological cases to trigger the cautiousness of the general user.

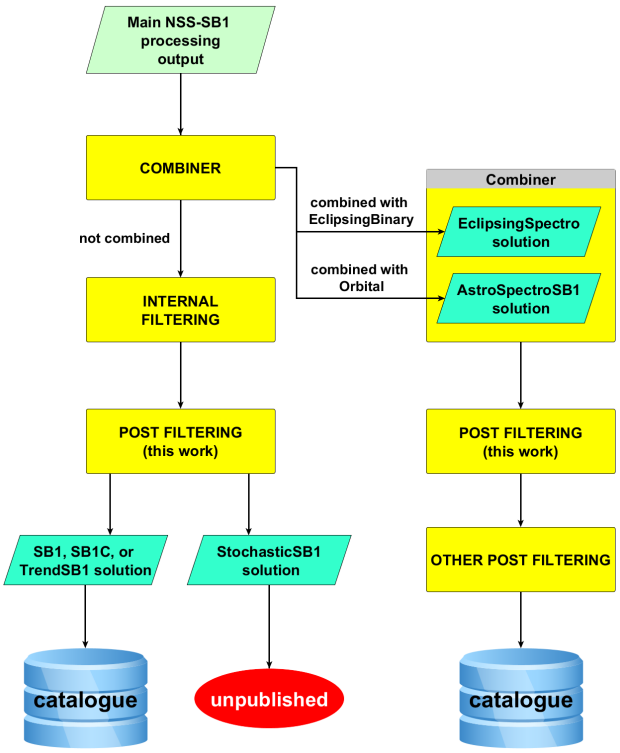

Fig. 5. Flowchart of the post-processing that is applied after the output of the flowchart of Fig. 4. The channel on the left side solely concerns the present work. The channel on the right side also concerns the combiner. The post-filtering (this work) concerns the post-filtering announced in Sect. 7.3 and detailed in Sect. 10.1. The other post-filtering only concerns the astrometric and/or the photometric channels.

statistical thresholds. The operational runs took place in summer 2021.

In order to maintain the quality level of the various solutions, we defined an internal filtering of these solutions that would allow us to retain the very best and most secure ones. The aim is also to spot out and reject solutions that are problematic. Most of the filtering is done via the definition of thresholds. Very often, the selection is made in a statistical way and by using probability thresholds based on the distribution of some statistical variables. As always, the selection near the borders is never perfect. We detail here the internal filtering for SB1 solutions.

Since the individual (per transit) RVs are searched for in the range between -1000 km s^{-1} and $+1000$ km s^{-1}, we found wise to test that all solutions have $|\gamma| \leq 1000$ km s^{-1}. This is the case but this sane test is not highly restrictive.

On the opposite, one of these important statistics is the F_2. The F_2 statistic represents a classical goodness of fit criterion that is directly related to the objective function of the fit and to the classical χ^2 statistic. The F_2 statistic is expected to have under the null-hypothesis a Gaussian distribution. Typically, above $F_2 = 5$, the solution is not satisfactory and presents extra-dispersion of the residuals, be it due to the data or to the model. In order to be conservative, we decided to put the threshold such that $F_2 \leq 3$.

In the course of the NSS-SB1 processing, various flags are raised, tracking the details of the decision tree and of the problems encountered. As explained in Sect. 6.5, some combinations of flags imply that the solution is not acceptable or is presumably inaccurate or too hazardous. The corresponding selection was of course considered. Since the internal filtering is supposed to protect the results from spurious solutions and that, in the present purely spectroscopic SB subcatalogue, no confirmation comes from other channels, we decided to be highly cautious for this step. For the SB1 solutions, we discarded objects having at least one of the three flags nr.14, nr.15, nr.18 raised. As a consequence, this suppresses all the SB1 solutions issued by the *SB1CRefiner*.

We further found that a conservative attitude is to restrict the palette of the solutions to those yielding a $K \leq 250$ km s^{-1} which eliminates a few cases of surprisingly large semi-amplitudes that must be further validated (a work beyond the scope of the present paper).

We also used a threshold on the solution efficiency. Any solution efficiency below 0.1 is rejected because the corresponding fit is considered as spurious (Eichhorn & Xu 1990). The solution efficiency is part of the catalogue, allowing the user to adopt a more restrictive attitude.

When the supposedly periodic RV data are folded in phase, it happens that a zone in the phase diagram is not populated by any data points. If this region is too wide, the corresponding solution could be pathological with respect to the sampling. Therefore, we choose to discard the solution with a too large $\delta\phi$ where this value represents the largest interval in phase with no data points. We reject solutions leading to $\delta\phi > 0.3$. The value of the gap for remaining solutions is not propagated to the SB subcatalogue but will be accessible in DR4 with the release of the epoch RVs. The threshold value itself has been optimised on the basis of the various validation runs (see Sect. 9 and in particular Sect. 9.2). Due to a small bug in our software, a small part of the objects presenting $\delta\phi > 0.3$ have not been properly rejected in DR3.

The mere existence of the periodicity in the data is a strict test of the validity of the orbital solution. It is reasonable to consider values for the probability `conf_spectro_period` from 0.95 to 0.995 or even stricter. During the validation, it appeared a strong population with short-period orbital solutions sometimes associated to relatively large eccentricities. This is somewhat in contradiction with the expected circularisation phenomenon (or these are not properly speaking orbital solutions). Therefore, we decided again to be conservative and to put a threshold that is period dependent. We hope to improve this procedure for DR4. For periods above 10 days, we selected a lower threshold of 0.95 on `conf_spectro_period` (significance level < 0.05). For periods less than one day, we adopted 0.995 as the lower limit. In-between, the lower threshold on `conf_spectro_period` is given by $0.995–0.045 \log P$. The solutions associated to a value less than 0.95 are filtered out. The solutions associated to a value above 0.995 are kept. Solutions in-between are kept but a flag (nr.13, i.e. 8192; see Appendix I) is raised if they are below the defined threshold.

Another natural statistic is the spectroscopic significance defined as K/σ_K. This ratio should not be confused with the probabilistic quantity canonically named significance level. A high significance is necessary to adopt a solution. However, we should not forget that σ_K is itself an estimator and that statistical fluctuations could give a value of σ_K that is exceedingly small, boosting the significance to very large values. The opposite case is also true. We decided to keep the significance as a possible sorting criterion for the future data applications. Therefore, we only rejected all the solutions with a significance lower than 5. This threshold has been fixed for the internal filtering, but it could be advisable to the general user to select a value of the order of 20 to 50 depending on the details of the scientific application. An example of this way of fine-tuning the selection is illustrated in Gaia Collaboration (2023a).

Finally, some solutions corresponding to values of the argument of the periastron ω are accompanied by a very large value of their estimated uncertainty. Conservatively, solutions characterised by a 1-σ uncertainty larger than 2π (in radians) are considered as problematic, since this actually means that the ω is properly undefined. We preferred to discard these objects. The filtering could have been more stringent but we preferred to rely on user's discernment. As expected, the majority of large uncertainties on ω correspond to very small eccentricities ($e < 0.05$) characterised by rather large uncertainties. However, some have eccentricities that can be considered as significant.

An internal filtering of the kind performed here above was also applied to the SB1C type of solutions (except of course for the last one on σ_ω). For the SB1C solutions, we discarded objects with either flags nr.16 or nr.19 raised.

An internal filtering was similarly defined for the TrendSB1 solutions. This last filtering was on F_2, on the bad flags and on the equivalent of the above-mentioned gap in phase. Concerning bad flags for TrendSB1 solutions, we should have discarded objects with either flag nr.17 or flag nr.20 raised. However, these flags were never activated. We decided to filter out on the flags nr.15 and nr.18, although there could be some unfortunate interactions with the SB1 channel. We rejected solutions presenting a gap (in time) in the time series larger than $0.3\Delta T$.

For reasons that will become evident in Sect. 8.2.3, we also decided to discard objects with flags nr.11 or nr.12 raised. The details on the various proportions of rejected objects are given in Sect. 8.1.

7.2. Combined solutions

After the spectroscopic processing described here above, the solutions are delivered to the combiner. The combiner and combinations are described elsewhere (see Sect. 7.7 by Gavras & Arenou in Pourbaix et al. 2022) and are not extensively addressed here. This combiner explores opportunities to unveil solutions common to other channels. In the framework of DR3, two possibilities were considered.

Either the NSS combines objects from the spectroscopic channel with those issued from the photometric channel and present in the class `EclipsingBinary`, i.e. presenting eclipses in photometry (see Sect. 7.6 by Siopis & Sadowski in Pourbaix et al. 2022; Gaia Collaboration 2023a); Siopis et al. (in prep.). The main criterion for the combination is the existence of a common value for the period. If the tentative combination is successful, a global fit over the photometric and the spectroscopic data is performed. In such a case, the new (improved) parameters supersede the old one and are introduced in the catalogue in the new class `EclipsingSpectro` (see Sect. 7.7 by Gavras & Arenou in Pourbaix et al. 2022). Most of the contributors to this class are issued from our set of SB1 solutions, although some other types of solutions could also be considered (SB1C, SB2, SB2C).

Alternatively, NSS can try to find objects for which the results of the astrometric channel (Halbwachs et al. 2023; Gaia Collaboration 2023a) can be combined with objects having a solution issued from the spectroscopic pipeline. The relevant combined solutions and the related objects are then redirected towards the new class labelled `AstroSpectroSB1` (see Sect. 7.7 by Gavras & Arenou in Pourbaix et al. 2022). New fits are here also performed and the old values of the parameters are also superseded. We recall that in this case, the fit in DR3 is done in the framework of the formalism of Thieles-Innes (Halbwachs et al. 2023; Gaia Collaboration 2023a). This formalism implies that in DR3 `AstroSpectroSB1`, the combined solutions include two different values (one astrometric and one spectroscopic) for the argument of periastron. The combination with astrometry accepts several types of possible spectroscopic input (SB1, SB1C, TrendSB1, StochasticSB1).

The objects corresponding to a successfully combined spectroscopic solution are removed from the spectroscopic output. Even after the combination process, it might be that an object appears in two classes. In the framework of DR3, these are the sole allowed and/or activated possibilities. The criteria to consider a combination opportunity are described in details in Sect. 7.7 of *Gaia* archive documentation[4] (Pourbaix et al. 2022).

If, for an object, the combination with the spectroscopic output turned out to be impossible, the sole spectroscopic solution (and the object) enters the newly created class that is eponymous of the spectroscopic solution type, either `SB1`, `SB1C`, `TrendSB1` or `StochasticSB1`. Contrary to the first three, the `StochasticSB1` class is not published in DR3. The classes `EclipsingSpectro` and `AstroSpectroSB1` are not further considered in the present paper.

To avoid losing any interesting objects (and in a purely prototypical context), the combination was performed before the internal filtering. It could have been more secure to perform it after the internal filtering. Also, the combination of the astrometric solution with the StochasticSB1 one should certainly be a matter of concern. The StochasticSB1 solutions by essence present no actual periodicity and are not expected to be further combined (in principle). On the opposite, solutions that are purely spectroscopic and are ultimately populating the classes `SB1`, `SB1C` and `TrendSB1` were necessitating increased caution and were submitted to the internal filtering.

7.3. Post-filtering

On the basis of the use of results from other origins, it is possible to further refine the set of solutions that is coming out of the chain. This is made by blacklisting spurious solutions. For a good understanding of the reader, this post-filtering is explained in details below in Sect. 10.1 after the validation sections.

8. Results

DR3 is the first release providing orbital solutions for the RV time series. As such, the present results must be considered preliminary even if the performances are already good.

8.1. Tallies

Some 33 812 183 objects have validated (non NULL) RV values in DR3. Among them, some 995 315 objects obey the input criteria. We recall them: RVs always between –1000 and +1000 kms^{-1}, $G_{\rm RVS}^{\rm int}$ brighter than 12, effective temperature between 3875 and 8125 K, and `rv_renormalised_gof` above 4. Only the objects suspected of variability enter the orbital solution chain. Thus, an additional cut-off of `rv_chisq_pvalue` at 0.01 (kept below 0.01) further decreases the sample. A little more than 866 kilo-objects (866 403) entered the present main NSS-SB1 processing. If we consider the SB1+SB2 processing, the total amount is 876 644. The repartition of the results among the different types of solutions introduced in Sect. 3 is given

[4] https://gea.esac.esa.int/archive/documentation/GDR3/

Table 1. Preliminary results of the spectroscopic orbital pipeline.

Solution types	From SB1 chain	Redirected from SB2 chain
SB1	367 328	107
SB1C	3533	2
TrendSB1	168 930	–
StochasticSB1	322 420	–
Total	862 211	109

Table 2. Tally of the results from the progressive application of the internal filtering for the objects leading to an SB1-type of solutions from the SB1 chain.

Internal filtering step	Number of objects
Entering objects	367 328
$F_2 > 3$	71 227 rejected
Bad flags	2840 rejected
$K > 250$ km s^{-1}	89 rejected
Bad efficiency	27 749 rejected
Phase gap > 0.3	25 656 rejected
Non-significant period	24 251 rejected
Significance < 5	15 829 rejected
Too large σ_ω	802 rejected
Remaining objects	198 885

Table 3. Tally of the results from the progressive application of the internal filtering for the objects leading to an SB1C-type of solutions from the SB1 chain.

Internal filtering step	Number of objects
Entering objects	3533
$F_2 > 3$	1920 rejected
Bad flags	61 rejected
$K > 250$ km s^{-1}	23 rejected
Bad efficiency	1 rejected
Phase gap > 0.3	588 rejected
Non-significant period	735 rejected
Significance < 5	2 rejected
Remaining objects	203

in Table 1. The missing objects correspond to a lack of possible solutions. An important contribution comes from the objects having $N_{\rm good}$ smaller than 10.

After this preliminary main NSS-SB1 processing, the internal filtering is applied, as explained in Sect. 7.1. The corresponding numbers are given in Table 2 for the SB1-type solutions, in Table 3 for the SB1C-type solutions and finally, in Table 4 for the TrendSB1-type solutions. We give the effect of each criterion on the tallies. The internal filtering has been actually performed in the same sequence as given in the table. These tallies are however only indicative because they can be dependent on the order of the adopted filtering. Only the number of entering objects and the number of remaining objects are not sensitive to the adopted order. As explained in Sect. 2, some objects can enter the SB1 processing coming from the SB2 processing (see also the details in Damerdji et al., in prep.). Table 5 deals with the very few objects coming from the SB2 chain. Due to the filtering

Table 4. Tally of the results from the progressive application of the internal filtering for the objects leading to a TrendSB1-type of solutions from the SB1 chain.

Internal filtering step	Number of objects
Entering objects	168 930
$F_2 > 3$	87 223 rejected
Bad flags	1381 rejected
Time gap > $0.3\Delta T$	12 103 rejected
Passage to 0.99865	10 503 rejected
Remaining objects	57 720

Table 5. Tally of the results from the progressive application of the internal filtering for the objects leading to an SB1- or SB1C-type of solutions and previously redirected from the SB2 chain.

Internal filtering step	Number of objects	
	SB1FromSB2	SB1CFromSB2
Entering objects	107	2
$F_2 > 3$	64 rejected	1 rejected
Bad flags	3 rejected	0 rejected
$K > 250$ km s^{-1}	0 rejected	0 rejected
Bad efficiency	3 rejected	0 rejected
Non-significant period	3 rejected	0 rejected
Significance < 5	4 rejected	0 rejected
Too large σ_ω	0 rejected	–
Remaining objects	30	1

Table 6. Repartition of the 198 885 SB1-type solutions among the various classes of the catalogue.

Class	Single	Duplicated	All
`SB1`	175 725	5602	181 327
`AstroSpectroSB1`	15 227	58	15 285
`EclipsingSpectro`	71	1	72
`Orbital`	55	5291	5346
Others	17	371	388
Total (all classes)	191 095	11 323	202 418
Total (first three classes)	191 023	5661	196 684
Not published	–	–	2201
Total (distributed)	–	–	198 885

Notes. Single: the solution appears in one class only; Duplicated: the solution appears in two classes and should thus be counted only once to avoid redundancy.

on flags nr.11 and nr.12, these objects do not pervade the final SB subcatalogue. No class is associated to these objects.

We have now to estimate the number of solutions from our pipeline that were removed because they were combined and are now appearing in the `EclipsingSpectro` and `AstroSpectroSB1` classes. Since we are only interested here in the disappearance of some solutions from our output due to the combination process, we can formally consider the tallies as if the combination were made after the internal filtering. The results of the combination mechanism for the 198 885 SB1-type solutions are detailed in Table 6. The column marked 'Single'

Table 7. Repartition of the 57 720 TrendSB1-type solutions among the `TrendSB1` and the `AstroSpectroSB1` classes of the catalogue.

Class	Count	Period	No period
`FirstDegreeTrendSB1`	24 083	8391	15 692
`SecondDegreeTrendSB1`	32 725	16 989	15 736
Total TrendSB1	56 808	25 380	31 428
`AstroSpectroSB1`	891	–	–
Not published	21	–	–
Total	57 720		

corresponds to the solutions appearing in one class only, whereas the column marked 'Duplicated' refers to the objects appearing in two classes. All the 198 885 objects are classified, except 2201 that will not be published due to a rejection during the process. Among these 2201 objects, 2107 correspond to a rejection at the level of the post-filtering which is detailed below. Some 94 other objects were rejected for other reasons. The majority of these 94 unpublished objects are corresponding to tentative combinations that finally failed and the objects were not re-injected in our SB subcatalogue. The `Orbital` and Others classes are given for the sake of completeness but they are not further considered either in the present work nor, for some, in DR3. Only the first three classes of the Table 6 are of importance for the present work. We will in addition restrict the present analysis of the SB1-type solutions to the `SB1` class. To illustrate the approach, we will further consider the case of `EclipsingSpectro`. According to Gaia Collaboration (2023a), some 155 objects are in this class. Over these 155 objects, 1 was originally attributed a solution of type SB2C, 1 of type SB1C, 2 of type SB2, and 151 of type SB1. Among these 151 SB1 solutions coming out of our main processing, 79 were internally filtered out. Therefore, some 72 solutions were dispatched to `EclipsingSpectro` and did not reach the `SB1` class.

The solutions of type SB1C were cross-correlated along the same lines as the SB1 ones. No object was classified in another category by the NSS combiner. One object was rejected at the level of the post-filtering. Therefore, all the 202 (i.e. 203–1) SB1C solutions entered the eponymous class `SB1C`.

Finally, the TrendSB1-type solutions were compared with the astrometric channel results and 891 turned out to be redirected into the `AstroSpectroSB1` class. They thus are not present in the `TrendSB1` class, leading to an amount of 56 808 objects. The details are explicited in Table 7. The majority of the 21 unpublished objects are corresponding to tentative combinations that finally failed and the objects were not re-injected in our SB subcatalogue. In this table, we also counted the `TrendSB1` objects that presented or not a significant period at the level of the main processing (flag nr.13 raised means no period found).

After the combination process, a further step of filtering took place. This one was called post-filtering, and the motivation and practical details are described below in Sect. 10.1. The post-filtering was done on the basis of data coming from outside the RVS channel and is basically a process of blacklisting aiming at improving the purity of the catalogue. However, it should be made clear that it is partly constrained by the fact that the spectroscopic solutions are further used by the combiner, and thus we should be careful not to eliminate good solutions that could be useful downstream in the process. In the present section we restrict ourselves to the effect of the applied post-filtering on

Table 8. Results of the application of the post-filtering.

Class	Countings before post-filter	Discarded at post-filter	Catalogue after post-filter
`SB1`	183 434	2107	181 327
`SB1C`	203	1	202
`EclipsingSpectro`	72	0	72
`AstroSpectroSB1`	36 082	2615	33 467
`TrendSB1`	56 808	0	56 808

Notes. The first four classes can be found in table `nss_two_body_orbit` whereas the last class can be found in table `nss_non_linear_spectro`.

Table 9. Various contributions to the class `AstroSpectroSB1` from the point of view of the spectroscopic solutions.

Original solution-types	Before filtering	After filtering
SB1	16 980	15 285
SB1C	0	0
TrendSB1	6104	891
StochasticSB1	10 383	10 383
Total	33 467	26 559

the countings of the objects among the various classes. Table 8 details the effect on the number of objects considered in the important classes. The amount of objects before the post-filtering and after it are mentioned. The column after post-filtering gives the countings that are corresponding to the final DR3 catalogue.

The spectroscopic channel presented here had a strong importance in the building of the `AstroSpectroSB1` class and a last information of interest must be given here. Namely, the origin of the various spectroscopic solutions that were combined with other types of data in order to build up the `AstroSpectroSB1` class. The detailed origins are explicited in Table 9 for the four kinds of solutions. We dwell again on the fact that none SB1C-type solutions were combined. We also notice that some combination has been performed on the basis of the StochasticSB1 solutions.

A few pie charts illustrating these tallies are available in Appendix D.

8.2. Some illustrative results

In the present section, we want to illustrate various solutions pointing out the efficiency of the code. Here below, we adopt the convention to express the uncertainties as error bars of $\pm 1\sigma$.

8.2.1. Good results

Here, we introduce a few examples of the solution fitted to selected RV time series. The selection is somewhat arbitrary and is essentially motivated by the illustration of the efficiency of the chain. All these illustrating good solutions are part of the DR3 spectroscopic orbit catalogue (SB subcatalogue). Figure E.1 illustrates various orbital solutions that are part of the `SB1` class. A small subset of these `SB1` solutions is also provided in Fig. 6. The displayed phase diagrams cover a palette of values

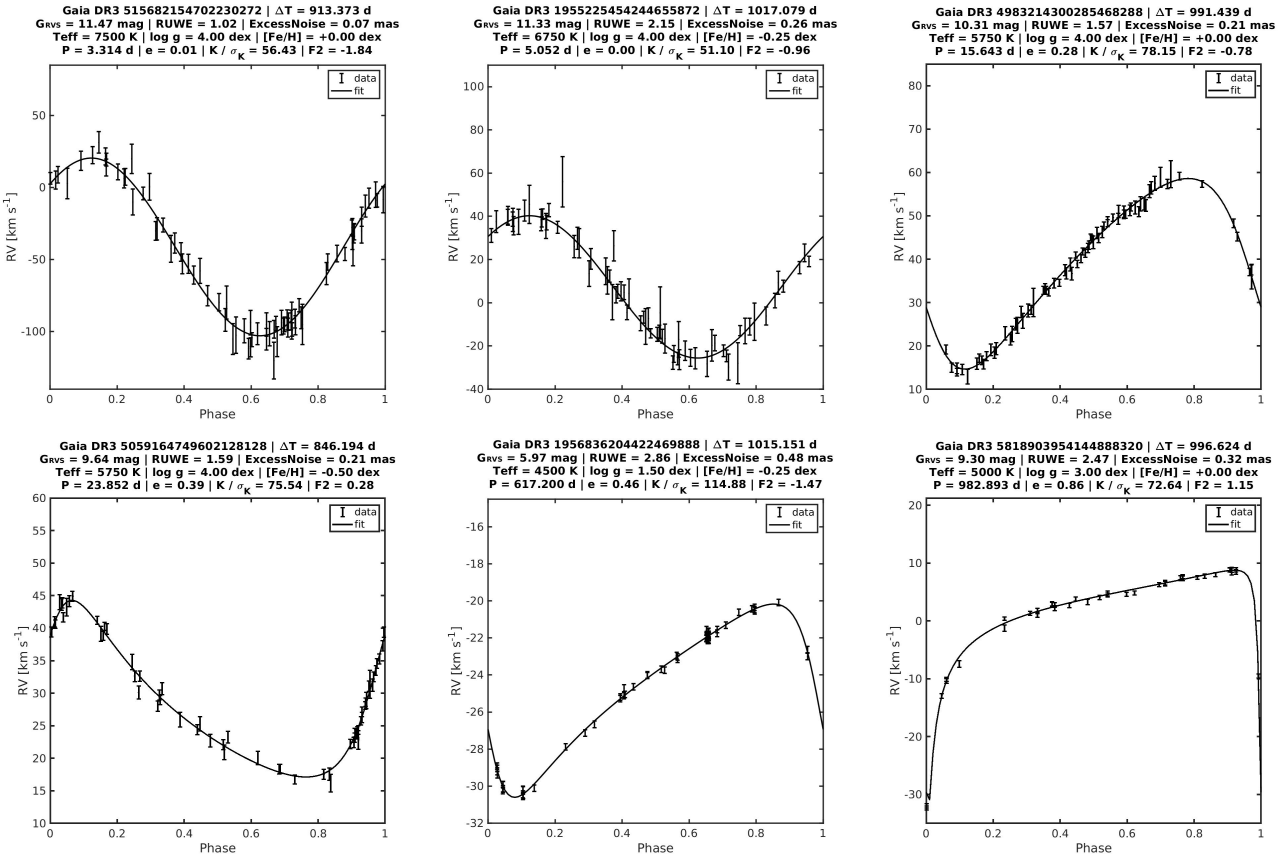

Fig. 6. Several examples illustrative of the output of the pipeline for SB1-type solutions. Each panel concerns an object whose *Gaia* name is given in the header. Each of the panels shows the phase diagram containing the folded RVs (data points at the centre of the $\pm 1\sigma$ error bars) along with the fitted orbital solution. The header gives also the time span covered by the time series (ΔT). The second line gives the $G_{\rm RVS}^{\rm int}$ magnitude of the object, the astrometric `ruwe` and the astrometric excess noise, whereas the third line concerns the physical parameters identifying the template used. The last line indicates the period, the eccentricity, the significance, and the statistic F_2 related to the adopted solution. The objects are ordered by increasing period.

for the period associated to the SB1 solution. It is a first positive point that the pipeline is able to treat such a large domain of periodicity. At mid to large period, we also present a few examples of highly eccentric orbits, up to $e \sim 0.9$. It is a matter of fact that a good fit to a very eccentric orbit necessitates a large number of data points in the time series, particularly near the periastron. An homogeneous distribution in phase is another necessity.

Since the vast majority of the analysed time series covers a span of time around 800–1000 d, the solutions with short periods are extremely convincing since the large number of cycles folded in the phase diagram do not mitigate the goodness of fit and this can be associated to a high coherency of the periodicity. This conclusion holds when more than three cycles are covered (i.e. $P \sim \Delta T/3$, or less). For larger periods, the coherency of the variations is not fully guaranteed by the quality of the goodness of fit. At larger values ($P \sim \Delta T$), we are no more able to guarantee the coherency of the variation, since only one cycle is observed and this could correspond to a transitory event, not to a periodic variation. In any case, a tentative fit is still more indicative than no fit at all. However, the SB1-type solutions with such large P up to $1.5\,\Delta T$ (the adopted upper value allowed) should be considered as not secure as periodic solutions and the reported period is at the very best an indicative timescale. Indeed, for given periods between ΔT and $1.5\,\Delta T$, the periods are not well determined and could even be potentially biased, due to the non-linear nature of the fit interacting with a possible bad phase coverage. This topic is further addressed in Sect. 9.4.1 and in Sect. 9.4.2. As illustrated below, the latter bias could induce biases on other parameters (see e.g. the case of the eccentricity in Fig. E.5). The objects *Gaia* DR3 6422103351757549056, *Gaia* DR3 2509001872917567360, and *Gaia* DR3 5818903954144888320 are certainly concerned by this possible pitfall.

It is interesting to notice that several objects in the sample correspond to a quasi-circular solution: the eccentricity is very small, even perhaps not significant. A perfect illustration (see Fig. 6) is the case of *Gaia* DR3 1955225454244655872, whose eccentricity is $e = 0.004$ with a 1-σ error of 0.021. This object most probably corresponds to a circular orbit, but the adopted solution is not an SB1C but an SB1. This SB1-type solution is associated to a solution efficiency of 0.16 and is thus on the verge to be processed with an SB1C model. This is due to the fact that we adopted the philosophy of not enforcing a circular solution even if the eccentricity turned out not to be significantly different from zero. We only adopt a circular solution when the structure matrix is such that it could not be properly inverted in the classical way and thus a full SB1-type solution is not possible. We are consequently constrained to decrease the number of free parameters. In fact, *Gaia* DR3 1955225454244655872 benefited from a large number of observations (good transits) and this large number of data points allowed us to fit an SB1 solution even for a very small apparent value of the eccentricity.

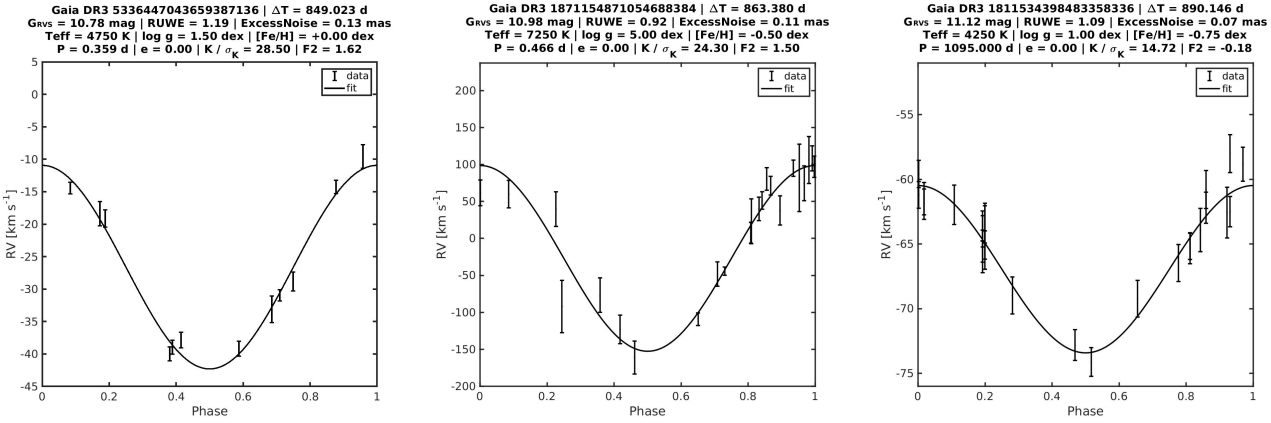

Fig. 7. Same as Fig. 6 but concerning the SB1C-type solutions.

Another example of truly circular orbit could be e.g. *Gaia* DR3 5156821547022302 72. Starting from these examples, we encourage the user of the catalogue to understand that the classes SB1 and SB1C differ by the fitted model and NOT by the true nature of the object. For example, someone interested in objects in circular orbits, i.e. to study only genuine circular orbit, should certainly consider both SB1 and SB1C classes. If a user wants to transform an SB1-type solutions into an SB1C one, they can certainly use the period, the centre-of-mass velocity and the semi-amplitude as given by the SB1-type solution. However, it should not be overlooked that the argument of periastron is thus modified, since the T_0 for SB1C solutions corresponds by definition to the time of maximum velocity. Therefore, the T_0 must be accordingly corrected for this change of ω. Otherwise, the SB1-type solution is applicable to objects in true circular orbits.

Another interesting and enlightening case is the one of *Gaia* DR3 4983214300285468288. As can be seen from Fig. 6, the fit of the SB1-type solution is very good and no individual data point deviates significantly from the model. However, an expert eye (not the pipeline) will detect that around phase 0.65, the observed RVs are systematically grouped below the model and then above the modelled curve. Both groups of data points contain some 10 points and the offset could thus be significant (at least marginally) although the grouping is observed a posteriori. In any case, if this anomaly at phase 0.65 is real (and this will become clear in future releases), it is hardly explained in the framework of the Keplerian motion. The time corresponds to a phase where the primary is approaching the maximum velocity and thus the possible lines of the secondary are close to the maximum separation and should not influence the determination of the primary RVs (primary meaning here the star that dominates the combined spectrum). We are also (most probably) far from a possible eclipse, and this could not be due to a Rossiter-McLaughlin effect. The reported anomaly, if confirmed, will necessitate exotic effects to be explained in the context of a Keplerian model. It is worrying that several intrinsic variables (pulsating stars at least) could exhibit such a phenomenon at approximately this location in their RV curve. It is beyond the scope of the present paper to make a decision for this particular object but we would like to take this opportunity to highlight the following statement. It is not because a beautiful fit can be performed with a Keplerian model by the pipeline that the true nature of the object is the one of a binary. The membership to the SB1 or SB1C classes of solutions does not automatically imply the binarity of the object. Additional clues are necessary to converge to a sound classification. The binary nature of SB2 objects is much more convincing (although problems could also exist; see Damerdji et al., in prep.). Concerning SB1 or SB1C objects, a conclusion could be drawn if, for example, the orbit is confirmed by the photometric or the astrometric channels (classes `EclipsingSpectro` or `AstroSpectroSB1`). The combined solutions are not part of the classes treated in the present paper. For the isolated spectroscopic orbital solutions reported here, we recommend at the very least, to build up a subcatalogue constituting some kind of gold sample, or to combine with additional data (e.g. positions in the HR diagram) in order to increase the veracity of the proposed classification.

The object *Gaia* DR3 1956836204422469888 is also found on the list of HIPPARCOS-*Gaia* proper motion anomaly of Kervella et al. (2019) and the list of Brandt (2021). All the objects included in Fig. 6 and in Fig. E.1 were not previously known to have periodic RV variations.

In Fig. E.2, we exhibit a few examples of orbital solutions of the class SB1C. A still smaller subset is provided in Fig. 7. A general finding is the fact that most of the concerned time series are characterised by a rather limited number of data points. They also correspond to objects characterised by $G^{\rm int}_{\rm RVS} > 9$. This is related to the decision we took to fit an SB1C-type solution only when the SB1 failed: we must be constrained to limit the amount of free parameters. We recall here that the convention has been adopted that, for SB1C-type solutions, the T_0 represents the time of passage at maximum velocity. We note from Fig. E.2 that we have circular solutions for a few large period objects. Here also, the majority of the objects were not previously known to present a periodic behaviour in their RV.

The object *Gaia* DR3 3384500857079232512 has been observed with KEPLER and a period $P = 3.74$ d (or twice that value) has been detected (Armstrong et al. 2015). This does not seem to be coherent with our results ($P = 0.297$ d). The solution of this discrepancy will certainly be automatically found in future releases thanks to the important increase of the number of data points in the time series under analysis. This conclusion is supported by the simulations reported in Sect. 9.1.

Another very interesting case is the one of *Gaia* DR3 1871154871054688384, also known as V2197 Cyg (NSVS 5761314). From ground-based observations, this star is a well-known SB2 star. From photometric data, the system seems to be an eclipsing system on a circular orbit with an inclination of about 80°. It is in near contact and has a period of 0.4657489 d (σ=0.0000001 d). This is in good agreement with the *Gaia*

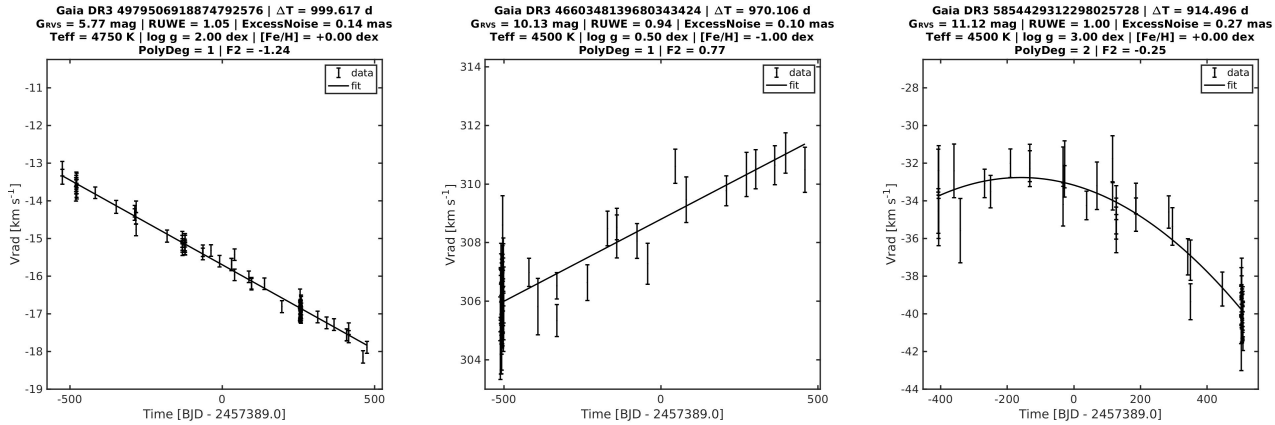

Fig. 8. Same as Fig. 6 but concerning the TrendSB1-type solutions. The order of the objects is arbitrary, except that the second degree trend has been placed at the end.

spectroscopic period $P = 0.4657410$ d (σ=0.0000034 d). The difference amounts to 2.3 σ. It is the first time ever that a spectroscopic period has been determined for this object. From the work of Nelson & Robb (2017), the primary star is of type F2-4 V and emits about 80 to 90% of the visible light of the system. With 7 spectra taken at mid-resolution, these authors derive a RV curve leading to $\gamma = -25.2$ ($\sigma = 0.6$) km s^{-1} and $K = 123.0$ ($\sigma = 0.8$) km s^{-1}. These values compare very well with the $\gamma = -27.0$ ($\sigma = 4.9$) km s^{-1} and $K = 125.7$ ($\sigma = 5.0$) km s^{-1} as derived from our pipeline. This illustrates the fact that the orbit of a relatively bright primary in an SB2 system observed as an SB1 one (here within the SB1C class) can be trusted. It is particularly interesting to look at the orbital solution in the phase diagram for this object (Fig. 7). From phase 0.5 to 1.0, the measured RVs are rather accurate: this corresponds to the primary being in front. On the opposite, from phase 0.0 to 0.5, the dispersion of the data points is larger; the unseen secondary is in front. The dispersion around the model is particularly important at phase 0.2–0.25. However, Nelson & Robb (2017) predict at these phases the presence of a Rossiter-McLaughlin effect and at phase 0.22, their model predicts a measured velocity of 40–50 km s^{-1}, again in good agreement with the *Gaia* observations. However, the agreement is not so good for the point at a slightly later phase. In any case, this could be the first time that this effect is detected by *Gaia*. This will be further tested in future releases. We would like to recall here that the measurement of RVs by cross-correlation with templates is not the best method to address this effect.

Finally, some illustrative examples of single-line trend solutions are proposed in Fig. E.3 with the small subset provided in Fig. 8. The various objects are ordered randomly except that the second-degree solutions occupy the last panels. None of these objects were previously known to exhibit RV variations. *Gaia* DR3 4979506918874792576 (≡ HD 4130) is also present in the lists of Kervella et al. (2019) and of Brandt (2021).

Gaia DR3 6070331129187853824 and *Gaia* DR3 5059000922370217728 have been observed a few times in the framework of the RAVE survey (Kunder et al. 2017); the first object of the two was also observed in the framework of the GALAH survey (Buder et al. 2018). No strong RV discrepancy with our orbital solutions is observed.

The velocities of the RV variable object *Gaia* DR3 4660348139680343424 (SV* HV5893) could be surprisingly large (see Fig. 8). Actually, this star is a known M supergiant in the Large Magellanic Cloud. It has been observed in photometry in the framework of the MACHO survey (Fraser et al. 2008) exhibiting a period of about 12.8 d. This is not directly detectable in *Gaia* RVs, but this is not either incompatible with the trend claimed here. Prevot et al. (1985) reported for this star a velocity of 310.4 km s^{-1} on HJD 2 445 662.77. The large difference in time between the two observed passages at this velocity is about 11975 d and suggests that the variation could be periodic with a scale of several thousands of days.

The solutions shown in Figs. E.1, E.2, and E.3 are just illustrative and much more objects with solutions of similar quality are present in the SB subcatalogue. This subcatalogue also contains solutions of slightly less good quality as well as solutions of intermediate quality. In Appendix E, we provide a few examples of SB1-type solutions selected from the SB subcatalogue at random but with F_2 around 0 (Fig. E.4). This is to have fits of similar quality but for which the number of data points is lower (10 to 45) than in Fig. 6. Although the quality of these solutions is sometimes less evident, only very few could be easily questionable. Since we decided to push the software to its limit, solutions of more limited quality are equally present. For that reason, we encourage the general user to further filter the data. This is possible thanks to the quality factors that we deliver for each object (see Sect. 6.4). For a general statistical use, this should generate no problem, provided the user adopts the filtering to its particular application. We encourage the user to be very careful about the results on individual objects that could not benefit from the statistics. Some further information and advice can be found below in Sect. 11.2.

8.2.2. Problematic cases: multi-variability

In the framework of DR3, the SB1 pipeline is intended to search for orbital solutions. It is however designed to treat the cases which present only one type of variability (either SB-type with a single period or pure trends). If one object presents two types of variability, most of the time the algorithm could fail to find an adequate solution. We present here a typical case of badly handled data. Figure 9 shows the time series of a star observed by *Gaia* RVS ($\Delta T = 845.177$ d). The data are unpublished and the solution is not part of the DR3 catalogue. The star is thus kept anonymous. The pipeline detected for this object a linear trend of about 1.15 km s^{-1} y^{-1}. The solution corresponds to $F_2 = 20.42$. This is easily understood because the RV curve exhibits extra-variations with a kind of oscillation with a period a little over

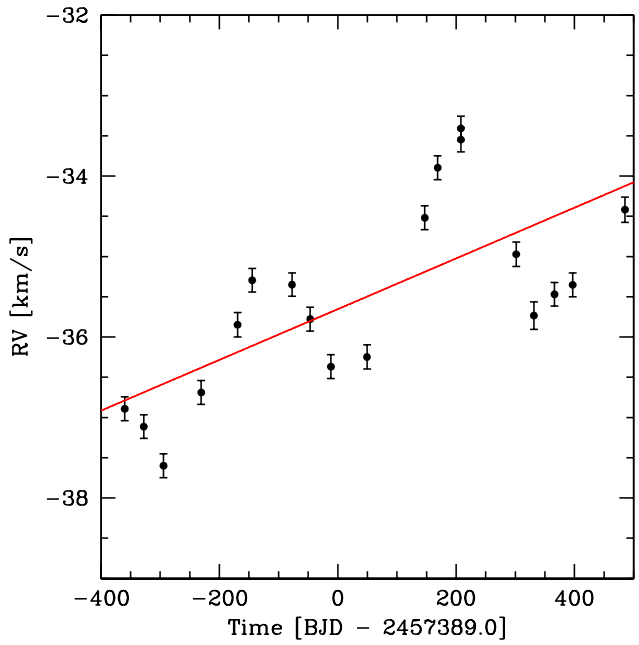

Fig. 9. Illustration of a time series of RVs for an anonymous object (unpublished). The dots represent the observed RVs with their error bars. The red straight line shows the trend fitted by the pipeline. It corresponds to $F_2 = 20.42$ and thus the corresponding solution is rejected.

300 d. The variation is larger than the error bars and is thus significant. It necessarily induces a large value for F_2 coming from the fit of the trend, and the corresponding solution has thus been rejected from the DR3 catalogue. This object actually presents two superposed, different, types of variability. A trend that could be associated to a long period binary and, in addition, a roughly sinusoidal variation that could be due to an SB1 system in a triple system, or to an intrinsic RV variation or to something else. If, by chance, the pipeline would have first detected and reported an SB1 solution, the existing trend would have also implied a large F_2 associated to the SB1 fit and the solution would also have been rejected. This is clearly a weak point of the DR3 chain.

For the coming releases, the problem will be corrected and the pipeline will be able to process properly this kind of multiple variability by simultaneously fitting two types of solutions. This will open the door to the treatment of objects like triple or quadruple SB1 stars, as well as of intrinsic variables in a binary system.

8.2.3. Problematic cases: fake SB1 and spectrocentre

We present here the illustrative case of HD 223323. This star ($G = 6.968$) is a well-known binary object of type SB2 that has been well studied from the ground; it has a period of 1175.1 d (Griffin 2007). The two components are rather similar, and only a merged spectral type in the range F2-4 IV-V is reported (see references in Griffin 2007). An accurate orbital solution has been published by the latter author, and the elements of interest for the present discussion are given in Table 10. Figure 10 illustrates the fitted orbital model published by Griffin (2007). The ground observations at phase 0.98 (almost at larger separation; see Fig. 5 of Griffin 2007) suggest that the lines of the two stars are almost deblended. The star has been observed with *Gaia* and is present in the DR3 catalogue under the name *Gaia* DR3 2852594583674129152. A total of 24 measurements ($N_{tot} = 24$) were collected by the RVS over a span of time of $\Delta T = 889.57$ d.

Table 10. A few parameters of interest from the Griffin orbital SB2 solution for HD 223323 as well as from the *Gaia* equivalent (SB1). Note the spurious values for *Gaia K* and mass function.

Parameter	Griffin	*Gaia*
P (d)	1175.1±1.3	977.9±104.3
γ (km s^{-1})$^{(*)}$	–9.56±0.06	–9.94±0.10
K_1 (km s^{-1})	16.38±0.11	1.188
K_2 (km s^{-1})	16.40±0.16	...
M_1/M_2	1.001	...
e	0.604±0.003	0.586±0.150
ω (degree)	77.8±0.60	56.0±11.0
$f(M_1)$ (solar mass)	0.272±0.006	9.07×10^{-5}
$f(M_2)$ (solar mass)	0.273±0.008	...
σ_{fit} (km s^{-1})	0.53	1.61

Notes. $^{(*)}$The systemic velocity is expressed in the Griffin RV reference system (see also Sect. 9.4.1).

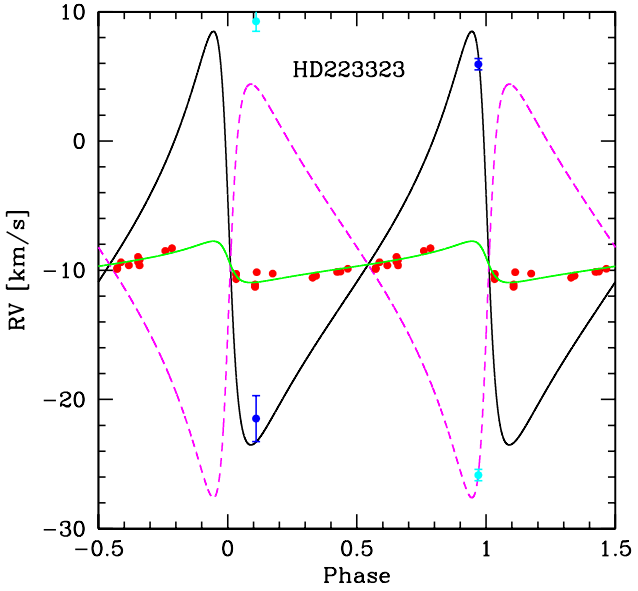

Fig. 10. Known ground-based orbital solution for the binary star HD 223323 (*Gaia* DR3 2852594583674129152). The primary orbit is illustrated with a black continuous line and the secondary with a magenta dashed line. The SB1 RV measurements are indicated by red dots, whereas the SB2 ones are marked in blue (for the primary) and in cyan (for the secondary). The expected, typical RV of the blend is plotted in green (see text for explanations).

Two of the transits exhibit a double-line spectrum, and RVs have been derived as explained in Damerdji et al. (in prep.). For the remaining 22 transits, the star has been measured as a single-line object. Owing to the small percentage of time when the star appeared as double-lined, the pipeline decided to discard the two corresponding transits (see the treatment described by Damerdji et al., in prep.). Therefore, the object entered the SB1 chain with $N_{good} = 22$ data points. An SB1 fitted model was presented in the catalogue with a period 977.9 ± 104.3 d. This value is in principle only indicative, since it is larger than the total span of time ΔT by a factor 1.11. In any case, the *Gaia* period is in good agreement with the precise Griffin period since it is only at 1.89 σ. The value of the period will automatically improve in the coming releases due to the increase of the ΔT.

The other parameters of interest derived by *Gaia* are available in Table 10. The measured *Gaia* RVs are plotted in Fig. 10. The red dots represent the RVs derived by the single-line channel (they are put in the Griffin reference system by adding 1 km s^{-1}; see Sect. 9.4.1 below). The other data points represent the primary (in blue) and the secondary (in cyan) RVs as derived in CU6 (Damerdji et al., in prep.). They were also corrected to the Griffin reference zero point. It is remarkable that the *Gaia* solution yields a very small value for the semi-amplitude. The interpretation is quite straightforward. Except for the phases near the extrema in RVs, the resolution of the RVS and the predominance of the broad Calcium lines of the triplet makes that the lines never deblend at the other phases. After the rejection of the two double-line transits, the star becomes a fake SB1 which is actually a never-deblending SB2. There is a tendency for the position of the blended lines to moderately follow the behaviour of the primary variation. According to Griffin (2007), the lines of the secondary are slightly broader and the lines of the primary slightly deeper. This thus implies that the primary lines are moderately dominating the blend. The secondary line depths are roughly 90 % of those of the primary one (we will call this ratio B = 0.9). The effect of the width of the lines of the secondary is attenuated since the lines are rather strongly blended. Tatum (1968) presented a computation of the position of the minimum of the blend of two Gaussian profiles (which is unique in the case of a strong blend) as a function of the widths and of the depths of the two profiles. Following their work, using the same notation as in their Eq. (5), and deciding to neglect the difference in broadness of the lines, we derived a position of the minimum which is given by

$$x = \frac{B}{1+B}. \tag{17}$$

According to the latter author, x adopts a value of 0 if the minimum is at the velocity of the primary and 1 at the secondary position. In the present case, we derive $x = 0.45$ which means that the minimum is indeed attracted by the primary. As an approximation, we consider that the position of the minimum and the one of the blend are similar. In the same approach, we consider that the position of the blend can be approximated by a weighted sum of the velocities of the primary and of the secondary. It follows

$$RV_{\rm blend} = 0.55\,RV_{\rm pr} + 0.45\,RV_{\rm sc} \tag{18}$$

consistently giving more weight to the primary velocity. Starting from the modelled RVs by Griffin, we computed the evolution of the position of the blend as a function of phase. This is illustrated in Fig. 10 by the green line. Despite the fact that this line is in no circumstances a fit, the agreement with the red dots is very good. This demonstrates the interpretation of the *Gaia* results as a fit of the blend. A more rigorous approach would be much heavier and is not in the scope of the present analysis which is intended to remain purely illustrative. By analogy with the notion of photocentre, we could introduce here the concept of *spectrocentre* which would correspond to the measurement of the RV of the blend (in the spectrum) constituted by the presence of two sets of undeblended lines. This could thus be termed a *spectrocentre* effect[5].

[5] It should however be pointed out that this effect is mainly dependent on the characteristics of both sets of lines (with a different contribution from line to line) and thus on the Astrophysical Parameters characterising the two stars.

It is interesting to notice that in the case of HD 223323, the most affected parameter is essentially the semi-amplitude. On the contrary, it is clear from Table 10 that the other parameters are not strongly perturbed. The eccentricity in particular is well determined, as well as the argument of periastron. This is a consequence of the slight domination of the primary on the secondary effect.

The case of the never-deblending SB2 stars is crucial. If the two object spectra are similar (except for the Doppler shift), the semi-amplitude associated to the blend will near zero. Therefore, the corresponding object will be characterised by a very-small K and also be associated to an abnormally small mass function. An extreme case with two identical stars will lead to a fake RV constant star. On the opposite, if the similarity of the two systems of lines for the two objects does not hold, the resulting semi-amplitude could of course be much larger and it is to be expected that such a case will be harder to detect. At the other extreme, if the secondary spectrum is negligible, the solution will tend towards a true SB1. The detection of the fake SB1 never-deblending SB2 is important, and nothing global has been implemented for DR3. We identify two avenues to solve the problem for future releases. The first one is to look at the evolution with phase of the broadness of the lines (even epoch by epoch). This could be done directly on the spectra. For DR3, an a posteriori look at the epoch `vbroad` values (Frémat et al. 2023) failed to produce results. An alternative would be to systematically reconstruct the mean spectrum at rest (according to the SB1 solution) and to detect line-profile variability through inspecting the dispersion of the reconstructed mean spectrum in the neighbourhood of the main spectroscopic lines. Due to the general architecture of the *Gaia* pipeline, these possibilities were impossible to implement in the DR3 pipeline. Similarly, we did not investigate in DR3 the possibility to use the information from astrometry to alleviate the problem.

8.2.4. Problematic cases: possible spuriously variable RVs

Some stars are associated with an orbital SB1-type solution that looks very convincing, and is identified as good according to the various quality tests. These stars thus enter the orbital solution catalogue. However, ground-based observations do not confirm these orbital variations. An illustrative case is the one of *Gaia* DR3 702159502868472832. Its RV time series suggests a significant period of 1.084400 ± 0.000058 d with a probability of 0.997 and a resulting folded RV curve that is given in Fig. 11. With a significance of 15.41 and an $F_2 = -0.36$, there is no solid reason to doubt the nature of this object. However, besides the *Gaia* RV curve and the fit, we plot also in the figure the few folded RVs that can be found in the literature from the APOGEE survey DR17 (Abdurro'uf et al. 2022). The two data sets are clearly not compatible. This suggests that some of the spectroscopic data from *Gaia* could not only present spurious RVs that are variable, but also that a period can be found from these measurements. This also suggests that the effect could be instrumental. The origin is currently not securely identified. A small number of objects enter this category, but the related sparsity could be due to the small percentage of objects having ground-based observations. This is supported by the validation performed in Sect. 9.4 below. Therefore, it is primordial to discover the origin in order to improve the robustness and purity of the future Data Release catalogues. We listed a few such doubtful cases that were detected here: they are *Gaia* DR3 665727189407951 36, *Gaia* DR3 286652890498864768, *Gaia* DR3 487550398091020160, *Gaia* DR3 164423034699251968,

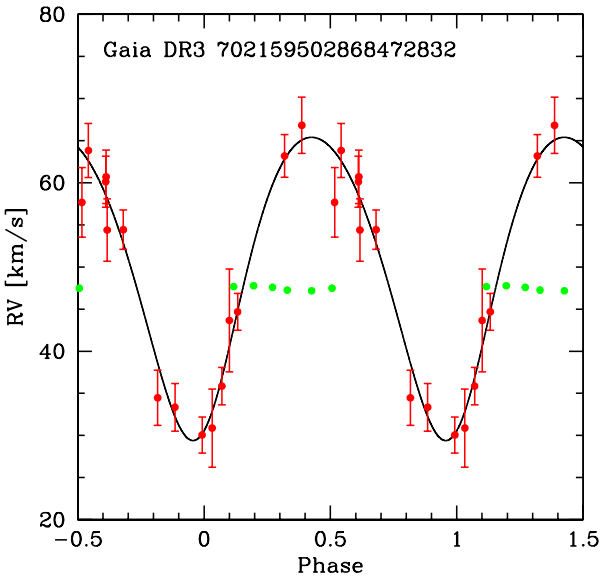

Fig. 11. Folded RV curve for the star *Gaia* DR3 702159502868472832 corresponding to the selected period of 1.084400 d. The observed RVs and their errors are plotted in red, along with in black the fitted RV curve. The green dots represent the ground-based RVs of the APOGEE survey folded with the same ephemeris.

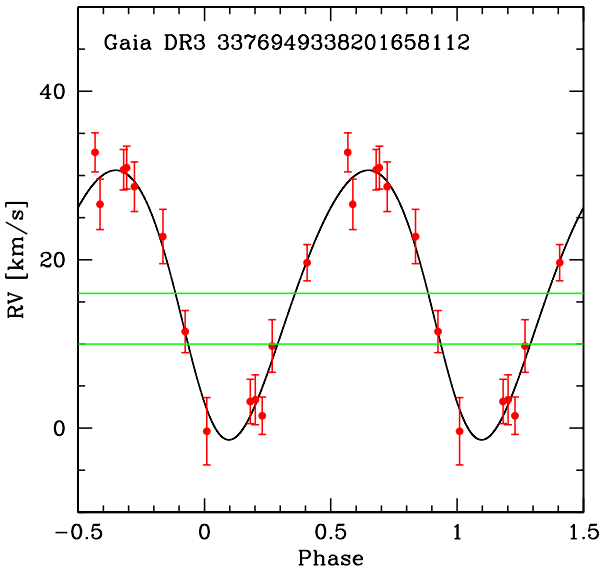

Fig. 12. Folded RV curve for the star *Gaia* DR3 3376949338201658112 corresponding to the selected period of 0.657582 d. The observed RVs and their error bars are plotted in red along with in black the fitted RV curve. The green lines delimit the domain where the ground-based RVs part of the LAMOST survey are distributed. See also the Figure A.3 of Bashi et al. (2022).

Gaia DR3 5776001255235744000, and the illustrative case *Gaia* DR3 702159502868472832. All of them have been compared to APOGEE DR17 data.

Another similar case has been detected by Bashi et al. (2022) but this time issued from comparison with the LAMOST RV survey (Cui et al. 2012[6]). The object is *Gaia* DR3 3376949338201658112 and a significant periodicity has here also been identified ($P = 0.657582 \pm 0.000022$ d) and the corresponding RV curve is shown in Fig. 12. Again the RV variations are convincing as well as the fit. The object is one of very few that have been observed at several epochs in the framework of the LAMOST survey. All the corresponding RVs are contained between the two green lines, further underlining the incompatibility of the ground-based data with the *Gaia* ones.

The fact that two ground-based surveys disagree with the results from *Gaia* underlines the fact that some instrumental effects are most probably present in the *Gaia* data. It is particularly intriguing that the RV time series exhibit what could appear as a convincing periodic variability characterised by timescale not known to be problematic. Clearly, the problem is linked to the production of the RV time series rather than to the behaviour of the algorithm of orbital parameter determination, except perhaps for the period.

9. Validation of the results

The analysis of the huge data set generated by *Gaia* and the establishment of the related catalogue necessitates some consolidation mechanism. This mechanism can be made of several approaches: they are described in the following subsections. The aim is evidently to test and fine-tune the pipeline itself in order to address its efficiency but also to discern the unavoidable biases. The understanding of the pipeline behaviour is directly related to the quality of the SB subcatalogue and looking at its characteristics from all angles is a necessity. This validation process can be made on the basis of the internal consistency of the solutions but also on the basis of a comparison between the parameters delivered from the *Gaia* data with those obtained from the literature. This last approach is limited to the dramatically smaller amount of objects that have been the subject of a previous study. The validation paves the way for the determination of the global selection function.

9.1. Simulations

As a very first step, it is necessary to study the ability of the pipeline to recover the period. The study described here is based on the generation of synthetic RV curves of eccentric systems (with a random e), and of circular systems. Both sets of RV curves are generated with a random phase. We considered a number of data points of $N = 10$, 17, 30, and 80. The times of observations are mimicking the *Gaia* time series and are drawn in the same manner as in Sect. 4.4. The adopted S/N ratio $s = 2\,K/\sigma_{\rm RV}$ is considered for values 5, 10, and 50. Figure 13 reveals the recovery rate (in percent) for various true periods. The period is assumed to be recovered if its value does not differ from the input period by more than $5\,\sigma$.

For s equal to or larger than 10, the recovery rate is mostly better than 80 percent except when the number of data points is too low (see e.g. $N = 10$). This suggests that a too low number N is a problem for the algorithm. It is also clear that an s of 5 is equally detrimental. In Table 11, we give the corresponding value for the equivalent `significance` given by

$$\sim \frac{\sqrt{N}}{2\sqrt{2}}\,s, \tag{19}$$

as derived from Eq. (12) of Lucy & Sweeney (1971).

As an alternative, a cut-off at a `significance` of about 10 is certainly a good attitude. It is interesting to note that the recovery rate is slightly better for circular orbits (see right panel of Fig. 13).

[6] https://dr6.lamost.org

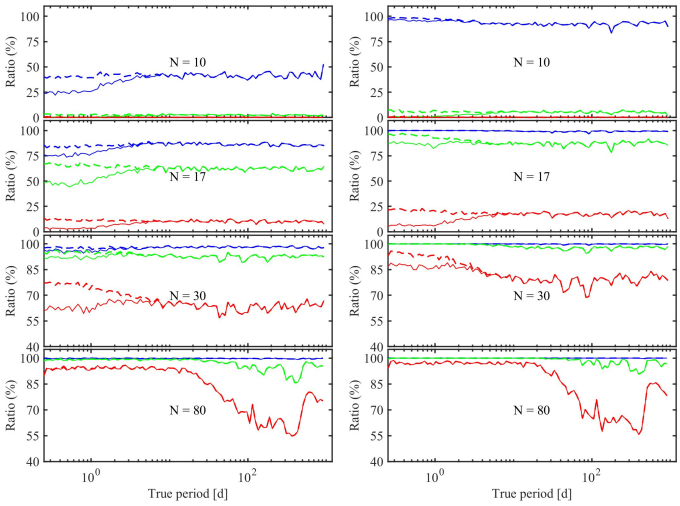

Fig. 13. Performance for the period recovery of the pipeline applied on simple simulated data. N is the number of data points in the time series. The colour curves correspond to different S/N ratios s (blue: 50, green: 10, red: 5). The dashed curves correspond to the recovery before the application of the statistical test on the period significance level; the continuous curve refers to the recovery rate after having applied the selection. Left: eccentric orbits; Right: circular orbits.

Table 11. `Significance` as a function of the number of data points N and of the S/N ratio, s.

s →	5	10	50
N = 10	5.6	11.2	55.9
N = 17	7.3	14.6	72.9
N = 30	9.7	19.4	96.8
N = 80	15.8	31.6	159.1

9.2. Internal verification of the results

9.2.1. SB1 and SB1C classes

The analysis presented here below is essentially based on the output of the spectroscopic pipeline after the internal filtering. The SB1 data set contains 183 434 objects. The basic reason for this choice is that the principal motivation is to first scrutinise the behaviour of the pipeline. The post-filtering is based on information from other channels and could thus be more complex to interpret, although it is necessary in order to constitute the DR3 SB subcatalogue. In any case, for DR3, the differences between the analysed data set and the data set after post-filtering that contains 181 327 objects, are almost negligible from a statistical point of view. The resulting data coming out of the pipeline were systematically inspected in order to verify the conformity to the expectation and the expected statistical distribution of the fitted parameters. Basically we are mainly treating the SB1 class excluding the SB1 solutions that were combined downstream with the solutions from other channels. The SB1C class contains too few objects to allow us to perform a fully sound statistical analysis. However, these solutions correspond to a very simple model that should not bear any particular problem.

The distribution of the number of good transits (i.e. data points) for the SB1-type solutions processed by the pipeline is illustrated in Fig. 14. The objects having less than 10 transits were not further considered as explained above. This cut-off below 10 transits should not discard too many objects. Typically,

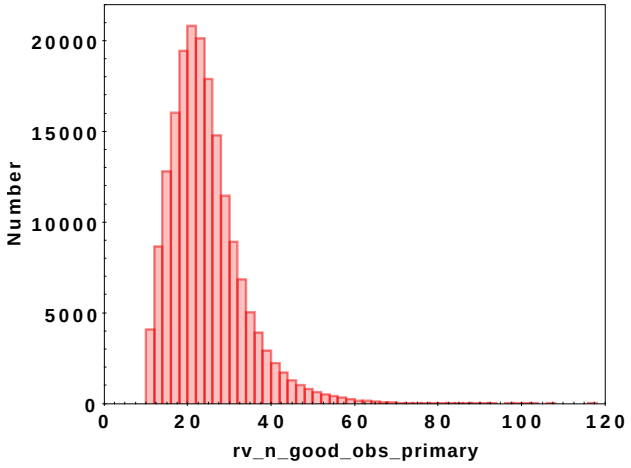

Fig. 14. Histogram of the distribution of the number of good transits (`rv_n_good_obs_primary` ≡ $N_{\rm good}$) associated to the different objects with SB1-type solutions processed by the pipeline. The typical number of data points is slightly larger than 20. The bin width is 2 transits.

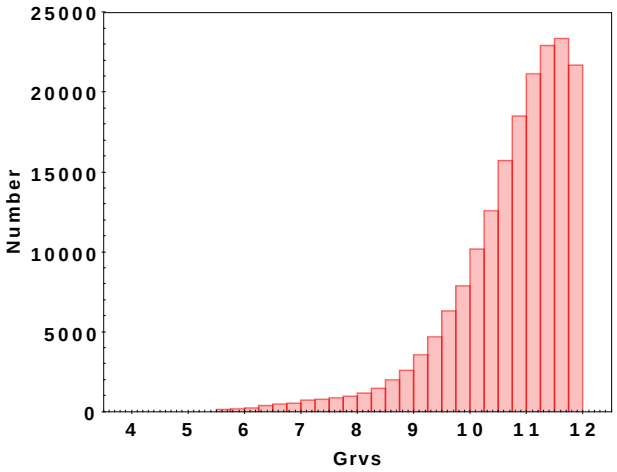

Fig. 15. Histogram of the distribution of $G^{\rm int}_{\rm RVS}$ for all the 183 434 SB1 solutions. The bin width is 0.25 mag.

we have slightly over 20 data points in our various time series. The largest value is 228 (not shown here) but these objects are exceptional and are corresponding to the EPSL observations. As a first test, we examine the distribution of the magnitude $G^{\rm int}_{\rm RVS}$ of all solutions in the SB1 class. Figure 15 exhibits the histogram of $G^{\rm int}_{\rm RVS}$ for all the 183 434 objects. We immediately notice that no object is present between $G^{\rm int}_{\rm RVS}$ = 2.76 and 5.5. This is an artefact of the selection based on the `rv_renormalised_gof` at the level of the input data. This quantity is defined when $G^{\rm int}_{\rm RVS} > 5.5$ and thus below this threshold, the value for these objects contains NaN even if these bright objects should be kept. This is a weakness of DR3 and it will be corrected in future releases. On the other hand, it is also clear that signs of incompleteness are starting to appear above magnitude 11.5. A minimum of ∼2000 objects are lost between magnitude 11.75 and 12.0. This is most probably an effect of the increased importance that noise could have for the fainter objects, thus leading to badly determined solutions. However, an incompleteness also appears in the global RV output of the CU6 pipeline (see Fig. 4 of Katz et al. 2023) for objects fainter than $G \sim 12.15$ (corresponding for a solar-like star to $G^{\rm int}_{\rm RVS} > 11.5$). Therefore, the difficulties to measure RVs on noisy spectra, the selection of variable objects in noisy

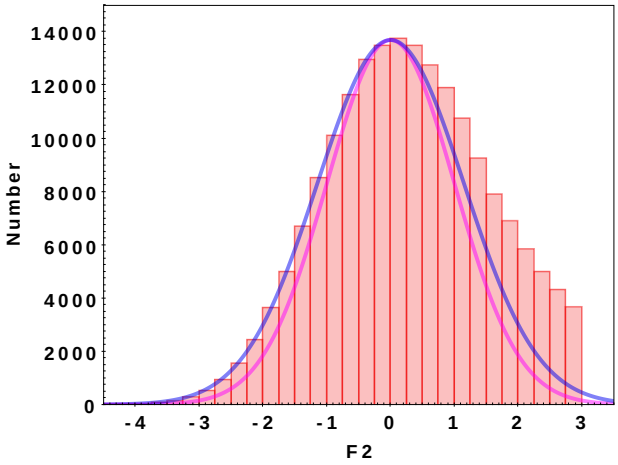

Fig. 16. Histogram of the distribution of the F_2 statistic of the different fits that are adopted as good ones. The bin width is taken to be 0.25. The cut-off at $F_2 = 3$ corresponds to the adopted threshold. The magenta curve is a Gaussian function corresponding to a zero mean and a dispersion $\sigma = 1$. This is the expected behaviour of the F_2 statistic. The blue curve represents a Gaussian with $\sigma = 1.15$. This latter function follows rather well the slope of the F_2 histogram on the left side. Broadly speaking, the F_2 function is behaving as expected except for this larger dispersion. Both Gaussian curves were adjusted on the maximum of the histogram and are consequently not normalised on their area. They are drawn to illustrate the explanation in the text.

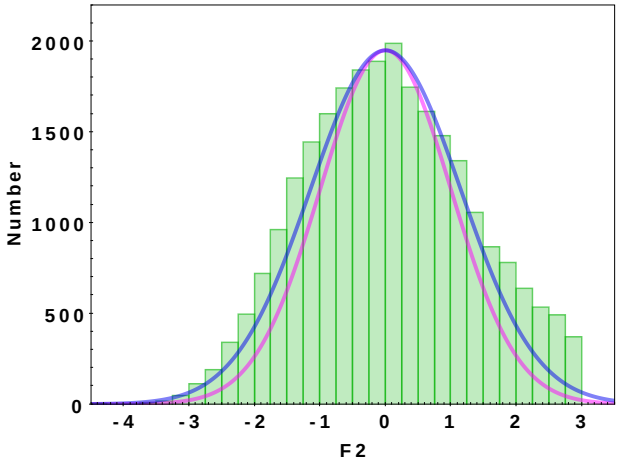

Fig. 17. Same as Fig. 16 but restricted to solutions with n in the range 4–9, i.e. $N_{\rm good}$ in the range 10 to 15.

time series and the subsequent fit of orbital solutions conspire to induce this incompleteness.

The observed distribution of the quality-of-fitting statistic deduced from the chi-square, the goodness of fit F_2, is given for the filtered data set of solutions (see Fig. 16). The recommended and applied cut-off at $F_2 = 3$ is clearly visible. Otherwise, the original decay of the histogram is going up to values of several tens. Large values of the statistic are to be associated with bad, unacceptable solutions. This is certainly so for F_2 larger than 5. The cases between 3 and 5 could be further discussed. Conservatively, we only kept values below 3. This induces the rejection of some acceptable solutions, and consequently of the related objects. Along with the observed distribution, we also plot a Gaussian around the mean zero and with a $\sigma = 1$ dispersion (in magenta). This curve represents the expected behaviour of F_2 under the null-hypothesis of a perfect Keplerian model with additive Gaussian noise (corresponding to the uncertainties associated to the individual RVs). Another Gaussian curve with a $\sigma = 1.15$ (in blue) is also shown. This latter better follows the histogram on its left tail. The side of the negative values of F_2 corresponds to solutions for which the realisation of the noise conspires to produce better fitting. On the other hand, the right tail is associated to the opposite effect where the solutions are noisier although still acceptable. It is interesting to notice that, in this region, the histogram reveals a remarkably too large number of objects. This is true even comparing to the blue Gaussian curve; an excess is thus present. They most probably correspond to solutions for objects following a Keplerian orbit (thus generating a rather good fit) but affected by additional variations (pulsations, activity, spots, ...), not due to the observational noise, but at a sufficiently low level such that the dominating effect is still the orbital variation. Among these objects, multiple systems and intrinsic variables could be the main contributors to this effect. A more sophisticated study, beyond the scope of the present paper, is necessary. However, we should not be misled by the normalisation of the comparison Gaussian curves. Actually, if we normalised with respect to the area, there would appear a deficit (roughly 25%) of objects with solutions characterised by an F_2 broadly located around zero (essentially between –1. and +1.). These solutions are displaced towards the right tail. The effect is present both on the individual χ^2 (with a selected value of n, the number of degrees of freedom) and on the global F_2 statistics. The transformation of the χ^2 values into the global F_2 statistic is not the cause of the deviation.

Coming back to the anomaly of the left tail, it could be due to a general over-estimation of the individual uncertainties on the measured RVs. This would displace the whole distribution by 15% to the left but the above-mentioned effect on the right tail would prevent any confirmation based on this right tail. This explanation would induce an effect independent of the number of data points.

Around $F_2 \sim -3.$, there is a deficit of solutions with large significance (roughly above 40). Consequently, any attempt to correct the shift to the left of the histogram by filtering out low-significance solutions, implies the appearance of a shift of the maximum of the peak in the histogram towards positive values of F_2. This is not the solution either. It is interesting to notice that, if we limit the population of the solutions to those with $N_{\rm good} \geq 21$ and thus to a number of degrees of freedom $n \geq 15$, the tail is only shifted by a factor 1.07 (7% of excess). Therefore, the observed shift seems to be more a matter of the amount of data points. The distribution of solution efficiencies is peaking at a value of 0.44. Indeed, if alternatively we select solutions with n between 4 and 9, we obtain the F_2 distribution exhibited in Fig. 17, further strengthening our conclusion. For low n, the efficiencies are peaking at values lower than 0.3. The behaviour of the left tail as a function of n could be due to the apparition for small amount of data points of overfitting or of effects due to the intrinsic non-linearities of the model that are inducing correlations among the parameters due to particularly bad distributions in phase of the observations and/or effects due to the fact that the individual uncertainties on RVs are estimators with a random distribution far from Gaussian.

We further inspected (not shown here) the distributions of the solutions as a function of two parameters, in particular F_2 versus `rv_n_good_obs_primary` or versus $G^{\rm int}_{\rm RVS}$. Nothing particular is deduced from these plots. The reason is that the 2D

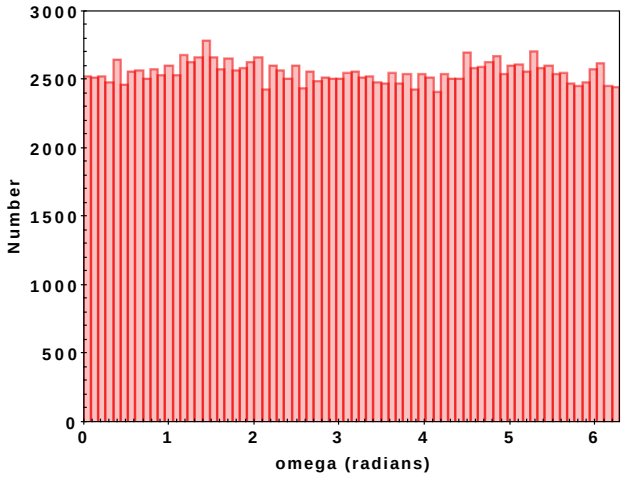

Fig. 18. Distribution of the argument of periastron over the whole SB1 sample. The distribution is pretty well homogeneous. The bin width is 0.0873 radian (5 °).

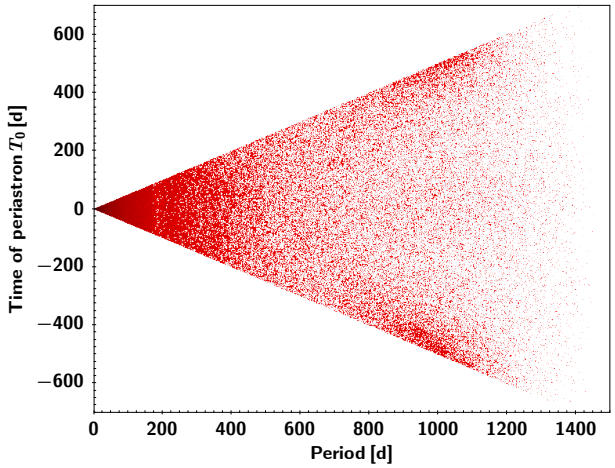

Fig. 19. Plot of T_0 versus the period for all the 183 434 objects.

distributions of the solutions are well interpreted as the product of the marginal 1D distributions.

The next step we adopted was to create the histogram of the argument of periastron distribution (see Fig. 18). The distribution is rather uniform, with a very slight tendency to have more objects at 90° ($\pi/2$) and at 270° ($3\pi/2$). This results from a very small bias because saw-tooth RV curves are easier to detect than shark-tooth RV curves, particularly at high eccentricity. The effect is practically negligible.

Another interesting plot is the distribution of T_0 versus the period. It is given in Fig. 19. Since T_0 is a cyclic variable of modulo P (see Sect. 3), an infinity of possible values are available. The data points being located around $t = 0$, the value presenting the best precision and a minimum correlation is necessarily situated between $-P/2$ and $+P/2$. It is clear that the graph is essentially as expected with a rather homogeneous density of data points slowly decreasing towards the largest periods. For periods around 900 to 1100 days, there is a slight tendency to have a larger density of solutions near the values $T_0 \sim \pm P/2$. A part of the solutions located there correspond to low significance and rather large eccentricities. This is dominantly an artefact of the decision we took to enforce an SB1 solution in this region of periods against the possibility to consider trends. Indeed, let us consider a time series of a noisy linear trend. The enforced SB1 solution will deliver a period roughly about $\Delta T/2$ (only one 'cycle' but this depends on the exact sampling). The trend folded in phase will very often mimic an eccentric curve with a jump between the last data point of the time series and the first one. Therefore, this jump, after folding, will be associated to the passage at periastron and, consequently, the returned T_0 will be attributed a value around $\pm P/2$. This conclusion is further supported by the 2D distributions of solutions (i.e. versus ω and T_0) shown in Fig. 20. We can clearly see that the excess of solutions is effectively due to the range of periods enforced to an SB1 solution (prominent zone in yellow in the bottom right panel, not present on the upper left and right ones). It is also clear that the effect is absent for quasi-circular orbits (which are fully correct) and is associated to larger eccentricities. In addition, we can clearly see that the solutions at the basis of the excess are confined to values of ω around $\pi/2$ and $3\pi/2$ (very same zone in yellow). This indicates that the trends enforced into an SB1 solution generate a RV curve that is predominantly of the saw-tooth type, further confirming our conclusions. The excess partly contributes to the excess at specific ω noticed in Fig. 18. This problem will be corrected for DR4; it will also occur at another period since ΔT will be larger. Some other effects might also be present. Note that not all the solutions in this region are problematic.

In the next figure, Fig. 21, we plot the histogram of the distribution of the RVs of the centres of mass. This histogram is very similar to the one built for the assumed constant stars for EDR3 (see Figure 4 of Seabroke et al. 2021) and for DR3 (see Figure 11 of Katz et al. 2023), which is a strong indication of the quality of these γ centre of mass RVs. The latter curve from DR3 is overplotted in Fig. 21 to facilitate the detailed comparison. The histogram for the constant stars is slightly more dispersed, but the agreement is actually good. The remaining discrepancy is a matter of sampling different stellar populations. Indeed, if we compare with a population restricted to stars cooler than 8125 K and with $G_{\rm RVS}^{\rm int}$ brighter than 12, the agreement becomes almost perfect.

The distribution of the derived semi-amplitudes is given in Fig. 22. The major part of the distribution is decreasing for increasing K. The decay law seems to be essentially linear on this logarithmic scale. Taking into account small number statistics (as seen above 150 km s^{-1}), no break is visible in the distribution function and we will retain that the distribution is rather smooth, which indicates a very good quality of the results. No anomaly is detected.

In the case of noiseless data, the difference between the two extrema of the individual RV curves is expected to be a good estimator of twice the semi-amplitude. We show in Fig. 23 a plot of the `rv_amplitude_robust` (Sartoretti et al. 2022) versus the K's. We conclude that the expected relation is certainly present. The distribution below the reference line (in blue) is essentially explained by the noise on the two values. The dispersion over the reference line is markedly larger. This is due to the fact that the `rv_amplitude_robust` has a larger dispersion, being the difference of the two extreme values (even after some cleaning). This means that it is more sensitive to persisting outliers or even slight extreme values (or noise) whereas the semi-amplitude is the result of a fit on all the data points. It appears that the discrepant data are more numerous at large K. This very small population of solutions (in red) at large K is associated to solutions with a large σ_K; the corresponding objects are characterised by large values of `vbroad` (Frémat et al. 2023). Based on a similar picture (not shown

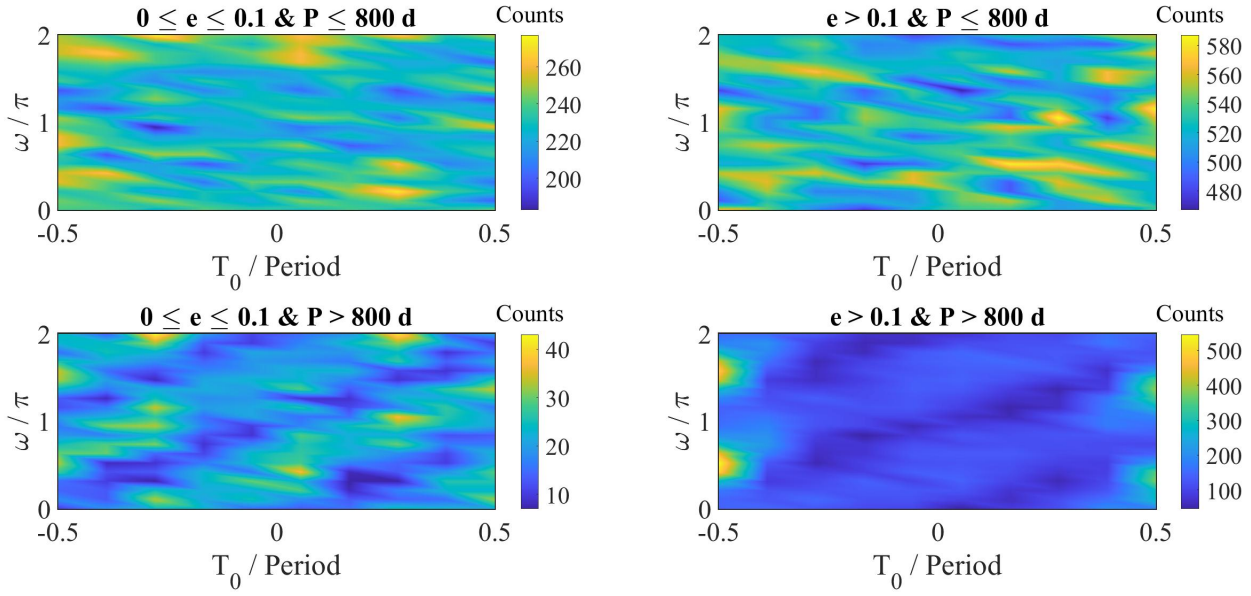

Fig. 20. Graph of the distribution of the density of solutions as a function of two parameters for various ranges in eccentricity and period. The abscissae denote T_0 normalised by the period whereas the ordinate denotes the value of ω over π. The four panels correspond to eccentricities below and above 0.1 (left and right, respectively) and periods below and above 800 d (up and bottom, respectively).

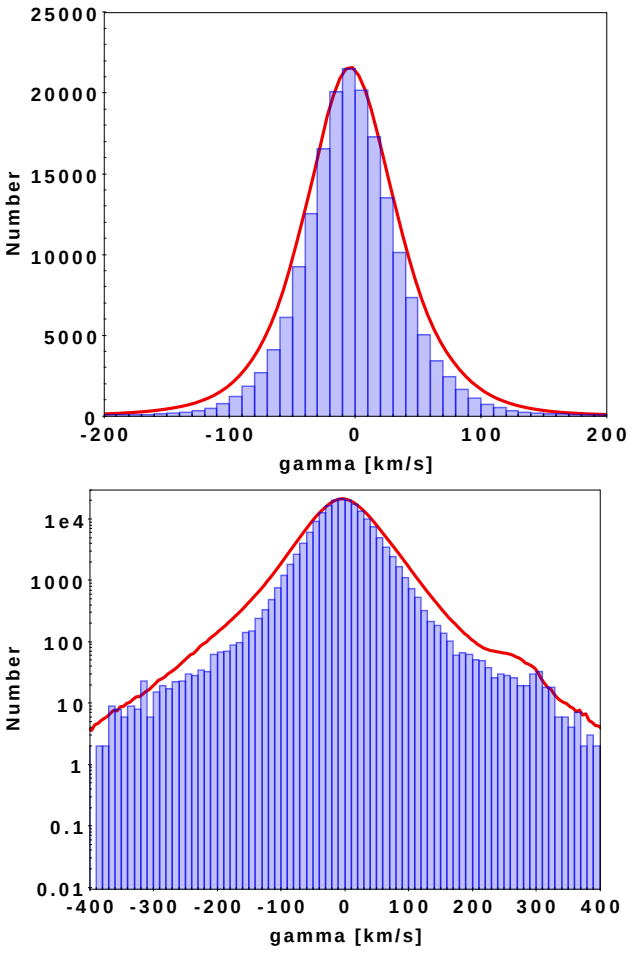

Fig. 21. Histogram (above) and histogram in log (below) of the RVs of the centre of mass γ for all the SB1 objects. The bin width is 10 km s^{-1}. The red curves represent the distribution observed for the RVs of the assumed constant stars (Katz et al. 2023). These red curves have been normalised to the maximum of the histogram.

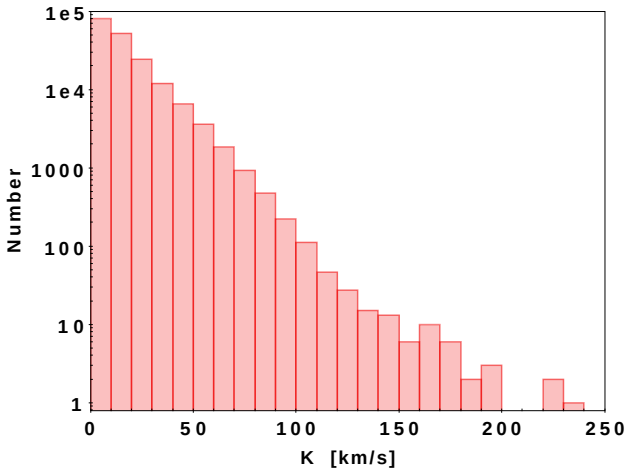

Fig. 22. Histogram of the various semi-amplitudes K (in km s^{-1}). The ordinates are given on a logarithmic scale. The number decreasing with K is to first order linear on this log scale. The bin width is 10 km s^{-1}.

here) for the objects in the SB1C class, the conclusions are quite identical.

Figure 24 presents the distribution of the uncertainty on the semi-amplitude K as a function of the uncertainty on the RV of the centre of mass γ. The ratio of the two is expected to be the square root of 2 (Lucy & Sweeney 1971) which is what is observed despite a rather large dispersion. No particular anomaly is present on the uncertainties of the considered parameters.

The critical point when fitting the orbital SB1 models is the determination of the period (see above). Figure 25 displays the histogram of the period (in log). Several features are visible. The log-distribution is rather flat, with a marked excess at very large periods. This excess can be due to the fact (see above) that the periods of the same order of magnitude as the span of time of the mission are rather indicative and certainly not a perfect estimator of the true period. This uncertainty implies the concentration of several solutions near these long periods. This excess is very

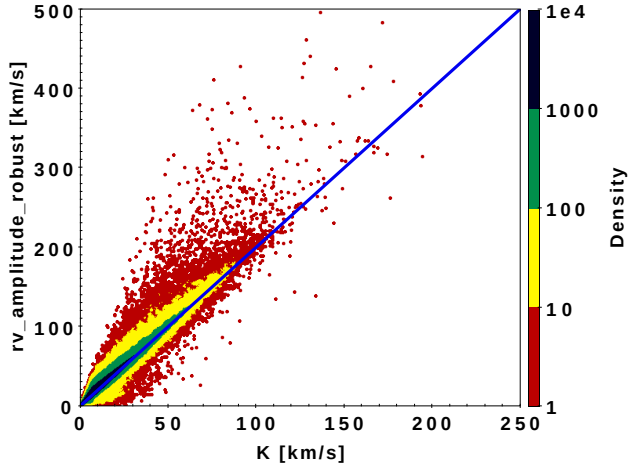

Fig. 23. `rv_amplitude_robust` measured by MTA as a function of the semi-amplitude K (in km s^{-1}). The blue line indicates the expected value (factor of 2) in absence of noise. Due to the large amount of objects (183 434), we preferred to adopt a colour-code scale that underlines the high density of solutions close to this blue line (see on the right).

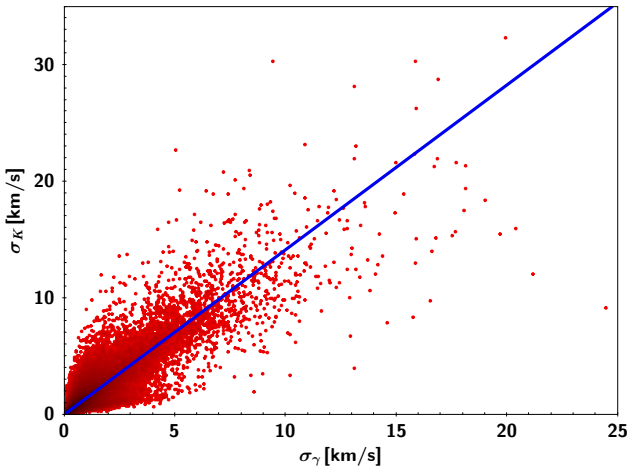

Fig. 24. Uncertainty (one standard deviation) on the semi-amplitude K as a function of the uncertainty on the RV of the centre of mass γ (systemic velocity). Their uncertainties should be in a ratio of the square root of two which is marked by the blue straight line. Although the dispersion is rather large, the ratio law is well respected.

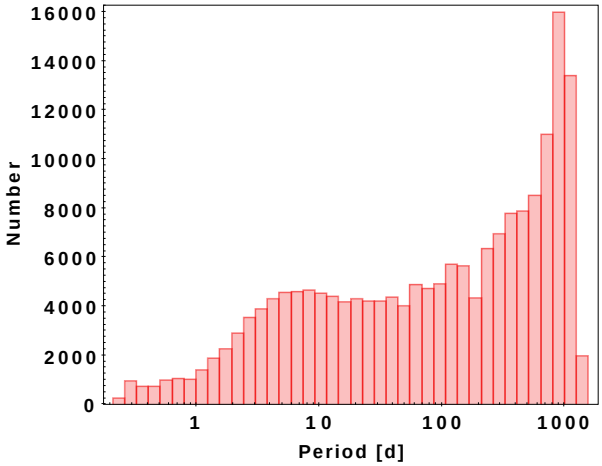

Fig. 25. Histogram of the observed period distribution. We note that the period is on a logarithmic scale.

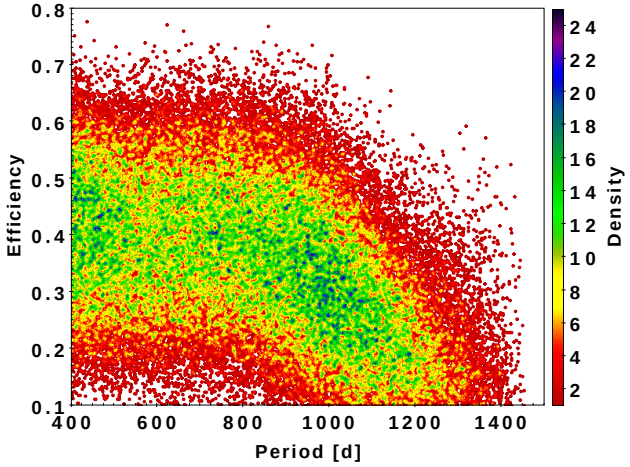

Fig. 26. Solution efficiency as a function of the period. The colour bar represents the density of objects on a linear, arbitrary scale.

sensitive to the threshold separating SB1 and TrendSB1 solutions. Another explanation is the fact that the orbital solutions are enforced in this region. On the other hand, the deficit of solutions around the period of 182.6 days is due to the yearly Earth orbit that the *Gaia* satellite follows by staying at L2. The satellite's avoidance of the direction of the Sun produces regular gaps with a timescale of half a year. These gaps presented by the sampling are associated to large gaps in phase leading to the removal of the solutions because they correspond to badly determined periods and orbits. Figure 27 is showing a diagram of the value of the largest gap in the solution as a function of the period. Exceptionally, here we use the data set of 367 328 objects constituting the set of SB1 solutions before the application of the internal filtering. It is clear from this figure that some of the solutions at periods near 182.6 d generate phase diagrams that are associated to the presence of vast gaps. This situation seems also to be the case, but to a lesser extent, for the period around one year. The phenomenon is perhaps also present for periods of 91.3 d. A detailed look at this Fig. 27 suggests that a conservative choice for the threshold on the gap of 0.3 removed all the potentially spurious or at least unsafe solutions. The decrease of density of solutions for lower periods, as seen in Fig. 25, is an artefact of the application of an ad hoc selection criterion (see the paragraph on the internal filtering of the period in Sect. 7.1) that prevents the proliferation of small periodicities. This effect should be further investigated in the next data releases.

In Fig. 26, we present a diagram showing the efficiency of the solutions as a function of the adopted period. We highlight here the long-period objects. For periods between 200 and 1000 d, the efficiency is essentially distributed between 0.2 and 0.8, independently of the period with a maximum of occurrence around 0.44 (as already reported above). The values of solution efficiency above 0.8 are basically absent, indicating that all the solutions present correlations among the final parameters. These values are part of the catalogue. There are very few solutions with efficiencies between 0.1 and 0.2, and all these solutions could be considered as doubtful, although we preferred to maintain them in the catalogue and to let the user the possibility to reject them or not. Above periods of 1000 d, the typical efficiencies are lessening and a population below 0.2 is appearing. This is a direct consequence of the fact that the associated period is larger

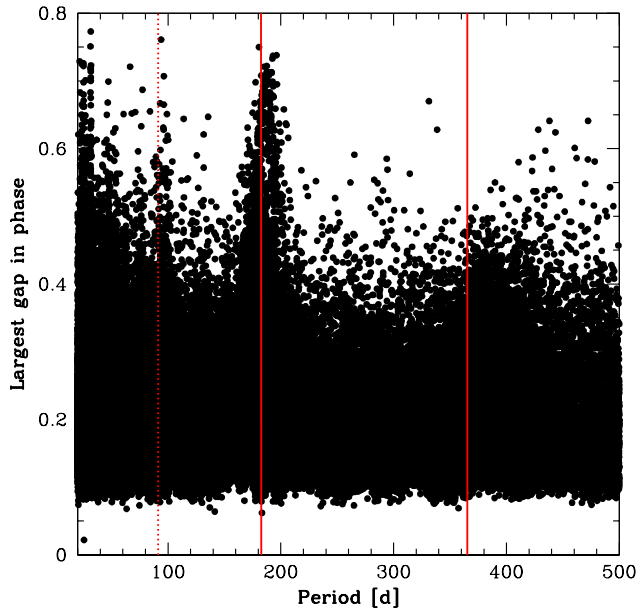

Fig. 27. Largest gap observed in the phase diagram as a function of the period. Periods near 365.25 d and particularly 182.6 d (vertical red lines) are corresponding to the presence of some large gaps. It could also be true for the period of 91.3 d (vertical red dotted line).

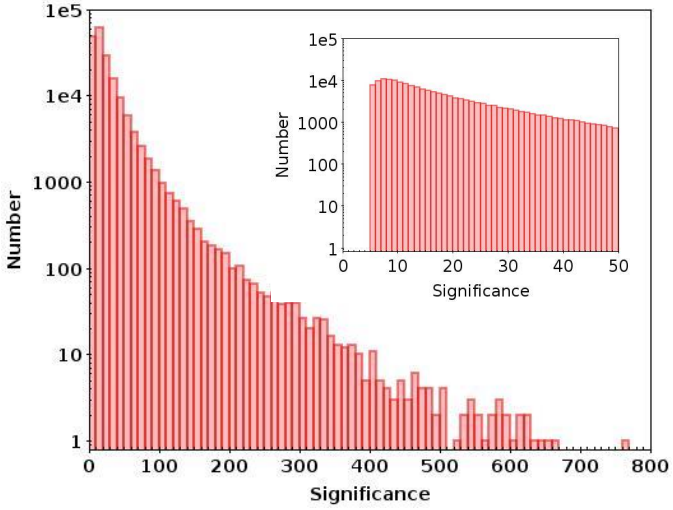

Fig. 28. Histogram of the significance $\frac{K}{\sigma_K}$. The ordinates are on a logarithmic scale; the bin width is 10. No break is seen in the distribution, which prevents us from defining a possible selection criterion. The insert presents a zoom on small significance: the related bin width is here 1.

than ΔT. It is evident that these solutions should not be blindly trusted, but in any case the reported periods can be considered as indicative. Again, the low value of the efficiency unveils existing correlations among the final parameters due to bad phase coverage. As mentioned above, for variabilities at still larger timescales, the trend solution is taking over. However, addressing solutions slightly above ΔT would necessitate a too large degree of the fitted trend polynomial, introducing a lack of robustness of the fit. The orbital solution is thus more informative despite the possible biases in the derived parameters. For future releases, the range of concerned periods will change accordingly with the increasing value of ΔT.

Figure 28 presents the histogram of the spectroscopic significance $\frac{K}{\sigma_K}$. The significance is distributed as a decreasing power

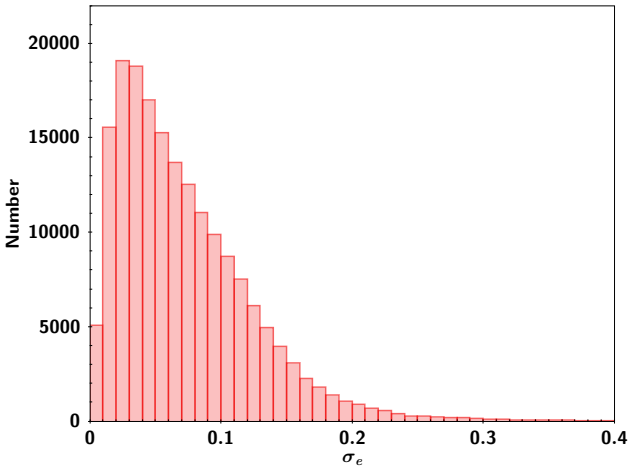

Fig. 29. Histogram of the derived standard deviations associated to the eccentricity. The bin width is 0.01.

law vanishing at 800. The maximum accepted value for K being 250 km s^{-1}, the maximum significance value corresponds to a typical uncertainty of about 0.3 km s^{-1} which is coherent with expectation. The distribution shows no inflection locus that could be useful in setting a selection threshold. In this context, we decided to introduce a threshold at 5 but not more restrictive ones. Below a value of 5, the corresponding solutions are clearly not trustworthy. We leave the possibility to the catalogue users to define a threshold that could be appropriate for their own work.

Figure 29 exhibits the distribution of the standard deviations on the eccentricity. The typical uncertainty on the eccentricity is of $\sigma \sim 0.03$–0.10. It is necessary to draw the attention to the fact that with such an uncertainty a derived eccentricity of 0.3–0.4 is still compatible with a circular orbit (this is within three σ_e keeping in mind that the distribution of eccentricities is not Gaussian). RVs are well-known to be a poor source of determination of eccentricities in contrast to the astrometric orbital solutions. Errors larger than 0.35–0.4 (there are a few) could be associated to insecure adjustments and we encourage the catalogue user to be careful with these questionable solutions.

As an example, Fig. 30 shows a plot of the eccentricities as a function of the corresponding period. The plot of all the objects (in red) includes short period solutions with some eccentricity up to 0.4. As mentioned just above, these values of e up to 0.3–0.4 are not particularly anomalous. However, the general aspect of the diagram is not the expected one in terms of the theory of circularisation of the orbits (close binaries are not expected to remain eccentric for long). If we restrict the data set to values with a significance larger than 30, we obtain the distribution in blue where the expected effects of the circularisation start to be well visible. A restriction to the significance values above 80 generates the distribution in orange. Another study of this topic can be found in Gaia Collaboration (2023a). Evidently, the impact of the selection on the basis of the significance is a delicate topic and we advise the greatest caution for the selection of the threshold to be applied. This remark can also be extended to other quality factors than just the significance.

Finally, Fig. 31 represents the plot of the derived SB1 mass function as a function of the adopted period. We can clearly see that the majority of the data points are between 0.001 and 0.5 $M_\odot$. Objects outside this range are very few, in particular on the side of the high mass function. From the statistical point

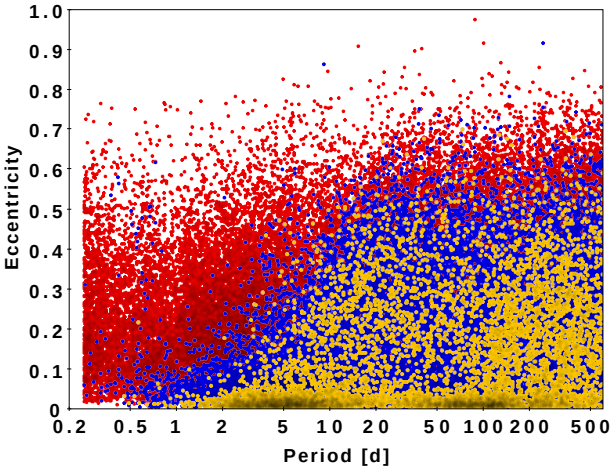

Fig. 30. Eccentricity as a function of the period. The red points correspond to the full sample (significance larger than 5). The blue (respectively orange) points correspond to a significance larger than 30 (respectively 80). The blue points already reveal the likely circularisation phenomenon that is expected to be present.

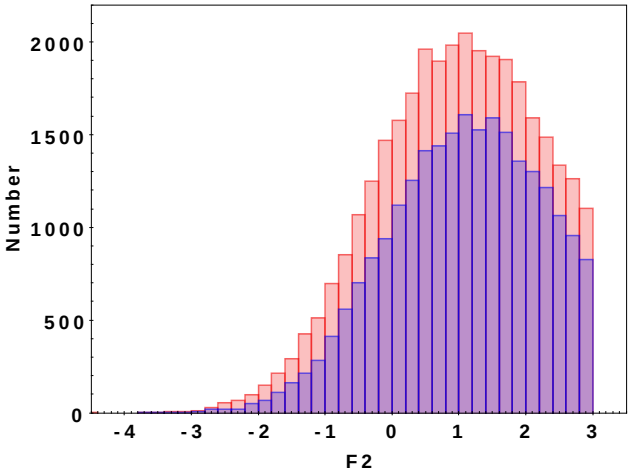

Fig. 32. Histogram of the goodness of fit F_2 for the trend fits. The bin width is 0.2. In blue the solutions are associated to the `FirstDegreeTrendSB1` class and in red with the `SecondDegreeTrendSB1` one. The cut-off above 3 is outstanding. The centring of the distribution is rather located near 1 instead of 0.

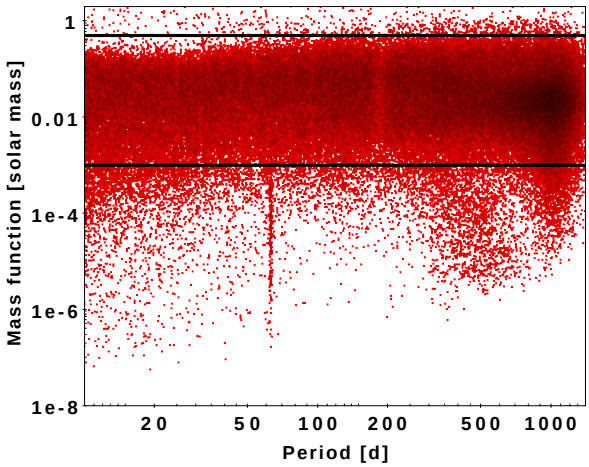

Fig. 31. Distribution of the observed SB1 mass function ($M_\odot$ in log scale) as a function of the period (days in log scale). Guiding black lines are given and correspond to 0.5 $M_\odot$ and 0.001 $M_\odot$. Clearly the majority of the objects is in-between the two lines, the maximum populated area is just half-way (in log) between the two lines.

of view, the diagram is rather close to the expected one. Data points at an unexpected location could be due to problems in the derivation of the period or of the eccentricity. In particular, we should be very careful while searching for abnormal mass function values or massive black holes. In addition, the overdensities tracing vertical lines at $P = 62.97$ days and at one year (much broader) are due to objects with a varying apparent position inducing systematic effects on the RVs associated to the variation of the scan-angle. These are objects presenting small amplitudes at these particular periodicities (bad astrometry inducing a wrong wavelength calibration). These objects could easily be filtered out but they were not before the application of the combiner with the astrometric channel in order not to miss interesting objects. The effect in astrometry is described in Lindegren (2019). A more detailed analysis of the scan-angle dependent signals is given in Holl et al. (2023a).

9.2.2. `FirstDegreeTrendSB1` and `SecondDegreeTrendSB1` classes

Concerning trend solutions, we show in Fig. 32 the distribution of the goodness of fit F_2 for all the TrendSB1 solutions. It appears that the expected distribution that should be centred on zero, is actually rather centred on 1. This points out to an anomaly of this distribution, tending to suggest that the fits are not perfect although still good. The reason of this shift is for the moment unknown; to understand it is very important because we should study if the cut-off at 3 is still justified in the context of the observed shift of the distribution. The characteristics of the distribution is similar for the two subtypes of solutions.

The trend is certainly not a purely physical model, since it is basically induced by the shortness of the time span of the observations compared to the true cycle of the star considered. Statistical decisions are adopted to select a model (orbital, linear trend, high-order trend) but it seems unlikely that the adopted thresholds could induce the shift of the F_2 distribution in such a global manner. Another explanation for the shift could be the presence of some intrinsic variability additional to the trend signal, as already evoked for orbital solutions. Again, it is very unlikely that this could affect such a large percentage of the objects. Thus, the last possibility and most likely hypothesis is that the individual uncertainties associated to RVs are underestimated. The trend models are certainly a good way to investigate this possibility. Indeed, the linear trend is a very simple model which is linear in the parameters and the fit should be robust and unambiguous at least compared to the non-linear orbital model whose behaviour could be more complex (see Sect. 9.2.1). The overfitting is here not a concern, and the number of degrees of freedom of the χ^2 of the fit is precisely established. We thus further investigated the global distribution (over all the objects) of the observed minus calculated residuals normalised by the associated individual uncertainties. To be strictly conservative, we only considered objects with an astrometric `ruwe` less than 1.4 for which no additional contribution to the uncertainties was applied; only the values delivered by CU6 are involved. The distribution of the normalised residuals follows very closely a Gaussian probability density function, centred on zero, but with a standard deviation larger than $\sigma = 1$: actually

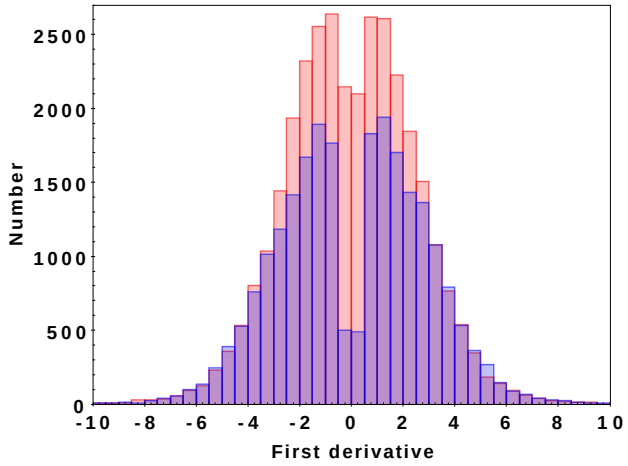

Fig. 33. Histogram of the first derivative normalised by its derived uncertainty for the trend fits. In blue the solutions are associated to the `FirstDegreeTrendSB1` class and in red with the `SecondDegreeTrendSB1` one. The bin width is 0.5. The central dip is due to fits where the first derivative is too small and thus to objects considered as constant in RVs.

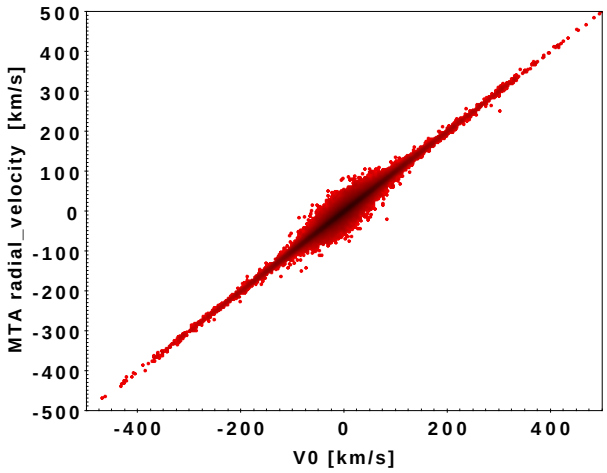

Fig. 34. Correlation between the V_0 derived in our pipeline and the mean RV as computed by MTA assuming the objects are constant.

we deduced $\sigma = 1.08$. Therefore, an increase by 8 percent of the uncertainties could solve the reported anomaly of the F_2 distribution except for a small effect that could be explained by some additional variability. This statement needs further confirmation.

Figure 33 illustrates the distribution of the first derivative normalised by its uncertainty. The distribution is very similar for the two classes and presents an almost perfect symmetry, as expected. The deficit of solutions around the zero derivative is due to the fact that these solutions are not included because they correspond to constant objects. A lot of the solutions are well above $\pm 3\,\sigma$. A few well-defined derivatives are further away up to $10\,\sigma$.

Finally, we try to validate the values of the derived V_0. They are plotted in Fig. 34 with in ordinates the mean velocity computed by MTA. We observe that the two parameters are very well correlated, providing further support to the V_0 values we are delivering.

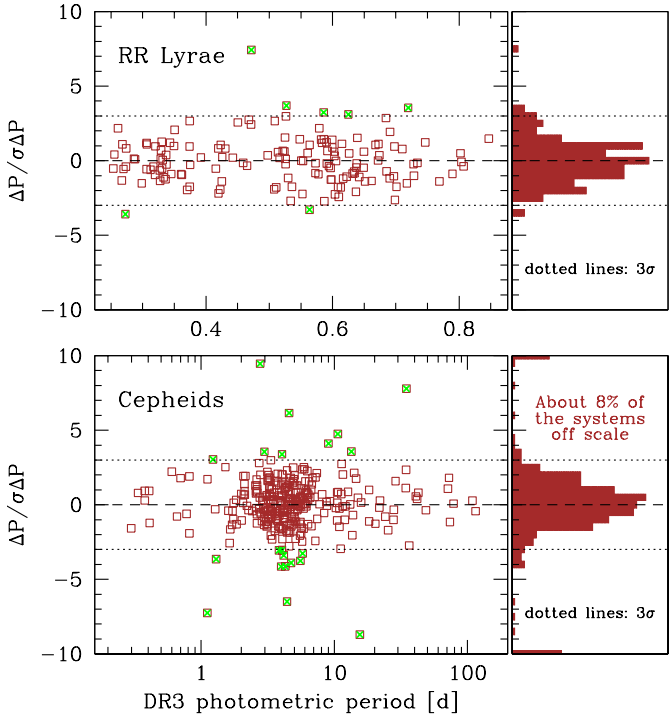

Fig. 35. Deviations in P expressed in σ units for the RR Lyrae (top panels) and Cepheid (bottom panels) variables excluded from the SB subcatalogue (see post-filtering in Sect. 10.1.1), as a function of the photometric period. Objects for which the reference period is not recovered to within 3σ are flagged with green crosses. We note that a small fraction of the Cepheids are off scale.

9.3. Validation against CU7 photometric periods

Some intrinsic variables can exhibit RV variations that could be mistaken as the effect of an orbital motion. As explained in Sect. 10.1.1, a sample of securely identified Cepheids (333, Ripepi et al. 2023) and RR Lyrae (147, Clementini et al. 2023) have also an SB1 solution. They should be eventually excluded from our SB subcatalogue but they also offer the unique opportunity to have an independent test of the spectroscopic periods by comparing them to the photometric ones. Indeed, we compared our spectroscopic periods to those determined from DR3 photometric data for the 147 RR Lyrae and the 333 Cepheids. The difference is expressed as [P(spectroscopy)–P(photometry)]/σ, where σ is the square root of the sum in quadrature of both uncertainties. We find a recovery success rate to within 3σ of ~95% for the RR Lyrae's. It is not only lower for the Cepheids (~85%), but the outliers are also much more severely discrepant (Fig. 35). We find that the solution efficiency is better for the RR Lyrae's than for the Cepheids.

9.4. Validation against other data sets

We present below a comparison between the DR3 orbital parameters and those from external catalogues, either relying on spectroscopic (Griffin, SB9, and APOGEE) or photometric (WISE and ASAS-SN) data. The most distinctive feature between the various validation samples is the widely different range of orbital periods that span about five orders of magnitude (see Fig. 36). The WISE and ASAS-SN catalogues are made up of candidate short-period binaries that are useful for assessing the reliability of our orbital solutions in this regime. We consider exactly the same DR3 data set as discussed in Sect. 9.2. In the following,

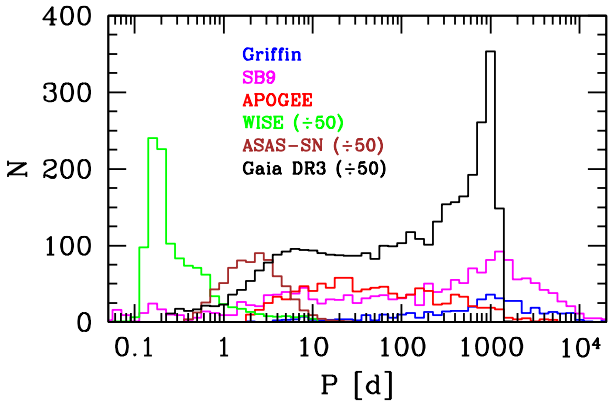

Fig. 36. Breakdown of the orbital periods for the DR3 (black), Griffin (blue), SB9 (magenta), APOGEE (red), WISE (green), and ASAS-SN (brown) binary samples. Only SB1s are shown for the first four catalogues that are based on spectroscopic data. The counts for DR3, WISE, and ASAS-SN are divided by 50.

except for ASAS-SN, literature results without an uncertainty given are ignored. The difference for a given orbital parameter is expressed as [value(NSS)–value(reference)]/σ, where σ is the square root of the sum in quadrature of both uncertainties.

9.4.1. Comparison against Griffin's results

We first considered the orbits determined in a series of papers (not only in The Observatory) by Roger Griffin and collaborators. The orbital elements were taken from the 2 March 2021 version of the SB9 catalogue[7] (original paper: Pourbaix et al. 2004), which is made up of 4021 systems. The association to the relevant DR3 ID was simply tracked from the DR2 alias available in SIMBAD[8] at that time. Since for validation we needed to be confident in our reference catalogue, we remained conservative at this level, avoiding adverse and insecure cross-matching. This procedure can now certainly be much improved. To avoid any ambiguities in this cross-match, or a possible bias in centre-of-mass velocity, γ, visual/multiple systems were excluded. They were identified as having a `component` field in SB9 not empty or to have a SB9 ID entry in the 23 June 2020 version of the Multiple Star Catalog (MSC)[9]. This validation catalogue contains the orbits of 414 SB1's and 101 SB2's.

The breakdown of DR3 spectroscopic solution types is as follows for Griffin's SB1 catalogue:

- `SB1: 83 systems,`
- `SB1C: 0 system,`
- `TrendSB1: 10 systems,`
- `SB2: 1 system,`
- `SB2C: 0 system,`
- `Without any DR3 spectroscopic deterministic solutions: 320 systems.`

The period of the ten systems with a trend solution is much longer than the time span of the RVS observations, which is as expected. Five and one systems among the SB2 validation catalogue have received a `SB1` and `TrendSB1` solution, respectively.

Figures F.1, F.2, F.3, and F.4 show the comparison between the DR3 and Griffin's SB1 orbital parameters, as a function of various DR3 parameters. The deviations in P expressed in σ units are shown in Fig. 37, as a function of the reference period. Systems for which the latter is difficult to find because of the limited length of the RVS observations (i.e. ΔT = time span observations < P) are indicated with a different colour. Systems for which this period is not recovered to within 3σ are also flagged.

Some deviating points can be noticed in Figs. F.1, F.2, F.3, and F.4, especially when the period is not recovered. The incidence of outliers is quantified below. In some cases, the agreement is satisfactory in absolute terms and the deviation is only due to (very) small uncertainties. There is clear evidence for a systematic γ offset, although the discrepancy only amounts to $\sim$1 km s^{-1} on average. It might be due to the fact that the RV values adopted by Griffin for his four reference stars back in the late sixties (Griffin 1969) appear to be larger by roughly the same amount than those in the recent literature, including Soubiran et al. (2018) for λ Lyr (Fig. 38).

Table 12 summarises the recovery rate of the SB1 parameters to within 3σ. Cases where the time span of the RVS observations is longer/shorter than the reference period are listed separately. As may be expected, the performance of the pipeline greatly improves when the observations cover more than one orbital cycle.

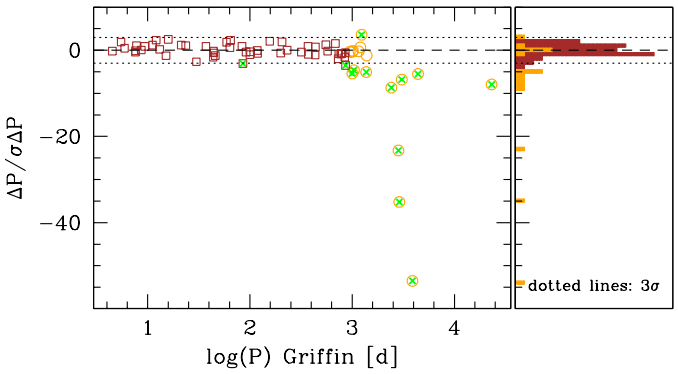

Fig. 37. Deviations in P expressed in σ units with respect to Griffin, as a function of the reference period. Systems for which the period, P, in the literature is unlikely to be found (i.e. $\Delta T < P$) are indicated with orange circles. Systems for which $\Delta T \geqslant P$ are shown as brown squares. Binaries for which the reference period is not recovered to within 3σ are flagged with green crosses.

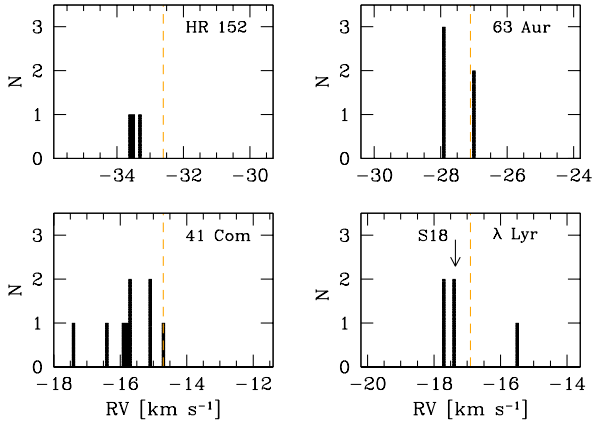

Fig. 38. Comparison for the reference stars adopted by Griffin (1969) between his RVs (dashed lines) and those in the more recent literature (histograms). The value of Soubiran et al. (2018) for λ Lyr is indicated.

[7] https://sb9.astro.ulb.ac.be/
[8] https://simbad.u-strasbg.fr/simbad/
[9] http://www.ctio.noirlab.edu/~atokovin/stars/

Table 12. Recovery rate to within 3σ of each parameter with respect to Griffin expressed in per cent.

	e	γ	K	P
$\Delta T \geqslant P$	96.4 (55)	12.0 (50)	94.1 (51)	96.1 (51)
$\Delta T < P$	88.2 (17)	17.6 (17)	58.8 (17)	35.3 (17)
All	94.4 (72)	13.4 (67)	85.3 (68)	80.9 (68)

Notes. The number in brackets is the number of systems the estimate is based on. Cases where the time span of the observations is longer/shorter than the reference period are listed separately.

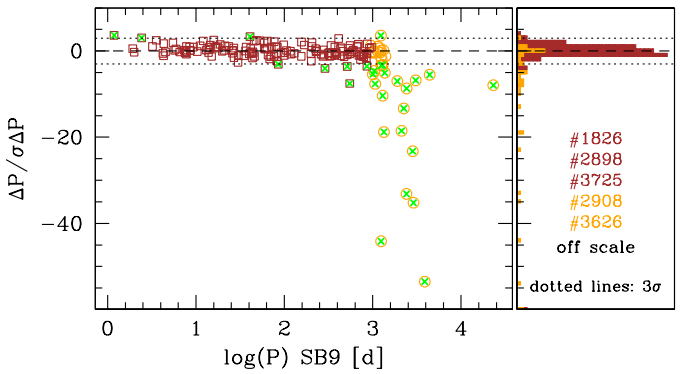

Fig. 39. Same as Fig. 37, but for the SB9. As indicated, a few systems have very negative $\Delta P/\sigma\Delta P$ values and are off scale. The SB9 IDs are given.

9.4.2. Comparison against SB9 catalogue

As a next step, the analysis above was repeated for the whole SB9. Although the sample is now considerably larger, it is much less homogeneous and contains solutions of generally lower quality. The most recent orbit was selected when a given system had multiple entries in the catalogue. As described above, visual and/or multiple systems were also excluded. Six SB2 orbits were replaced by those determined by Lester (2020). This validation catalogue contains the orbits of 1750 SB1's and 647 SB2's. The breakdown of DR3 spectroscopic solution types is as follows for SB1s in the SB9 catalogue:

- `SB1: 226 systems,`
- `SB1C: 0 system,`
- `TrendSB1: 45 systems,`
- `SB2: 3 systems,`
- `SB2C: 0 system,`
- `Without any DR3 spectroscopic deterministic solutions: 1476 systems.`

All (but one) of the 45 systems with a trend solution have very long periods. Furthermore, 18 and 2 systems in the SB2 validation catalogue have received a SB1 and a TrendSB1 solution, respectively.

Figures F.5, F.6, F.7, and F.8 show the comparison between the orbital parameters. The deviations in P expressed in σ units are shown in Fig. 39, as a function of the reference period.

The results are consistent with those obtained for Griffin's sample. Some features are expected. For instance, the difficulty to find the correct period for (very) wide, low-amplitude systems: the orbital periods are generally systematically underestimated (Fig. F.8). Once again, the large discrepancies may be misleading in some cases. For instance, some short-period systems are outliers in P because the uncertainties are very small ($\sim 10^{-4}$ d).

Table 13. Same as Table 12, but for the whole SB9.

	e	γ	K	P
$\Delta T \geqslant P$	89.2 (167)	53.9 (152)	90.4 (156)	92.7 (150)
$\Delta T < P$	81.1 (37)	43.2 (37)	64.9 (37)	37.8 (37)
All	87.7 (204)	51.9 (189)	85.4 (193)	81.8 (187)

The same conclusion holds for SB2's (see Damerdji et al., in prep.). The slight systematic discrepancy and large dispersion for γ is likely due to the fact that a significant fraction of the orbits are actually taken from Griffin (see discussion about offset in Sect. 9.4.1). The γ deviations show some evidence for a bimodal distribution (Fig. F.6), and it indeed appears that the smaller peak is largely dominated by results from Griffin. The RV bias of up to 75 m s^{-1} affecting the faint end (Sect. 2) is negligible in this respect. Furthermore, there are very few objects with $G_{\rm RVS}^{\rm int} >$ 11 mag (Fig. F.6).

Table 13 summarises the recovery rate of the SB1 parameters to within 3σ. Once again, not surprisingly, the incidence of outliers decreases when the time span of the observations exceeds one orbital cycle.

9.4.3. Comparison against APOGEE

Orbital parameters for ~20 000 SB1's have been released as part of APO Galactic Evolution Experiment (APOGEE)[10] DR16 (Price-Whelan et al. 2020, hereafter PW20). Our validation sample is their 'Gold Sample' that contains 1032 systems with high-quality, unimodal posterior samplings.

The breakdown of DR3 spectroscopic solution types is as follows for the SB1s in PW20:

- `SB1: 126 systems,`
- `SB1C: 0 system,`
- `TrendSB1: 4 systems,`
- `SB2: 3 systems,`
- `SB2C: 0 system,`
- `Without any DR3 spectroscopic deterministic solutions: 899 systems.`

All the systems with a trend solution have a period much longer than ΔT.

Figures F.9, F.10, F.11, and F.12 show the comparison between the orbital parameters. The deviations in P expressed in σ units are shown in Fig. 40, as a function of the reference period. For the period, two distinct situations are encountered: either there is a reasonable correspondence or, in about ~25% of the cases, the values are dramatically discrepant. For those, systematic differences are not observed. The solutions in PW20 for the majority of these binaries rely on much fewer RVs (less than ten). It is likely that the paucity of measurements prevents a robust determination of the orbital period. We caution that these systems are off scale in Figs. 40 and F.12.

Table 14 summarises the recovery rate of the SB1 parameters to within 3σ. The lower incidence of outliers when a full orbital cycle is not covered may be attributed to small number statistics.

9.4.4. Comparison against WISE

A total of ~56 000 short-period (<10 d) binary systems observed by the Wide-field Infrared Survey Explorer (WISE)[11] have

[10] http://www.sdss3.org/surveys/apogee.php
[11] https://www.jpl.nasa.gov/missions/wide-field-infrared-survey-explorer-wise

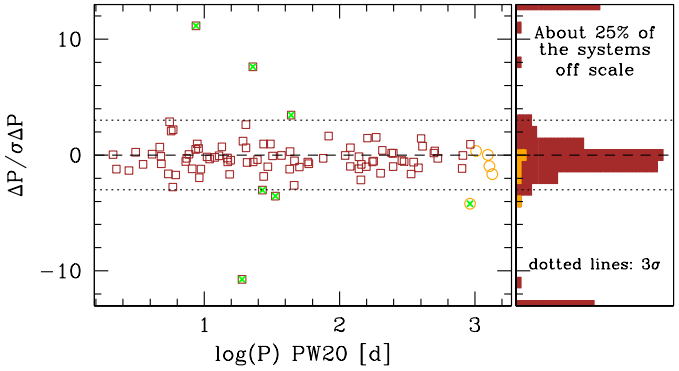

Fig. 40. Same as Fig. 37, but for the APOGEE sample. We note that a significant number of systems are off scale.

Table 14. Same as Table 12, but for the APOGEE gold sample of PW20.

	e	γ	K	P
$\Delta T \geqslant P$	87.6 (121)	73.6 (121)	85.1 (121)	69.4 (121)
$\Delta T < P$	100.0 (5)	100.0 (5)	100.0 (5)	80.0 (5)
All	88.1 (126)	74.6 (126)	85.7 (126)	69.8 (126)

been discussed by Petrosky et al. (2021). Most of them are (near-)contact EBs. However, as will be shown below, a non-negligible amount are classical pulsators, especially RR Lyrae stars. The infrared photometric measurements are made through the W1 filter centred at about 3.4 µm. A subset of 49 465 systems can be unambiguously associated to a DR3 source[12]. This sample allows us to test the reliability of our orbital periods for very close binaries, although other orbital parameters are not available.

The breakdown of DR3 spectroscopic solution types is as follows for the SB1s in Petrosky et al. (2021):

- `SB1: 704 systems,`
- `SB1C: 3 systems,`
- `TrendSB1: 12 systems,`
- `SB2: 69 systems,`
- `SB2C: 23 systems,`
- `Without any DR3 spectroscopic deterministic solutions: 48 654 systems.`

We find a close correspondence in most cases with either the main WISE period or its second harmonic (Fig. 41). As discussed by Petrosky et al. (2021), the orbital period they found for the binaries is usually half the true value because of the nearly symmetrical nature of the WISE light curves. The WISE period is recovered to within 1% for 25.3% and 55.4% of the cases for the first and second harmonic, respectively. The overall recovery rate at this level is thus about 81%. It remains quite high for lower tolerances (e.g. ∼78% for 0.1%). The deviations lie in the vast majority of cases within five times the mutual uncertainties (Fig. 42).

Most of the WISE sources were assigned a classification in DR3 based on their light variations (Mowlavi et al. 2023; Ripepi et al. 2023; Clementini et al. 2023). As can be seen in Fig. 43, a (nearly) circular solution is found for the majority of the EBs, as may be expected for such close binaries. In contrast, the systems with high eccentricities are mostly classical pulsators (Cepheids or RR Lyrae) mistaken as binaries. As described in

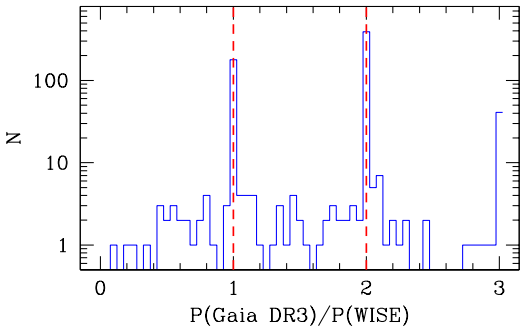

Fig. 41. Ratio between spectroscopic DR3 and photometric WISE periods. DR3 values that can be identified with the first and second WISE harmonics are indicated as vertical, dashed lines.

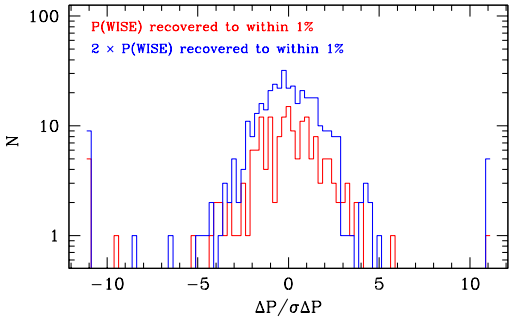

Fig. 42. Deviations in σ units between the DR3 and WISE periods when the latter (or its harmonic) are recovered to within 1%.

Sect. 10.1.1, those pulsating variables were rejected from the final DR3 catalogue during the post-filtering process.

9.4.5. Comparison against ASAS-SN

The physical parameters of 35 464 detached eclipsing binaries (DEBs) observed in the framework of the All-Sky Automated Survey for Supernovae (ASAS-SN; Kochanek et al. 2017)[13] have been determined by Rowan et al. (2022). A total of 35 307 systems have a *Gaia* DR3 ID. The data are based on the analysis of g- and V-band light curves. The orbital periods nicely bridge the gap between the WISE sample of close binaries and other more general catalogues, such as APOGEE or the SB9 (Fig. 36). Contrary to WISE, other orbital elements (e.g. eccentricity) are provided besides the period, although without any uncertainties. For P, we assume a typical relative uncertainty of 10^{-6} (e.g. Holanda et al. 2018).

The breakdown of DR3 spectroscopic solution types is as follows for the binaries in Rowan et al. (2022):

- `SB1: 626 systems,`
- `SB1C: 3 systems,`
- `TrendSB1: 6 systems,`
- `SB2: 92 systems,`
- `SB2C: 35 systems,`
- `Without any DR3 spectroscopic deterministic solutions: 34 545 systems.`

Most objects have $G \gtrsim 13$ mag and are fainter than our magnitude cut-off.

The majority of the 626 objects are classified in DR3 as EBs based on their photometric behaviour (Mowlavi et al. 2023). Conversely, none are identified as classical pulsators

[12] http://cdn.gea.esac.esa.int/Gaia/gedr3/cross_match/

[13] https://www.astronomy.ohio-state.edu/asassn/

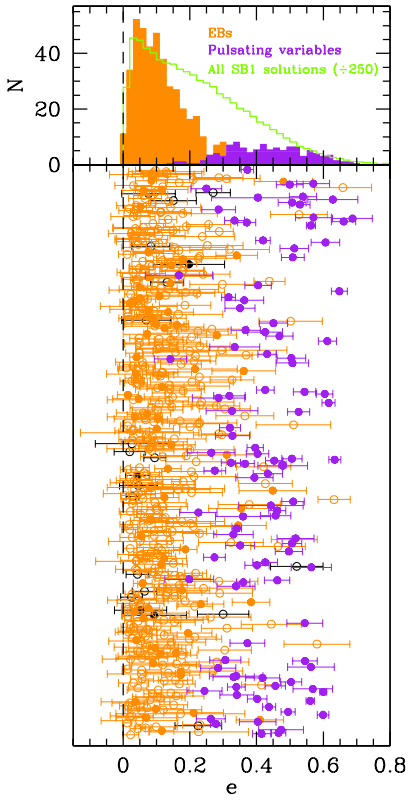

Fig. 43. Distribution of eccentricities for the WISE sources for which the period (filled symbols) or its harmonic (open symbols) is recovered to within 1%. Lower panel: eccentricities for the systems classified as EBs (orange; Mowlavi et al. 2023), Cepheids or RR Lyrae (purple; Ripepi et al. 2023; Clementini et al. 2023). Those without a classification are shown in black. The systems are displayed in no particular order along the y-axis. Upper panel: histogram of the distributions for the EBs, pulsating variables, and all SB1 solutions (the last one is divided by 250).

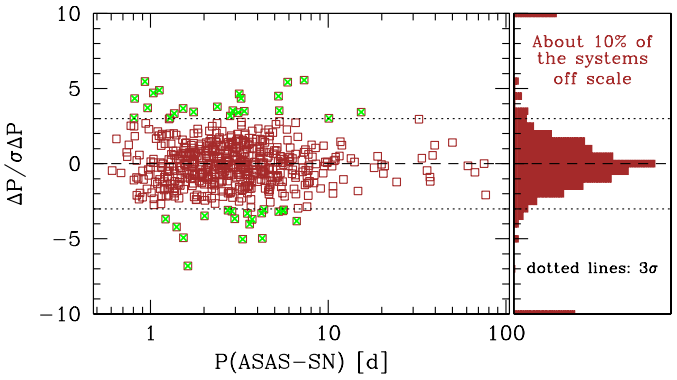

Fig. 44. Same as Fig. 37, but for the ASAS-SN sample. We note that a significant number of systems are off scale.

(Clementini et al. 2023; Ripepi et al. 2023). For this sample with SB1-type solutions, the overall recovery rate of the photometric period to within 3σ amounts to ~83% (Fig. 44). This estimate includes one system where the ASAS-SN period is actually the harmonic (as discussed for WISE above) or for nine objects twice the true value because one of the two eclipses is too shallow to be detected (see figure 6 of Rowan et al. 2022 for an example). Visual inspection of the ASAS-SN light curves for the nine DEBs supports the latter explanation.

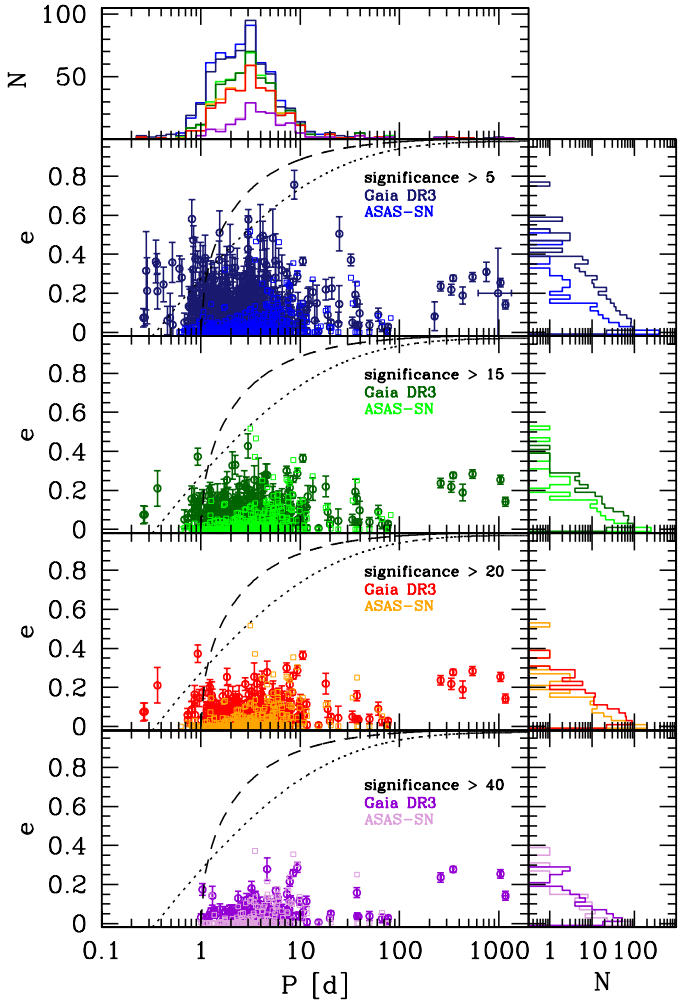

Fig. 45. e-log P diagrams for the 626 systems in common between DR3 and ASAS-SN. Each panel corresponds to a different threshold on the solution significance. The dotted and dashed lines show the upper envelopes defined by Mazeh (2008) and Halbwachs et al. (2005), respectively. A cut-off period of one day was adopted for the latter. The position of the curve is only illustrative as this choice is somewhat arbitrary.

The relevance of filtering out short-period systems with a low solution significance was discussed in Sect. 9.2. We show in Fig. 45 the DR3 and ASAS-SN data in a e-log P diagram for various cuts in solution significance (see also Fig. 30). The upper envelopes defined by Mazeh (2008) and Halbwachs et al. (2005, modified to $P_{\rm cut-off}$ at 1 d) are overplotted. The former is based on an empirical fit of the SB9 data, but the latter rests on theoretical arguments. While P is recovered in ~67% of the cases for K/σ_K values up to 15, the proportion sharply increases for higher significances (success rate of ~90% on average). The differences are generally small in absolute terms, which explains the close match in the period distributions shown in Fig. 45. However, an excess of seemingly eccentric systems with respect to Rowan et al. (2022) is readily visible. The eccentricity distributions become more and more similar when the significance cut-off increases until they become statistically indistinguishable above about 50. Rather than a poor determination of the period for close binaries, the issue thus seems to rather lie in solutions that are too eccentric when they are ill-constrained (see also Jayasinghe et al. 2023).

9.4.6. Summary

As a preamble to summarising the main conclusions of our analysis, it is worth emphasising that some of the validation data sets discussed above are of limited usefulness in the sense that they do not offer stringent tests of our results. In particular, the APOGEE orbits from PW20 are based on fewer epochs (about 14 on average concerning the 126 objects) than most NSS sources (see Fig. 14). As a consequence, their orbital solutions are not necessarily more robust and precise. We indeed find a poorer match for the periods with respect to the other validation catalogues, even at high significance. It might be no coincidence that, except for γ, the highest recovery rates for systems monitored over a sufficiently long baseline are found for the high-quality Griffin solutions (Table 12).

As discussed above, the success rate for the main orbital parameters (P, K, e, and γ) is usually well above 80% for the whole set of validation catalogues considered here (Griffin, SB9, APOGEE, WISE, and ASAS-SN). As expected, the performance increases for relatively short-period systems for which at least one orbital period is covered by the observations (about 2.5 years for the RVS spectroscopic processing data set considered here). In this case, the recovery rate commonly exceeds 85–90%. The WISE sample of (near-)contact binaries allows us to investigate the performance of our period search at high frequencies. We recover the WISE photometric periods at a level of about 80–85%. It increases to about 90% for the DEBs with wider orbits discussed by Rowan et al. (2022). We can also note the excellent success rate for the pulsation periods of the RR Lyrae stars that are below one day (Sect. 9.3). The good agreement on the period with both EB data sets advocates for a deeper study using these data. Their combination will certainly result in an extensive analysis that could reveal various kinds of objects such as EW binaries and/or ellipsoidal variables. Such a task is beyond the scope of the present paper. Our analysis globally suggests that the orbital periods below about 10 days are generally robust provided the solution significance is not too low. We find a roughly constant success rate at the level of about 90–95% once the significance exceeds 15–20. However, although they can often be regarded as compatible with zero within the uncertainties, it appears that the eccentricities may be overestimated even if the significance is larger.

We ascribe the poor correspondence between the DR3 and Griffin/SB9 γ's to a slight zero-point offset at the $\sim$1 km s^{-1} level. There is indeed no evidence for a significant bias with respect to APOGEE (Fig. F.10). For the other parameters, no obvious systematic discrepancies are observed.

In order to obtain an order-of-magnitude estimate of the period whenever possible, a choice is made to systematically favour an `SB1` solution over a `TrendSB1` one when the period is estimated to be in the range 0.5 to 1.5 ΔT. As a result, an `SB1` solution could have been assigned despite the fact that the length of the observations is insufficient. However, although the proportion strongly depends on the validation data set, it concerns less than 15% of the cases. Conversely, overall about 79% of the `TrendSB1` solutions are associated to binaries with periods significantly longer than the observations.

Our ability to recover the orbital parameters quite significantly hinges on the retrieval of the true period. When only counting the binaries for which this condition is satisfied in the reference catalogues discussed in Sects. 9.4.1 to 9.4.3, we find an increase of $\sim$5–10% in the recovery rate of e, γ, and K. For e and K, it typically reaches $\sim$95%. It is evident that other limitations, either related to the characteristics of the observations (e.g. number of RVS transits) or the very nature of the binary system (e.g. velocity semi-amplitude) can also potentially play a significant role in the quality of the solutions. Some clues about the sensitivity (or lack thereof) to the combination of some parameters can be gauged from plots such as Fig. 46. Besides suggesting a lack of dependency between the recovery rate of the period and some quantities (e.g. `ruwe`), it illustrates, as discussed above, the difficulty in successfully finding the period at low significance values. However, there may be more complex dependencies lurking in the data that cannot be easily and reliably quantified given the quite limited number of cross-matches with the validation catalogues (a few hundreds systems at most).

In this respect, the fact that only a relatively small fraction of the SB1s in the validation catalogues have a spectroscopic deterministic DR3 solution is the consequence of the various filtering steps, both at CU4 and CU6 levels, described previously. In addition, we note that for the SB9 and Griffin's sample, only a subset of the reference catalogue was used. The case of the SB9, and the steps that led to the relatively small overlap with respect to DR3 (see Sect. 9.4.2) are discussed in detail in Appendix G. Similar arguments apply to the other external catalogues and explain the limited number of benchmark binaries with a spectroscopic deterministic solution in DR3, although the proportion of systems lost at every filtering step may markedly vary. As an illustration, only the very brightest ASAS-SN eclipsing binaries get a DR3 solution. While the $T_{\rm eff}$ range is not an issue, we find that $\sim$90% have a (combined) G magnitude fainter than 12 mag. Similar figures apply to the WISE sample. Regardless of these intrinsic limitations, it is naturally expected that the completeness level will substantially improve in DR4 thanks to the larger amount of RVs collected and improvements in data analysis.

We finally note that a few systems securely reported as double-lined in the literature have received an SB1-like solution (either eccentric or trend). The causes for this are multiple. For instance, the spectral signature of the secondary could be too weak to be detected by *Gaia* or the lines in the composite spectrum cannot be fully deblended at the moderate resolution of the RVS, as discussed in Sect. 8.2.3. The various cases which can be encountered are discussed in Damerdji et al. (in prep.).

10. Additional considerations

10.1. Post-filtering

The results of the sole NSS-SB1 processing (see Sect. 8) are providing a plethora of orbital solutions covering a large domain of the values for the orbital parameters. The pipeline nature of the processing necessarily implies that part of the solutions could be wrong or spurious (whatever could go wrong will go wrong at least once). Therefore, we scrutinised a subset of the results in order to detect surviving problems. Since the DPAC policy with respect to the final production of the DR3 catalogue is intended to favour purity of the list rather than completeness, it is necessary to blacklist solutions that are most certainly spurious. We call this procedure post-filtering and it intervenes between the internally filtered output of the pipeline and the publication of the catalogue. This post-filtering can be based on information coming from outside our internal validation. A clear flaw of this procedure is that this analysis will never identify all the pitfalls. Conversely, some objects could be blacklisted whereas the solution was correct. In the following, we address the identified problematic cases in the context of the classes associated to the sole spectroscopic processing.

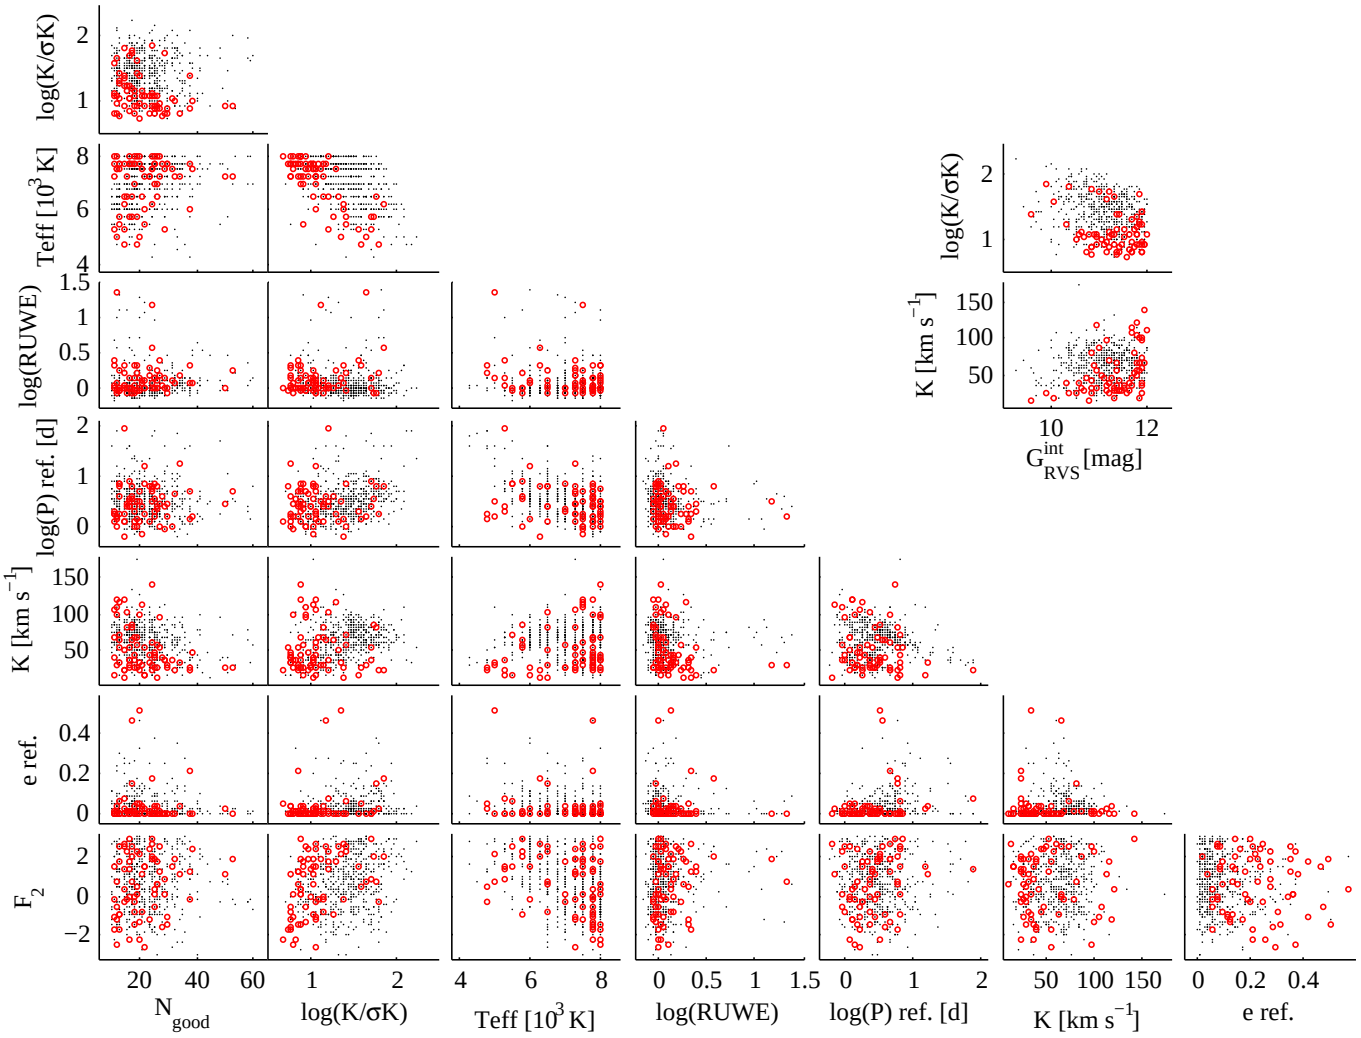

Fig. 46. Corner plot for the ASAS-SN sample, with the strong outliers in P deviating by more than 5σ highlighted in red. They constitute about 11% of the whole sample.

10.1.1. Intrinsic variables

Since the NSS-SB1 processing is applied to time series of variable RVs, it is expected to also treat some intrinsic variables that are exhibiting well-defined RV variations. This could be the case for example for Cepheids (actually the bright ones). The only firewall preventing the processing from delivering spurious orbital solutions for intrinsic variables is the fact that the processing only considers an harmonic content of the RV curve that is Keplerian (or nearly so). Unfortunately, Cepheids RV curves, for example, are renowned for their harmonic content being quite similar to the Keplerian one. This has been pointed out by Imbert et al. (1985); Imbert (1987); Imbert et al. (1989). The problem generated by intrinsic variables has been pointed out in Sect. 9.4.4 and in Sect. 9.4.5 where some of them were detected in the binary catalogue, mainly at short periods and high eccentricities.

As part of DR3, Ripepi et al. (2023) published a complete analysis of the *Gaia* photometric curves for a little more than 15 000 Cepheids (of all classes), 799 having time series of RVs. This list was cross-correlated with our SB1/SB1C list. They found 338 objects in common. Among them, the Cepheid nature is confirmed from the literature for 276 objects. After an eye inspection of photometric and RV curves, five of them were found doubtful Cepheids and have been rejected. Finally, 333 spurious SB1 solutions due to the Cepheid nature of the objects have been blacklisted from the SB subcatalogue.

Along these lines, we also considered the case of RR Lyrae stars (of all classes too). Clementini et al. (2023) gathered a list of a little more than 270 000 such objects, among which some 1096 have time series of RVs. The cross-match with our SB1 list yields 152 objects in common. Among these, 5 turned out to have a discrepancy between the period from our SB1 pipeline and the one derived by Clementini et al. (2023). Their RR Lyrae classification were thus considered as insecure. Therefore, we rejected (blacklisted), from the SB1 list, 147 objects considered as bona fide RR Lyrae. In summary, we rejected from our SB1 list some 480 sources as being intrinsic variables. Although a few pulsating stars (Cepheids, RR Lyrae) are rejected from our SB subcatalogue, they are nevertheless studied as pulsators elsewhere (see e.g. Eyer et al. 2023).

Another class of intrinsic variable objects exhibiting large photometric amplitudes are the long-period variables (LPVs). No similar secure list exists in DR3 for the LPVs. A tentative preliminary list exists that is based on the analysis of the photometry: the second *Gaia* catalogue of long-period variable candidates (Lebzelter et al. 2023). However, the related results

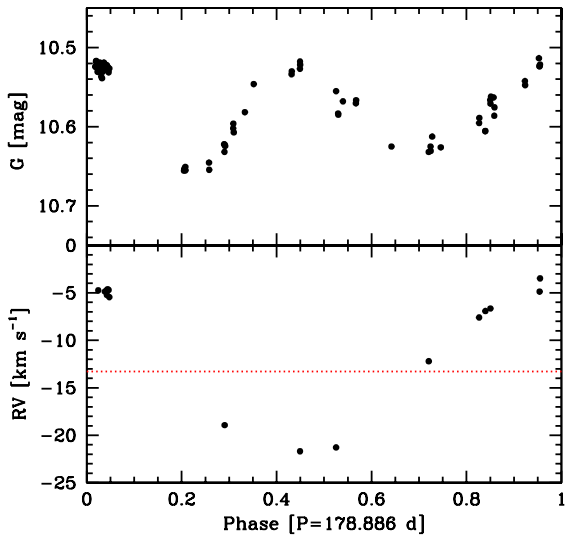

Fig. 47. Variability of *Gaia* DR3 5829160851462523008. Upper panel: *G* lightcurve of this object plotted as a function of the phase. Lower panel: RV curve of the same object. The red dotted line represents the systemic velocity. The zero of phase is arbitrary and has been chosen to improve the readability of the graph. Error bars are not given since they have sizes of the order of the symbols.

were not entirely accessible at the time of the execution of our pipeline. In addition, from the later provisional analysis, it became clear that this preliminary list most probably also contains a lot of other objects and in particular ellipsoidal variables. We illustrate here our concerns by further studying the case of *Gaia* DR3 5829160851462523008. This object was not known to vary before *Gaia*. It appears in the catalogue of LPV candidates with a derived photometric period of 90.5 d. This object also appears in the SB subcatalogue, produced in the present work, as a binary with a period of 178.886 d. Figure 47 exhibits both the RV curve and the photometric light curve of this object with the phase computed according to the RV period. From this figure, it is clear that the RV curve is unimodal, as it should be, whereas the *G* light curve presents two cycles in good agreement with the period attributed by Lebzelter et al. (2023). If the two maxima are very similar, the two minima are significantly different. In addition, the minima happen at phases where the RV curve crosses the systemic velocity. All this strongly suggests that this star is an ellipsoidal variable. Therefore, at the time of DR3, it appeared hazardous to filter our solutions on the basis of the existing LPV catalogue. Objects like *Gaia* DR3 5829160851462523008 would have been unduly rejected and we cautiously decided to not reject LPVs on this basis. Certainly, a sound analysis necessitates the concomitant study of RV and photometry, which was not foreseen for DR3. This should be done in the subsequent FPR (see Gaia Collaboration 2023c) where some 20 % of the LPV candidates are expected to be ellipsoidal variables. We encourage the user of the SB subcatalogue to filter out true intrinsic variables thanks to these subsequent results.

Yet another source of contaminants are main-sequence pulsators. The main periods of δ Scuti stars are typically below 0.2 d (Gaia Collaboration 2023b) and therefore lie outside the frequency range explored by our pipeline (even considering aliasing). However, γ Doradus stars are a more serious concern because they pulsate with periods in the range 0.3–1.4 d (Gaia Collaboration 2023b) and exhibit RV variations of the order of a few km s^{-1} (e.g. Mathias et al. 2004).

The problem of the rejection of intrinsic variables is certainly a difficult one, and much progress is still to be made on this topic. As mentioned above, the measurement of RVs by cross-correlating the spectra with template prevents us from identifying the origin of the RV variations.

10.1.2. Double-line transits

Since the RVs are originally measured by the STA processing on the basis of a per transit approach, there is initially no constraint from transit to transit. In particular, the same object could exhibit a single-line spectrum at a given time and a composite spectrum at another moment. As part of the spectroscopic processing, a value representing the percentage of transits leading to double-line detection is computed. If the star is constant or assumed non-variable, the median velocity is computed, a cut-off is then applied on this percentage in order to avoid the propagation of the effect of suspicious secondary RVs. If no transit exhibits a composite spectrum, the object presenting variations is processed as a single star. On the opposite, if a sufficient number of transits show a double-line spectrum, the object is processed as an SB2 (see Damerdji et al., in prep.). If the assumed constant object presents only a minority of double-line transits, the latter are ignored and the object is analysed on the basis of the sole single-line data. The presence of 10 % or more double-line transits is considered as a source of difficulties for the processing of single constant objects. It has been proposed a posteriori that the same rule (as the one applied to constant stars) should apply to the NSS-SB1 chain. This led to the rejection of 1475 `SB1` objects and of one `SB1C` object.

10.1.3. Scan-angle dependent signals

The RVS being a slitless spectrometer, the knowledge of the exact position of a star is necessary to have a correct calibration of the wavelength for its observed spectrum (Sartoretti et al. 2018). The astrometric solution assumes that all the objects are point-like, and the objects presenting incompatibility with this hypothesis have potentially wrong RVs. They are thus associated to large values of the `ruwe` parameter. This point is also mentioned in Sect. 2. To these deviating objects is associated an excess noise compared to the global astrometric solution (labelled `astrometric_excess_noise`). One of the main categories of the non-point-like objects are close pairs (physical or not) unresolved or partially resolved. A blend of objects that presents some elongated image will induce the following effect. The measured position of the star will have a position dependent on how the elongated axis of the pair enters the astrometric CCD field. Thus, the resulting RVs are dependent on the orientation of the main axis of the elongated objects with respect to the scan. It depends on the scan angle in a rather smooth manner. This potentially could generate, from transit to transit, RV variations that are a simple function of the scan angle. In particular, solutions for RV-constant objects could be associated to solutions mimicking a spectroscopic orbit at a period of 62.97 d; these orbits are consequently spurious. During the precession cycle, the scan angles of the various transits are performing a complete round. The phenomenon manifests itself as an excess of solutions with small K and thus small mass function (because the period is fixed). It is visible in Fig. 31 and is further illustrated in Fig. 48. Clearly, most of the solutions are spurious and are associated to large values of the astrometric excess noise. The suspicious K span the approximate interval between 0.8 and 16 km s^{-1}. Another illustration can be found in Gaia Collaboration (2023a),

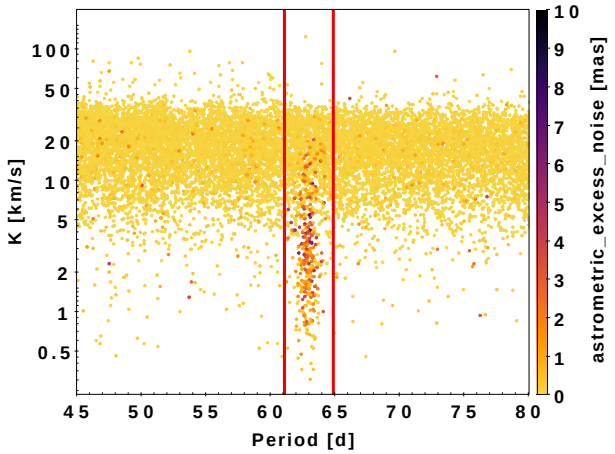

Fig. 48. Semi-amplitude K as a function of the period for the SB1 solutions. The data points are coloured as a function of the `astrometric_excess_noise`. A population of solutions very close to the 62.97 d (within ± 3 % as underlined by the vertical red lines) precession period exhibits small K and large such noise. The polluted domain spans values of K between 0.8 and 16 km s^{-1}.

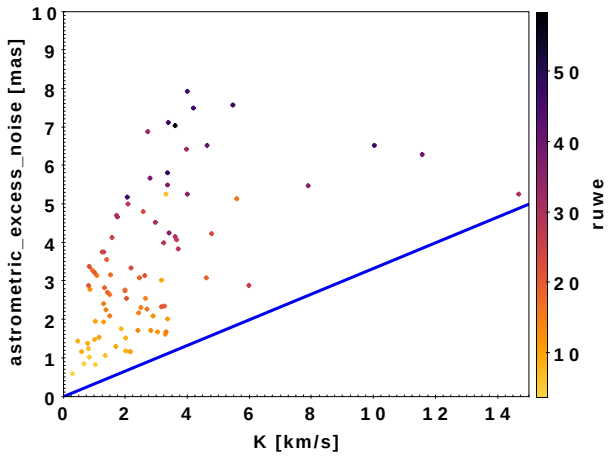

Fig. 49. `astrometric_excess_noise` as a function of the semi-amplitude, K, for the selected subsample. The data points are coloured according to the `ruwe`. The blue line indicates the separation above which the K is lower than three times the excess noise.

where it is also shown that the phenomenon could be related to a particular behaviour of `ipd_gof_harmonic_amplitude`. A detailed investigation of this behaviour is to be found in Holl et al. (2023a).

Thus, the epoch RVs are affected by unmodelled astrometric shifts along the scan direction. In order to alleviate the problem, the objects in the class `SB1` were first selected on the basis of an astrometric `ruwe` larger than 1.4 and a solution leading to a period within 3% of the precession period (62.97 d). In addition, we selected objects obeying the condition:

$$\texttt{astrometric\_excess\_noise}\,(\mathrm{km\,s^{-1}}) = 0.146 \times \texttt{astrometric\_excess\_noise}\,(\mathrm{mas}) > 0.5 \times \epsilon, \quad (20)$$

where ϵ is the mean of the uncertainties on the individual epoch RVs over the full time series. Figure 49 illustrates the distribution of the `astrometric_excess_noise` as a function of the parameter K for the selected sample. We can clearly see that all the K values are lower than three times the `astrometric_excess_noise` suggesting that these orbital solutions are spurious and are due to a scan-angle effect. The semi-amplitude is too small relatively to the error induced by the astrometric solution. It has been decided to reject and blacklist the 87 objects presenting these characteristics.

A similar approach concerning the satellite revolution period around the Sun (i.e. within 10% of one year) led to the rejection of another 77 sources.

10.1.4. Further remarks

All in all, 2107 objects were blacklisted for the above-mentioned three reasons. The mere sum of the individual rejected objects amounts to 2119, but 12 objects appear in two lists. The post-filtering described here above does not concern classes that have been combined. It could happen that a successful combination leads to a removal of some objects from the purely spectroscopic classes. If later the global fit is not successful, the object will disappear from the catalogue and will not be re-injected in any purely spectroscopic classes. Thus, some countings could be difficult to perform.

Although the post-filtering has been applied with great caution, it also happens that some problems persist. One of these cases is the following. Some objects exhibit a good SB1 or SB1C solution with a well established significant orbital semi-amplitude for *Gaia*. However, these solutions could sometimes be doubtful because the same object observed in the framework of the APOGEE DR17 survey (Abdurro'uf et al. 2022) or LAMOST survey (Cui et al. 2012; Bashi et al. 2022) appears to be constant in RV as detailed in Sect. 8.2.4. This anomaly has not been fully explained but it could be related in some cases to the presence of a nearby neighbour (see below). These objects have not been rejected. We encourage the catalogue user to be very cautious when using the present SB subcatalogue despite the fact that persisting problems should only represent a minority. All the remaining and identified problems should be corrected in the future DR4.

10.2. An explanation for gaia DR3 3376949338201650112

In Sect. 8.2.4, we report the existence of SB1 solutions presenting a marked periodic variation, whereas ground-based data conclude to the constancy of the related object. Here below, we propose an explanation for this discrepancy by investigating in details the case of *Gaia* DR3 3376949338201650112. This example has the advantage of being enlightening and will allow us to correct the pipeline for future releases.

The astrometric solution for this object is characterised by a `ruwe` of 25.35, pinpointing to a problem with the fit. Referring to the IPD parameters (image parameter determination; see e.g. Holl et al. 2023a), we remark that `ipd_gof_harmonic_amplitude` is equal to 0.164 which is certainly different from zero. This is a clue that the object is badly treated as a single object. The value of `ipd_frac_multi_peak` turns out to be 84 suggesting that the object is expected to be double or multiple on the plane of the sky: most probably a resolved or, more precisely, a partially resolved pair (according to the definition of Holl et al. 2023a). As explained by these authors, such pairs (or higher order) have a predisposition to suffer from scan-angle effects that induce a bias on the measured RVs that is function of the relative orientation of the scanning direction on the sky at the time of the observations. This is confirmed for *Gaia* DR3 3376949338201650112 by Fig. 50 where it becomes clear that the measured RVs are function of the

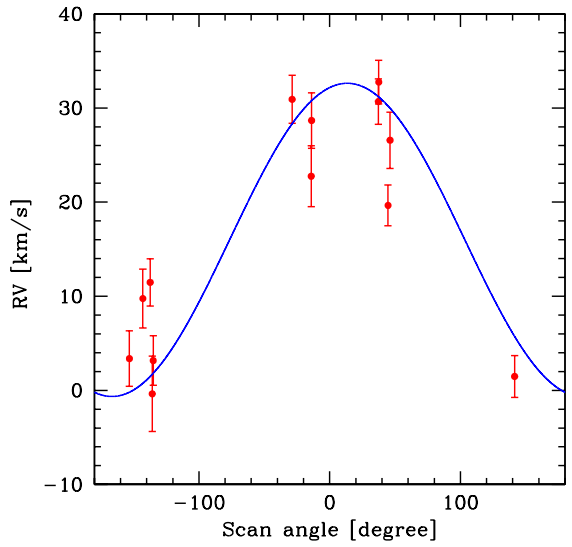

Fig. 50. RVs of *Gaia* DR3 3376949338201650112 as a function of the scan angle at the time of the observations. The blue curve represents the function given by Eq. (2) in Holl et al. (2023a) with the derived K taken from the catalogue.

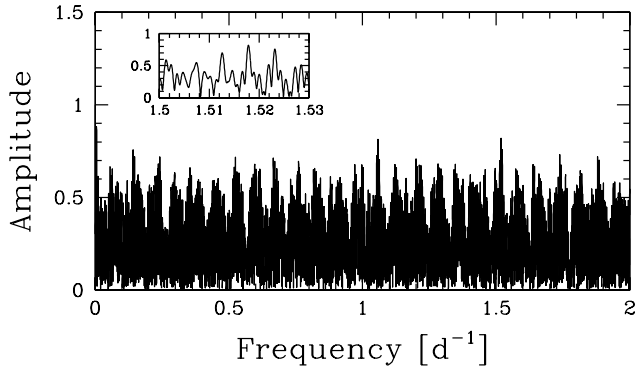

Fig. 51. Amplitude spectral window as a function of frequency as computed from the acquired time series for *Gaia* DR3 3376949338201650112. The insert is a zoom on the important second-highest peak, the first being at the origin as usual.

scan angle. The parameter `ipd_gof_harmonic_phase` adopts a value of 103.4° that suggests that the line joining the two stars of the pair has a position angle on the sky of 13.4°, after the application of a 90° correction as explained in Sect. 3.1.2 of Holl et al. (2023a). The semi-amplitude K of the fake orbital solution is about 16.6 km s^{-1} which translates according to Eq. (2) of Holl et al. (2023a) into a separation of the pair of 114 mas. The existence of such a pair is confirmed by a preliminary run of AGIS (AGIS 4.1) linked to the future DR4 that actually detects two objects *Gaia* DR4 3376949342501945344 and *Gaia* DR4 3376949342501945472 that are separated by some 200 mas and are exhibiting an orientation of 9°, thereby supporting our interpretation.

In principle, the bias due to the scan-angle effect is intended to produce variability with a period either of 62.97 d or of 1 y. This property has been used in Sect. 10.1.3 to apply post-filtering. However, in the present case, the period is 0.657582 d which does not appear to be compatible and could shed some doubt on our interpretation. The spectral window associated to the DR3 RV time series corresponding to this object is given in Fig. 51. Besides the main peak at $\nu = 0\,\mathrm{d}^{-1}$ as usual, it exhibits an important peak at $\nu = 1.51784\,\mathrm{d}^{-1}$ that is exceeding a value of 0.8; this is an indication that isolated aliasing could be important for this particular time series. An aliasing with $\delta\nu = 1.51784\,\mathrm{d}^{-1}$ explains the ambiguity between a period of 0.657582 d and that of 1 year. The combination of the aliasing phenomenon and of the scan-angle effect explains the existence of such fake SB1 objects in our SB subcatalogue: such objects were not post-filtered, since the post-filtering was based on the 62.97 d or 1 y timescales. We conclude that it would be more appropriate to include a statistical test inspired by graphs like the one shown in Fig. 50 rather than on a criterion based on the period. However, a true periodicity with these timescales will imply a similar behaviour of the RV dependency on the scan angle even if no scan-angle effect is present.

The explanation given here does not apply to the other object cited in Sect. 8.2.4 (*Gaia* DR3 702159502868472832).

10.3. The case of astrometric binaries

In Sects. 10.1.3 and 10.2, we discussed the effect of astrometric close pairs on the determination of RVs. A close pair fixed on the sky could often generate a scan-angle effect and lead to corrupted RV values but also could generate fake orbital motions. An additional degree of complexity is reached if we consider astrometric binaries. Due to their binarity, these close pairs are not fixed on the sky and the main axis of the pair is rotating due to the orbital motion. Therefore, there is interaction between the above-mentioned scan-angle effect and the astrometric orbital motion itself. The situation still increases in complexity if the astrometric orbital motion and the spectroscopic RV changes correspond to two aspects of the very same orbital motion and thus have the same period. In this latter case, the correction presents a timescale equal to the phenomenon we are studying. The correcting scheme is explained in Sect. 3.3.1 of Holl et al. (2023a) and is also illustrated there with the (somewhat extreme) example of *Gaia* DR3 6631710607341412096. For details, we refer to this paper.

The effect addressed here was not anticipated and was not fully understood at the time of the final delivery of the NSS code for DR3 processing. As a consequence, some of our solutions are damaged by this phenomenon. In addition, those pertaining to the `AstroSpectroSB1` class should be the most affected. For DR4, a preliminary run of the NSS astrometric channel has been included in order to compute a correction before the CU6 processing performing the RV measurements.

11. The SB subcatalogue

11.1. Description of the output

The present first catalogue of pure spectroscopic solutions is part of the *Gaia* Data Release 3. The orbital solutions (in class `SB1` or `SB1C` as indicated in the entry `nss_solution_id`) are part of an auxiliary table named `gaiadr3.nss_two_body_orbit`. We output in this table the classical Campbell orbital parameters introduced in Sect. 3.1 for the `SB1` class (P, γ, K, e, ω, T_0) along with their estimated uncertainties (1σ). The same is done for circular orbit parameters (P, γ, K, T_0) introduced in Sect. 3.2 and corresponding to the `SB1C` class. Beyond the orbital parameters themselves, this table also contains some indicators of the quality of the solution (see Sect. 6.4). The flags associated to the solution (see Sect. 6.5) are also provided. The boolean mask `bit_index` defines the parameters that are actually fitted, and

corresponds for the SB1 class (respectively SB1C) to the value 127 (respectively 31). In addition, is also provided the upper triangle of the correlation matrix of the parameters. It is expressed as corr_vec that contains the cross-coefficients folded into a vector and corresponding to the ordering of the adopted list of parameters. The first term of the vector $C(0)$ gives the element (0,1) of the matrix, the element $C(1)$ gives the element (0,2) of the matrix and element $C(2)$ the (1,2) one. Elements $C(3)$, $C(4)$, $C(5)$ correspond respectively to the matrix elements (0,3), (1,3), (2,3), and so on. This is thus a column-major storage scheme. The details of the output are described in Table H.1 for the class SB1 and in Table H.2 for the SB1C one. The table contains 181327 objects (with solutions) in the SB1 class and 202 in the SB1C one. The marked difference in the counting is due to our decision to consider eccentric solutions down to as small eccentricities as possible before going circular. This was described in Sect. 6.3.

We also released the trend solutions as discussed in Sect. 3.3 and in Sect. 6.3.5. The objects belonging to the TrendSB1 class are part of the auxiliary table gaiadr3.nss_non_linear_spectro. The subtype of trend solutions (linear or second-order trends) are indicated in nss_solution_id (respectively FirstDegreeTrendSB1 and SecondDegreeTrendSB1). They are characterised by a bit_index of 7 and 15, respectively. A few statistics about the quality of the fit are also given. The table contains 24 083 objects with solutions of type FirstDegreeTrendSB1 and 32 725 objects with SecondDegreeTrendSB1 ones, thus totalling 56 808 objects with solutions in the class TrendSB1. The details of the content of the latter table are described in Table H.3. We recall that the StochasticSB1 solutions are not delivered in DR3.

We should also point out that the auxiliary file gaiadr3.nss_two_body_orbit also contains the orbital parameters of the classes AstroSpectroSB1 and EclipsingSpectro that result from the action of the combiner (see Sects. 3.5 and 7.2).

11.2. Weaknesses, caveats, and recommendations

The present section gathers the caveats dispersed all over the paper and delivers a summary of the characteristics of the catalogue that could be beneficial to the general user. The RVs studied here have the quality and shortcomings of the measurements of individual (epoch) RVs for constant stars. The properties and precision of these RVs for SB1 studies should be of the same order. Therefore, we refer the reader to the work of Katz et al. (2023) concerning the characteristics and condition of utilisation of these data. Seabroke et al. (in preparation) is also a work to be consulted.

The information extracted from variations of the RVs could be delicate to use. Firstly, it must be clear that the way the measurements are performed (cross-correlation with a template) makes the implicit assumption that the line is supposed to vary without changing its shape. This is not necessarily true. Two textbook cases at least come to mind. Some stars, notably the intrinsically variable ones, could exhibit dominant variations of their line shapes. In such a case, the variations do not represent a Doppler shift of the object as a whole and thus are not representative of an SB1 (see Sect. 10.1.1). In principle, the fact we are fitting models where the harmonic content is typical of a Keplerian motion is certainly a protection against variable stars appearing in the SB1 class. However, some variable stars, like Cepheids are well-known to display RV changes that closely mimic the Keplerian motion (e.g. Imbert et al. 1989). As explained in Sect. 10.1.1, efforts have been made to post-filter the Cepheid and the RR Lyrae stars, but this post-filtering is certainly not complete for the time being. In addition, no such step was made for other intrinsic variables in DR3 (notably Long-Period Variables).

Another tricky case is that of true SB2 stars, where the two lines never fully deblend. They are appearing most probably in the SB1 channel. The RV variations are then due to the corresponding change of the full profile made by the blended lines. This has little effect on the period determination but the impact on the eccentricity and mainly on the semi-amplitude might be huge (Sect. 8.2.3). The related solutions are characterised by a too small value of K and of the mass function. They could be considered as suspicious but we cannot have any certitude that all of these types of solutions are necessarily wrong.

The algorithm used is intended to fit orbital solutions to the RV curves that it analyses. It is doing so but is not protected against some anomalies. The RV data set suffers from all the problems not attenuated by the CU6 processing and is not particularly protected against outliers, miss-classified double-line objects, intrinsic variables, emission-line objects and other problematic situations (e.g. badly normalised continuum, etc...). In addition, it is not (for the moment) able to properly deal with higher degree systems. It is possible that the code detects a binary motion, but this one could be quantitatively corrupted by the presence of additional components (see Sect. 8.2.2). In the future (DR4), we hope to include a cleaning (whitening) of the dominant signal followed by a new search for a second period or a trend. However, the presence of two signals renders the statistical analysis much more complicated.

Instrumental effects are also present. The scan-angle effect notably induces periodic RV variations for constant stars that do not fulfil the assumptions used in astrometry (objects are assumed to be isolated and thus single). Close astrometric pairs could thus mimic SB1 RV variations (see the work by Holl et al. 2023a), and are polluting the present sample (even if some post-filtering took place; see also Sect. 10.1.3). If, in addition, the close pair is rotating on the plane of the sky as is the case for astrometric binaries, the measured RV could be corrupted (see Sect. 10.3). This problem will be corrected for DR4.

In addition, it should be noted that all decisions are taken in a statistical way; this means that the output is statistically correct but that some incorrect individual fits can survive near the borders where the decisions are taken (e.g. between an orbital fit and a trend solution for example).

We recommend to the user of the catalogue to keep in mind the weaknesses and caveats described here above. We also strongly recommend, when possible, the use of the various quality factors accompanying each solution. If the catalogue is expected to be well-behaved from the statistical point of view, the situation could be more intricate for individual objects or particular set of objects. It remains that biases in the catalogue could be present from the point of view of stellar population studies. Such an analysis is beyond the scope of the present paper. For the time being, the genuine selection function of the catalogue is poorly known and any user should be aware of this and behave accordingly. The user is strongly encouraged to adapt the selection of solutions to their particular objective instead of using it blindly. They should combine the various delivered quality factors according to their aim. We advise a cautious approach. An example of interesting complex combinations is detailed in Bashi et al. (2022). To some extent, other information can be considered as quality factors (such as flags but, mainly, the number of

points in the time series N_{good}), which could also be included in the selection process.

The global weak point of the algorithms could be the determination of the period even if here above it is shown that the success rate is at the very least 80% and is more often located in the range 85–95%. There is evidently at a very low number of transits, a difficulty to correctly determine the period that forced us to neglect data with N_{good} less than 10. However, at slightly larger N_{good}, the situation is not so simple. From the simulations described in Sect. 9.1, we conclude that a value of N_{good} smaller than 17 is somewhat problematic (although it depends on the signal-to-noise ratio) for the period recovery. We also show, from the validation (see Sect. 9.2), that an overfitting phenomenon characterises part of the solutions associated to an N_{good} in the range 10–15. We should further add that, at small N_{good}, the periodogram has a larger propensity to produce isolated aliasing peaks (generating ambiguity due to aliasing). While interpreting the simulations, it appears that it is unclear if the determining parameter is the N_{good} or the significance. According to these simulations, significance lower than 10–15 can be associated with difficulties to have the proper recovery rate. This seems to be confirmed by the validation in Sect. 9.4, in particular in the corner plot of Fig. 46. It is also clear, from Fig. 30, that the circularisation phenomenon is not outstandingly visible for significance much lower than 30. Therefore, we consider that depending on the studies performed, the significance should be cut between 15 and 25, or more conservatively at 30. If 10 is still affordable in some cases, the solutions characterised by a significance between 5 (our cut-off) and 10 are most probably hazardous to accept.

The spectroscopic solutions could be merged with photometric or astrometric ones, and in these cases the resulting output is certainly very secure for the period. If the user of the catalogue only considers the sole results of the spectroscopy fits, it is advisable to restrict the sample to some kind of gold sample. This can be done by performing additional selections in the catalogue: for example, a larger number of transits, a higher significance, periods in a dedicated range (between ten days and ΔT, or allowing it to go down to smaller periods but at the expense of a selection on the basis of a larger significance), periods with small uncertainties, solutions with negligible phase gaps (although the gap values are not provided in the DR3 catalogue), and so on. In Appendix J, we present indicative but simple examples of extraction of a gold sample from the catalogue.

12. Conclusions

Within the framework of the *Gaia* Data Release 3, we present the first version of the *Gaia* catalogue of spectroscopic orbital solutions (SB subcatalogue). We present 181 327 objects that have been found to be characterised by class SB1 and 202 objects by class SB1C. We also present objects exhibiting trends of variations of their RVs over time.

The present catalogue constitutes the largest collection of objects with a spectroscopic orbital solution by far, based on a comparison with the notorious SB9 compilation[14]. The improvement in terms of number of objects is more than one order (nearly two) of magnitude. The present SB subcatalogue has also the major advantage of being homogeneous from the point of view of both data acquisition (properties) and data processing, surpassing the characteristics of existing compilation. The only exception is the typical number of measurements that constitutes the individual time series, which remains limited (at least for the present DR3 catalogue). The situation will undoubtedly improve for future data releases.

After a description of the data entering the NSS-SB1 processing, we describe in detail the logic and structure of the related code. We reveal the general philosophy behind, along with the mathematical formalism used. We also describe, following the main processing details, the filtering of the results performed to discard bad solutions and improve the catalogue. We give a number of examples illustrating the quality of the output, as well as a few examples of detected problematic cases. A full validation of the output is performed both on the basis of internal consistency and on the basis of external existing catalogues from the literature. We further analyse the characteristics of the catalogue and also describe the way the SB subcatalogue has been included in the *Gaia* Data Release 3. We additionally scrutinises the weaknesses of the approach and provide the necessary caveats and recommendations for scientific applications of the catalogue.

Data availability

The epoch RVs constituting the time series analysed here have not been published in the context of the *Gaia* Data Release 3. The orbital solutions and trends we derived are available since 13 June 2022, in the ESA *Gaia* DR3 archive accessible at https://gea.esac.esa.int/archive. The same results can also be found at CDS Strasbourg via the link https://cdsarc.cds.unistra.fr/viz-bin/Cat?I/357. Appendices D to J are only available on Zenodo (see https://zenodo.org/records/13990211).

Acknowledgements. The authors are particularly thankful to Alain Detal who maintained over the years the Liège computer park with great efficiency. EG is also thanking F. Mélen for her consequent help in the manipulation of various computer codes. The authors are acknowledging the numerous scientific discussions that took place with Christos Siopis. The participation in this endeavour has been made possible through an Impulsion Credit from the University of Liège to EG. This work presents results from the European Space Agency (ESA) space mission *Gaia*. *Gaia* data are being processed by the *Gaia* Data Processing and Analysis Consortium (DPAC). Funding for the DPAC is provided by national institutions, in particular the institutions participating in the *Gaia* MultiLateral Agreement (MLA). The *Gaia* mission website is https://www.cosmos.esa.int/gaia. The *Gaia* archive website is https://archives.esac.esa.int/gaia. The *Gaia* mission and data processing have, in Liège, been continuously supported by the Belgian Politique Scientifique Fédérale (BELSPO) through PRODEX contracts *Gaia DPAC: Binaires, étoiles extrêmes et objets du système solaire* (convention PEA 4200090289; years 2007–2013), *Gaia DPAC Early Mission Belgian consolidation (Gaia-be)* (convention PEA 4000110152: years 2014–1016), *Gaia DPAC Early Mission Belgian consolidation (Gaia-be2)* (convention PEA 4000119840: years 2017–1019; years 2020–2022). TM acknowledges financial further support from BELSPO for contract PLATO mission development. We acknowledge support from an ARC grant (13/17-06; *Massive Stars: drivers of the evolution of the Universe*) for concerted research actions financed by the French community of Belgium (Wallonia-Brussels Federation). We are also largely indebted to the Fonds de la Recherche Scientifique (F.R.S.-FNRS, Belgium) for various supports. YD acknowledges support to the Algerian Centre de Recherche en Astronomie, Astrophysique et Géophysique of Bouzareah Observatory. GS acknowledges support from BELSPO/PRODEX under various contracts and in particular convention PEA 40000119826. AS acknowledges support from the Italian Space Agency (ASI) under contract 2018-24-HH.0 "The Italian participation to the Gaia Data Processing and Analysis Consortium (DPAC)" in collaboration with the Italian National Institute of Astrophysics.

References

Abdurro'uf, Accetta, K., Aerts, C., et al. 2022, ApJS, 259, 35
Andrae, R., Fouesneau, M., Sordo, R., et al. 2023, A&A, 674, A27
Antoja, T., Helmi, A., Romero-Gómez, M., et al. 2018, Nature, 561, 360

[14] https://sb9.astro.ulb.ac.be/

Armstrong, D. J., Kirk, J., Lam, K. W. F., et al. 2015, A&A, 579, A19
Bashi, D., Shahaf, S., Mazeh, T., et al. 2022, MNRAS, 517, 3888
Bevington, P. R., & Robinson, D. K. 2003, Data Reduction and Error Analysis for the Physical Sciences, 3rd edn. (MA: McGraw-Hill)
Binnendijk, L. 1960, Properties of Double Stars; A Survey of Parallaxes and Orbits (Philadelphia: University of Pennsylvania Press)
Bloomfield, P. 1976, Fourier Analysis of Time Series: An Introduction (New York: Princeton University, John Wiley and Sons, Inc.)
Boulkaboul, A., Damerdji, Y., Morel, T., et al. 2022, MNRAS, 517, 1849
Brandt, T. D. 2021, ApJS, 254, 42
Buder, S., Asplund, M., Duong, L., et al. 2018, MNRAS, 478, 4513
Clementini, G., Ripepi, V., Garofalo, A., et al. 2023, A&A, 674, A18
Cropper, M., Katz, D., Sartoretti, P., et al. 2018, A&A, 616, A5
Cui, X.-Q., Zhao, Y.-H., Chu, Y.-Q., et al. 2012, Res. Astron. Astrophys., 12, 1197
Cumming, A. 2004, MNRAS, 354, 1165
David, M., Blomme, R., Frémat, Y., et al. 2014, A&A, 562, A97
Deeming, T. J. 1975, Ap&SS, 36, 137
Deeming, T. J. 1976, Ap&SS, 42, 257
Delchambre, L. 2010, Recherche d'une évaluation robuste de l'excentricité dans des systèmes binaires d'étoiles, Master Thesis Liège University (available at https://orbi.uliege.be/handle/2268/322803)
Eichhorn, H. K., & Xu, Y.-L. 1990, ApJ, 358, 575
Endl, M., Kürster, M., Els, S., et al. 2002, A&A, 392, 671
ESA, ed. 1997, ESA Special Publication, 1200, The HIPPARCOS and TYCHO catalogues. Astrometric and photometric star catalogues derived from the ESA HIPPARCOS Space Astrometry Mission
ESA, 1997, VizieR Online Data Catalog: I/239
Evans, D. W., Riello, M., De Angeli, F., et al. 2018, A&A, 616, A4
Eyer, L., Audard, M., Holl, B., et al. 2023, A&A, 674, A13
Fabricius, C., Luri, X., Arenou, F., et al. 2021, A&A, 649, A5
Ford, E. B., & Gregory, P. C. 2007, in Astronomical Society of the Pacific Conference Series, 371, Statistical Challenges in Modern Astronomy IV, eds. G. J. Babu, & E. D. Feigelson, 189
Fraser, O. J., Hawley, S. L., & Cook, K. H. 2008, AJ, 136, 1242
Frémat, Y., Royer, F., Marchal, O., et al. 2023, A&A, 674, A8
Fukushima, T. 1996, AJ, 112, 2858
Gaia Collaboration (Brown, A. G. A., et al.) 2016a, A&A, 595, A2
Gaia Collaboration (Prusti, T., et al.) 2016b, A&A, 595, A1
Gaia Collaboration (Brown, A. G. A., et al.) 2018a, A&A, 616, A1
Gaia Collaboration (Katz, D., et al.) 2018b, A&A, 616, A11
Gaia Collaboration (Arenou, F., et al.) 2023a, A&A, 674, A34
Gaia Collaboration (De Ridder, J., et al.) 2023b, A&A, 674, A36
Gaia Collaboration (Trabucchi, M., et al.) 2023c, A&A, 680, A36
Gaia Collaboration (Vallenari, A., et al.) 2023d, A&A, 674, A1
Golub, G., & Pereyra, V. 2003, Inverse Probl., 19, R1
Gosset, E., Royer, P., Rauw, G., Manfroid, J., & Vreux, J. M. 2001, MNRAS, 327, 435
Griffin, R. F. 1969, MNRAS, 145, 163
Griffin, R. F. 2007, The Observatory, 127, 113
Halbwachs, J. L., Mayor, M., & Udry, S. 2005, A&A, 431, 1129
Halbwachs, J.-L., Pourbaix, D., Arenou, F., et al. 2023, A&A, 674, A9
Heck, A., Manfroid, J., & Mersch, G. 1985, A&AS, 59, 63
Holanda, N., Barros, P. P., & Silva, A. N. 2018, New A, 65, 21
Holl, B., Fabricius, C., Portell, J., et al. 2023a, A&A, 674, A25
Holl, B., Sozzetti, A., Sahlmann, J., et al. 2023b, A&A, 674, A10
Imbert, M. 1987, A&A, 175, 30
Imbert, M., Andersen, J., Ardeberg, A., et al. 1985, A&AS, 61, 259
Imbert, M., Andersen, J., Ardeberg, A., et al. 1989, A&AS, 81, 339
Jayasinghe, T., Rowan, D. M., Thompson, T. A., Kochanek, C. S., & Stanek, K. Z. 2023, MNRAS, 521, 5927
Katz, D., Sartoretti, P., Cropper, M., et al. 2019, A&A, 622, A205
Katz, D., Sartoretti, P., Guerrier, A., et al. 2023, A&A, 674, A5
Kawata, D., Baba, J., Ciucă, I., et al. 2018, MNRAS, 479, L108
Kervella, P., Arenou, F., Mignard, F., & Thévenin, F. 2019, A&A, 623, A72
Khoperskov, S., Gerhard, O., Di Matteo, P., et al. 2020, A&A, 634, L8
Kochanek, C. S., Shappee, B. J., Stanek, K. Z., et al. 2017, PASP, 129, 104502
Koppelman, H. H., Helmi, A., Massari, D., Roelenga, S., & Bastian, U. 2019, A&A, 625, A5
Kunder, A., Kordopatis, G., Steinmetz, M., et al. 2017, AJ, 153, 75
Lebzelter, T., Mowlavi, N., Lecoeur-Taibi, I., et al. 2023, A&A, 674, A15
Lester, K. V. 2020, PhD thesis, Georgia State University, United States
Lindegren, L. 2019, Gaia Data Processing and Analysis Consortium (DPAC) technical note GAIA-C3-TN-LU-LL-125, http://www.cosmos.esa.int/web/gaia/public-dpac-documents
Lindegren, L., Hernández, J., Bombrun, A., et al. 2018, A&A, 616, A2
Lindegren, L., Klioner, S. A., Hernández, J., et al. 2021, A&A, 649, A2
López-Corredoira, M., & Sylos Labini, F. 2019, A&A, 621, A48
López-Corredoira, M., Garzón, F., Wang, H. F., et al. 2020, A&A, 634, A66
Lucy, L. B., & Sweeney, M. A. 1971, AJ, 76, 544
Marchetti, T., Rossi, E. M., & Brown, A. G. A. 2019, MNRAS, 490, 157
Mathias, P., Le Contel, J. M., Chapellier, E., et al. 2004, A&A, 417, 189
Mazeh, T. 2008, in EAS Publications Series, 29, EAS Publications Series, eds. M. J. Goupil, & J. P. Zahn, 1
Monari, G., Famaey, B., Siebert, A., Wegg, C., & Gerhard, O. 2019, A&A, 626, A41
Mowlavi, N., Holl, B., Lecoeur-Taïbi, I., et al. 2023, A&A, 674, A16
Nelson, R. H., & Robb, R. M. 2017, Inform. Bull. Variable Stars, 6203, 1
Nitschai, M. S., Cappellari, M., & Neumayer, N. 2020, MNRAS, 494, 6001
Perryman, M. 2009, Astronomical Applications of Astrometry: Ten Years of Exploitation of the Hipparcos Satellite Data (Cambridge University Press)
Petrosky, E., Hwang, H.-C., Zakamska, N. L., Chandra, V., & Hill, M. J. 2021, MNRAS, 503, 3975
Pourbaix, D., Arenou, F., Gavras, P., et al. 2022, Gaia DR3 Documentation. Chapter 7: Non-single stars, Gaia DR3 Documentation, European Space Agency; Gaia Data Processing and Analysis Consortium. Online at https://gea.esac.esa.int/archive/documentation/GDR3/index.html, 7
Pourbaix, D., Tokovinin, A. A., Batten, A. H., et al. 2004, A&A, 424, 727
Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 2007, Numerical Recipes: The Art of Scientific Computing, 3rd edn. (Cambridge University Press)
Prevot, L., Andersen, J., Ardeberg, A., et al. 1985, A&AS, 62, 23
Price-Whelan, A. M., Hogg, D. W., Foreman-Mackey, D., & Rix, H.-W. 2017, ApJ, 837, 20
Price-Whelan, A. M., Hogg, D. W., Rix, H.-W., et al. 2020, ApJ, 895, 2
Ramos, P., Antoja, T., & Figueras, F. 2018, A&A, 619, A72
Recio-Blanco, A., de Laverny, P., Palicio, P. A., et al. 2023, A&A, 674, A29
Riello, M., De Angeli, F., Evans, D. W., et al. 2018, A&A, 616, A3
Riello, M., De Angeli, F., Evans, D. W., et al. 2021, A&A, 649, A3
Ripepi, V., Clementini, G., Molinaro, R., et al. 2023, A&A, 674, A17
Rowan, D. M., Jayasinghe, T., Stanek, K. Z., et al. 2022, MNRAS, 517, 2190
Sartoretti, P., Katz, D., Cropper, M., et al. 2018, A&A, 616, A6
Sartoretti, P., Blomme, R., David, M., & Seabroke, G. 2022, Gaia DR3 Documentation. Chapter 6: Spectroscopy, Gaia DR3 Documentation, European Space Agency; Gaia Data Processing and Analysis Consortium. Online at https://gea.esac.esa.int/archive/documentation/GDR3/index.html, 6
Sartoretti, P., Marchal, O., Babusiaux, C., et al. 2023, A&A, 674, A6
Schlesinger, F. 1910, Publ. Allegheny Observ. Univ. Pittsburgh, 1, 33
Seabroke, G. M., Fabricius, C., Teyssier, D., et al. 2021, A&A, 653, A160
Sharma, S. 2017, ARA&A, 55, 213
Soubiran, C., Jasniewicz, G., Chemin, L., et al. 2018, A&A, 616, A7
Sterne, T. E. 1941, PNAS, 27, 175
Tatum, J. B. 1968, MNRAS, 141, 43
Trick, W. H., Coronado, J., & Rix, H.-W. 2019, MNRAS, 484, 3291
Wilson, E. B., & Hilferty, M. M. 1931, PNAS, 17, 684
Zechmeister, M., & Kürster, M. 2009, A&A, 496, 577

Appendix A: Historical developments

Once the period has been evaluated, the general model to fit (Eq. 1) remains non-linear in some of the parameters (mainly the eccentricity, e). This implies that the χ^2 of the fit could present a few or even numerous minima in the parameter space, and to find the global one is a true challenge. In particular, the derivation of the eccentricity is a notorious pitfall. The search for the global minimum might be performed empirically by guessing its approximate position and then by descending the χ^2 valley by classical numerical methods (such as the Levenberg-Marquardt algorithm). However, in a pipeline as the one developed for *Gaia*, identifying the location of the global minimum is uncertain and a specific type of search is necessary. This is particularly frustrating because a well-trained expert eye is able to estimate a rough value for the eccentricity with great efficiency. In this context, we explored the possibilities to adopt alternative approaches. We acted along two main axes. On the one hand, we tried to invent a method based on pattern recognition in the phase diagram of the RVs folded with the candidate periods. Such an approach would allow us to take a first guess for the eccentricity, which could then be used for further classical derivation. In order to apply pattern recognition techniques, it is mandatory to define an interpolation function in the phase diagram to obtain a robust description of the shape of the RV curve. Several tools have been explored such that the nearest-neighbour interpolation, the linear interpolation, the interpolation by cosines, the local Akima periodic interpolation, the periodic interpolation by cubic splines, and the periodic interpolation by Catmull-Rom splines. We also explored the possibilities of performing classical adjustment by least-square methods. In particular, we considered classical polynomials, Fourier series, Bézier curves (using De Casteljan algorithm) and finally, cubic B-splines in their open uniform version. Once the interpolation curve is chosen and the shape is defined, we tested various pattern recognition methods such as algorithms for regression based on the nearest neighbourhood, algorithms for regression based on trees and in particular on extremely randomised trees and, finally, algorithms for regression based on neural networks. Full details of this exploration and detailed references are given in the work of Delchambre (2010).

On the other hand, we developed an approach based on genetic algorithms. These algorithms have the advantage to circumvent any preliminary guess of the solution. This is at the expense of being obliged to explore (even cleverly) the space of the parameters using the generation of random numbers (Monte Carlo methods). This work is described in details in Delchambre (2010). The use of genetic algorithms is a trade-off between the probability to effectively detect the global χ^2 minimum and the consumption of computer time. The codes using genetic algorithms were written, delivered and tested. The detection of the global minimum was possible but was implying a prohibitively large amount of computer time despite strong efforts made to accelerate it. The large amount of objects to treat in the *Gaia* processing suggests using much more efficient algorithms. The genetic algorithm code was not activated for the operational runs.

Other methods are now available to derive orbital solutions. They became widespread in the domain of exoplanets, where the orbits are particularly incompletely covered and complex. The majority of these methods are based on Monte Carlo simulations. Very often, they use the Markov chain Monte Carlo (MCMC) method (see e.g. Sharma 2017) often utilised in conjunction with Bayesian approaches. These algorithms necessitate a dense coverage of the full space of priors, at the expense of the computer time. Comparing the works of Price-Whelan et al. (2017) and Price-Whelan et al. (2020) pinpoints the strong impact that could have the choice of priors on the resulting posteriors. The Bayesian methods have admittedly the big advantage to potentially provide marginal posterior probabilities for a series of models and thus estimate how much the observational data could favour a particular model. According to Ford & Gregory (2007), the challenge of estimating marginal posterior probabilities is still a crucial problem and obvious estimators are suffering from poor convergence properties. In addition, for the preliminary DR3 catalogue, the directives were that we should output one and only one model per treated object, thus mitigating the interest of using Bayesian methods. In addition, these Monte Carlo methods are computer time-consuming and were not acceptable for the DR3 processing. Finally, we adopted a method based on fully analytical and classical numerical developments.

Appendix B: Practical details on the processing

The Data Processing Centre at CNES (DPCC[15]), located at the Centre National d'Etudes Spatiales in Toulouse (France), ran the spectroscopic processing chain (CU6) and the object processing chain (CU4). Globally, it ran the whole set of NSS codes and, in particular, the spectroscopic orbital solution chain described in the present paper. Our NSS-SB1 processing chain delivers the results to the combiner. The Main Data Base (MDB) is fed with solutions from this combiner. If the combiner finds no other solutions to combine with, the spectroscopic results are directly transmitted to the MDB. If a combination is possible, only the output of the combination is saved.

In the framework of the *Gaia* endeavour, the DPCC is hosting a cluster of 250 computers (6000 cores, 40 TBytes of RAM and 7 PBytes of HDFS disk storage). The resources are managed by an Hadoop system. The codes are integrated into a framework SAGA (System of Accomodation of *Gaia* Algorithms) developed by Thales. The spectroscopic processing was executed over 2500 cores and occupied 17 TBytes of RAM. Some 3×10^6 hours of CPU were necessary, equivalent in real time to 120 days. Concerning the NSS chain, the full treatment (Ingestion, Processing, and Post-Processing) was executed on the Hadoop cluster over 1400 cores and necessitated 16 105 days of CPU corresponding to an effective duration of 17 days. All the codes were written in Java 8.0. The version of the code tagged and executed for the operational run was 20.10.

[15] https://www.cosmos.esa.int/web/gaia/dpac

Appendix C: Separable variable technique applied to the Keplerian model fit

Appendix C.1. A reminder of the Levenberg-Marquardt method

Eq. (10) links the N dependent variables $RV(t_i)$ (the RVs of the model) to the independent ones t_i (the observing times). Hereafter, we represent these two variables by the column arrays **RV** and **t** respectively. f represents the Keplerian function, such that

$$\mathbf{RV} = f\left(\mathbf{t}, \mathbf{p}\right), \tag{C.1}$$

with **p** being the array of 6 parameters $[\gamma, \alpha, \beta, \nu_k, e, T_0]^T$ where the superscript T denotes the transpose operation. The Newton-Raphson algorithm is based on the Taylor first order development

$$\mathbf{RV} = f\left(\mathbf{t}, \mathbf{p}^{\{n\}}\right) + \mathbf{J}^{\{n\}}\, \mathbf{\Delta p}^{\{n\}}, \tag{C.2}$$

where $\mathbf{\Delta p}^{\{n\}} = \mathbf{p}^{\{n+1\}} - \mathbf{p}^{\{n\}}$, and $\{n\}$ denotes the iteration number. **J** is the Jacobian matrix which elements are

$$J_{i,j}^{\{n\}} = \frac{\partial f\left(t_i, p_j^{\{n\}}\right)}{\partial p_j^{\{n\}}} \quad \text{for} \;\; 1 \leq i \leq N,\; 1 \leq j \leq 6. \tag{C.3}$$

The objective function to be minimised is the quadratic sum of residuals as

$$\chi^2 = \left[\mathbf{RV}_{\rm obs} - f\left(\mathbf{t}, \mathbf{p}\right)\right]^T \mathbf{W} \left[\mathbf{RV}_{\rm obs} - f\left(\mathbf{t}, \mathbf{p}\right)\right], \tag{C.4}$$

with **W** being the inverse of the data variance-covariance matrix. The χ^2 is given at the $n^{\rm th}$ iteration by

$$\chi^{2,\{n\}} = \left[\mathbf{RV}_{\rm obs} - f\left(\mathbf{t}, \mathbf{p}^{\{n\}}\right) - \mathbf{J}^{\{n\}}\, \mathbf{\Delta p}^{\{n\}}\right]^T \mathbf{W} \left[\mathbf{RV}_{\rm obs} - f\left(\mathbf{t}, \mathbf{p}^{\{n\}}\right) - \mathbf{J}^{\{n\}}\, \mathbf{\Delta p}^{\{n\}}\right], \tag{C.5}$$

where $\mathbf{RV}_{\rm obs}$ contains the observed RVs. The optimal parameter step $\mathbf{\Delta p}^{\{n\}}$ is obtained by solving

$$\frac{\partial \chi^2}{\partial \mathbf{\Delta p}^{\{n\}}} = 0. \tag{C.6}$$

In the case of Newton-Raphson, this leads to the following algebraic system:

$$\left(\mathbf{J}^{\{n\}T}\, \mathbf{W}\, \mathbf{J}^{\{n\}}\right) \mathbf{\Delta p}^{\{n\}} = \mathbf{J}^{\{n\}T} \mathbf{W} \left[\mathbf{RV}_{\rm obs} - f\left(\mathbf{t}, \mathbf{p}^{\{n\}}\right)\right]. \tag{C.7}$$

The matrix $\mathbf{H}^{\{n\}} = \mathbf{J}^{\{n\}T}\, \mathbf{W}\, \mathbf{J}^{\{n\}}$ is a positive semi-definite matrix called the hessian matrix (it is not a positive definite matrix because it might be singular).

The Levenberg-Marquardt algorithm consists in rescaling the diagonal elements of $\mathbf{H}^{\{n\}}$ by the damping factor $1 + \lambda$. The system (Eq. C.7) becomes

$$\left[\mathbf{H} + \lambda\, \mathrm{diag}\left(\mathbf{H}\right)\right]^{\{n\}} \mathbf{\Delta p}^{\{n\}} = \mathbf{J}^{\{n\}T} \mathbf{W} \left[\mathbf{RV}_{\rm obs} - f\left(\mathbf{t}, \mathbf{p}^{\{n\}}\right)\right]; \tag{C.8}$$

where λ is the damping factor, chosen to be large when **H** approaches the singularity.

Appendix C.2. Separable variable technique

The parameters $\mathbf{p_l} = [\gamma, \alpha, \beta]^T$ are linear because they satisfy the following property:

$$f\left(\mathbf{t}, \mathbf{p}\right) = \sum_{k=1}^{3} \frac{\partial f\left(\mathbf{t}, \mathbf{p}\right)}{\partial \mathbf{p_l}(k)}\, \mathbf{p_l}(k) = \mathbf{J_l}\, \mathbf{p_l}. \tag{C.9}$$

Hereafter, subscripts **l** and **nl** refer to linear and non-linear respectively. Then, $\mathbf{J_l}$ is the Jacobian matrix whose columns are the derivatives of f with respect to $\mathbf{p_l}$.

While in both Newton-Raphson and the (commonly used) Levenberg-Marquardt (LM) algorithms, the linear and non-linear parameters are treated the same way, the separable variable technique has the advantage that $\mathbf{p_l}$ can be obtained for every non-linear parameter $\mathbf{p_{nl}}^{\{n\}} = [\nu_k, e, T_0]^{\{n\}T}$ by a linear least-square regression, such that [16]

$$\mathbf{H_l}^{\{n\}}\, \mathbf{p_l}^{\{n\}} = \mathbf{Z_l}^{\{n\}}, \tag{C.10}$$

[16] The inversion of the normal equations is known to be numerically unstable and should be avoided in favour of orthogonal decomposition approaches.

with

$$\mathbf{H_l}^{\{n\}} = \mathbf{J_l}^{\{n\}T}\,\mathbf{W}\,\mathbf{J_l}^{\{n\}}, \tag{C.11}$$
$$\mathbf{Z_l}^{\{n\}} = \mathbf{J_l}^{\{n\}T}\,\mathbf{W}\,\mathbf{RV}_{\rm obs}. \tag{C.12}$$

Therefore, Eq. C.2 becomes a function of $\mathbf{p_{nl}}^{\{n\}}$ only, such that

$$\mathbf{RV} = f\left(\mathbf{t}, \mathbf{p}^{\{n\}}\right) + \tilde{\mathbf{J}}_{\mathbf{nl}}^{\{n\}}\,\mathbf{\Delta p_{nl}}^{\{n\}}, \tag{C.13}$$

where the step of the non-linear parameters is $\mathbf{\Delta p_{nl}}^{\{n\}} = \mathbf{p_{nl}}^{\{n+1\}} - \mathbf{p_{nl}}^{\{n\}}$.
The Jacobian matrix $\tilde{\mathbf{J}}_{\mathbf{nl}}^{\{n\}}$ encompasses both the derivatives of f with respect to $\mathbf{p_{nl}}^{\{n\}}$ and $\mathbf{p_l}^{\{n\}}$

$$\tilde{\mathbf{J}}_{\mathbf{nl}}^{\{n\}} = \frac{\partial f\left(\mathbf{t}, \mathbf{p}^{\{n\}}\right)}{\partial \mathbf{p_{nl}}^{\{n\}}} + \frac{\partial f\left(\mathbf{t}, \mathbf{p}^{\{n\}}\right)}{\partial \mathbf{p_l}^{\{n\}}}\frac{\partial \mathbf{p_l}^{\{n\}}}{\partial \mathbf{p_{nl}}^{\{n\}}} \tag{C.14}$$
$$= \mathbf{J_{nl}}^{\{n\}} + \mathbf{J_l}^{\{n\}}\mathbf{A}^{\{n\}}. \tag{C.15}$$

$\mathbf{A}^{\{n\}} = \frac{\partial \mathbf{p_l}^{\{n\}}}{\partial \mathbf{p_{nl}}^{\{n\}}}$ is a 3×3 matrix containing the derivatives of the linear parameters $\mathbf{p_l}^{\{n\}}$ with respect to the non-linear ones $\mathbf{p_{nl}}^{\{n\}}$. It is obtained by computing the derivative of Eq. (C.10) with respect to $\mathbf{p_{nl}}^{\{n\}}$

$$\mathbf{A}^{\{n\}} = \left[\mathbf{H_l}^{\{n\}}\right]^{-1}\left[\frac{\partial \mathbf{Z_l}^{\{n\}}}{\partial \mathbf{p_{nl}}^{\{n\}}} - \frac{\partial \mathbf{H_l}^{\{n\}}}{\partial \mathbf{p_{nl}}^{\{n\}}}\mathbf{p_l}^{\{n\}}\right]. \tag{C.16}$$

Using Eqs. (C.11) and (C.12), we deduce the derivatives as

$$\frac{\partial \mathbf{H_l}^{\{n\}}}{\partial \mathbf{p_{nl}}^{\{n\}}} = \frac{\partial \mathbf{J_l}^{\{n\}T}}{\partial \mathbf{p_{nl}}^{\{n\}}}\,\mathbf{W}\,\mathbf{J_l}^{\{n\}} + \mathbf{J_l}^{\{n\}T}\,\mathbf{W}\,\frac{\partial \mathbf{J_l}^{\{n\}}}{\partial \mathbf{p_{nl}}^{\{n\}}} \tag{C.17}$$
$$\frac{\partial \mathbf{Z_l}^{\{n\}}}{\partial \mathbf{p_{nl}}^{\{n\}}} = \frac{\partial \mathbf{J_l}^{\{n\}T}}{\partial \mathbf{p_{nl}}^{\{n\}}}\,\mathbf{W}\,\mathbf{RV}_{\rm obs}. \tag{C.18}$$

Therefore, the system (Eq. C.8) becomes

$$\left[\tilde{\mathbf{H}}_{\mathbf{nl}} + \lambda\,\mathrm{diag}\left(\tilde{\mathbf{H}}_{\mathbf{nl}}\right)\right]^{\{n\}}\mathbf{\Delta p_{nl}}^{\{n\}} = \tilde{\mathbf{J}}_{\mathbf{nl}}^{\{n\}T}\mathbf{W}\left[\mathbf{RV}_{\rm obs} - f\left(\mathbf{t}, \mathbf{p}^{\{n\}}\right)\right] \tag{C.19}$$

with $\tilde{\mathbf{H}}_{\mathbf{nl}} = \tilde{\mathbf{J}}_{\mathbf{nl}}^{\{n\}T}\mathbf{W}\,\tilde{\mathbf{J}}_{\mathbf{nl}}^{\{n\}}$ being the new hessian matrix of the system. We have then split the 6×6 system (Eq. C.8) into two smaller systems, each one of 3×3, in Eqs. (C.10) and (C.19), respectively.

Appendix C.3. Analytical expressions of the involved matrices

To perform the above computations, we need the expressions of the Jacobian matrices $\mathbf{J_l}^{\{n\}}$ and $\mathbf{J_{nl}}^{\{n\}}$. The former, relative to the linear system of Eq. (C.9) is given by

$$\mathbf{J_l}^{\{n\}} = [1, \cos(\mathbf{v}) + e, \sin(\mathbf{v})]^{\{n\}}. \tag{C.20}$$

Its derivatives with respect to the non-linear parameters are

$$\frac{\partial \mathbf{J_l}^{\{n\}}}{\partial \nu_k} = [0, -\sin(\mathbf{v}), \cos(\mathbf{v})]^{\{n\}}\,\frac{\partial \mathbf{v}^{\{n\}}}{\partial \nu_k}, \tag{C.21}$$
$$\frac{\partial \mathbf{J_l}^{\{n\}}}{\partial e} = \left[0, 1 - \sin(\mathbf{v})\,\frac{\partial \mathbf{v}}{\partial e}, \cos(\mathbf{v})\,\frac{\partial \mathbf{v}}{\partial e}\right]^{\{n\}}, \tag{C.22}$$
$$\frac{\partial \mathbf{J_l}^{\{n\}}}{\partial T_0} = [0, -\sin(\mathbf{v}), \cos(\mathbf{v})]^{\{n\}}\,\frac{\partial \mathbf{v}^{\{n\}}}{\partial T_0}. \tag{C.23}$$

Very similarly, the Jacobian matrix $\mathbf{J_{nl}}^{\{n\}}$ is given by

$$\frac{\partial \mathbf{J_{nl}}^{\{n\}}}{\partial \nu_k} = [0, -\alpha\,\sin(\mathbf{v}), \beta\,\cos(\mathbf{v})]^{\{n\}}\,\frac{\partial \mathbf{v}^{\{n\}}}{\partial \nu_k}, \tag{C.24}$$
$$\frac{\partial \mathbf{J_{nl}}^{\{n\}}}{\partial e} = \left[0, \alpha\left(1 - \sin(\mathbf{v})\,\frac{\partial \mathbf{v}}{\partial e}\right), \beta\,\cos(\mathbf{v})\,\frac{\partial \mathbf{v}}{\partial e}\right]^{\{n\}}, \tag{C.25}$$

$$\frac{\partial \mathbf{J_{nl}}^{\{n\}}}{\partial T_0} = [0, -\alpha \sin(\mathbf{v}), \beta \cos(\mathbf{v})]^{\{n\}} \frac{\partial \mathbf{v}^{\{n\}}}{\partial T_0}. \tag{C.26}$$

The derivatives of the true anomalies $\mathbf{v}$ with respect to the three non-linear parameters are

$$\frac{\partial \mathbf{v}}{\partial \nu_k} = 2\pi (\mathbf{t} - T_0)(1 + e\cos(\mathbf{v}))^2 / (1 - e^2)^{3/2}, \tag{C.27}$$

$$\frac{\partial \mathbf{v}}{\partial e} = \sin(\mathbf{v})(2 + e\cos(\mathbf{v})) / (1 - e^2), \tag{C.28}$$

$$\frac{\partial \mathbf{v}}{\partial T_0} = -2\pi \nu_k (1 + e\cos(\mathbf{v}))^2 / (1 - e^2)^{3/2}. \tag{C.29}$$

Appendix C.4. Computation steps

The parameter step $\mathbf{\Delta p_{nl}}^{\{n\}}$ in Eq. C.19 is computed as follows:

1. Given the preliminary non-linear parameters $\mathbf{p_{nl}}$, we compute the mean and eccentric anomalies given by Eq. 4. This transcendental equation is solved following the algorithm proposed by Fukushima (1996), and implemented in a common *Gaia* tool library.
2. The true anomaly trigonometric functions are computed through Eq. 2 and Eq. 3.
3. The Jacobian matrix $\mathbf{J_l}^{\{n\}}$ is filled given Eq. C.20. The linear parameters $\mathbf{p_l}^{\{n\}}$ are then derived by the linear least-square regression given by Eq. C.10.
4. The derivatives of the true anomalies are computed using Eq. C.27, Eq. C.28, and Eq. C.29.
5. We fill the Jacobian matrix $\mathbf{J_{nl}}^{\{n\}}$ following Eq. C.24, Eq. C.25, and Eq. C.26.
6. The matrix $\mathbf{A}$ is computed thanks to Eq. C.17, Eq. C.18, and Eq. C.16.
7. Finally, the total Jacobian matrix $\tilde{\mathbf{J}}_{\mathbf{nl}}^{\{n\}}$ is obtained by Eq. C.15, and the linear system (Eq. C.19) is solved for $\mathbf{\Delta p_{nl}}^{\{n\}}$.