

PSR J0537–6910: Exponential recoveries detected for 12 glitches

E. Zubieta^{1,2,*}, C. M. Espinoza^{3,4}, D. Antonopoulou⁵, W. C. G. Ho⁶, L. Kuiper⁷,
F. García^{1,2}, and S. del Palacio^{1,8}

¹ Instituto Argentino de Radioastronomía (CCT La Plata, CONICET; CICPBA; UNLP), C.C.5, (1894), Villa Elisa, Buenos Aires, Argentina

² Facultad de Ciencias Astronómicas y Geofísicas, Universidad Nacional de La Plata, Paseo del Bosque, B1900FWA La Plata, Argentina

³ Departamento de Física, Universidad de Santiago de Chile (USACH), Av. Víctor Jara 3493, Estación Central, Chile

⁴ Center for Interdisciplinary Research in Astrophysics and Space Sciences (CIRAS), Universidad de Santiago de Chile, Av. Víctor Jara 3493, Estación Central, Chile

⁵ Jodrell Bank Centre for Astrophysics, Department of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, UK

⁶ Department of Physics and Astronomy, Haverford College, 370 Lancaster Avenue, Haverford, PA 19041, USA

⁷ SRON – Space Research Organisation Netherlands, Niels Bohrweg 4, 2333 CA Leiden, The Netherlands

⁸ Department of Space, Earth and Environment, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden

Received 3 October 2025 / Accepted 8 December 2025

ABSTRACT

Context. Pulsar rotational glitches are unresolved increments of the rotation rate that sometimes trigger an enhancement of the spin-down rate. On occasion, the augmented spin-down decays gradually over days or months in an exponential manner. This is observed particularly after the largest glitch events. Glitches and their exponential recoveries are attributed to the presence of a neutron superfluid inside pulsars. The young pulsar PSR J0537–6910 exhibits the highest known glitching rate, with 60 events detected over nearly 18 years of monitoring.

Aims. Despite most PSR J0537–6910 glitches being large, only one exponential recovery has been reported, following its first discovered glitch. This is puzzling as pulsars of similar age, rotational properties, and glitching behaviour typically present significant exponential recoveries. We wish to understand if this is an intrinsic difference of PSR J0537–6910 or a detectability issue due, for example, to its high glitch frequency.

Methods. The full dataset, including recent NICER observations, was systematically searched for evidence of exponential relaxations. Each glitch was first tested for evidence of recovery across a broad range of trial timescales. Good candidates were further investigated by identifying the best recovery models with and without an exponential term, then comparing them using Bayesian evidence.

Results. We discovered six new glitches, bringing the total number of glitches in this pulsar to 66. Our search and selection criteria strongly indicate the presence of 11 additional, previously undetected, exponential recoveries. We provide updated glitch and timing solutions. Exponential recoveries were detected only for the largest glitches ($\Delta\nu > 20 \mu\text{Hz}$), though not for all of them. The inferred exponential timescales vary between 4 and 37 d, with the decaying frequency increment being close to or below 1% of the total in general. We find that the second derivative of the spin frequency ($\ddot{\nu}$) can remain fairly stable across several glitches, and only some glitches are associated with persistent $\ddot{\nu}$ changes. In particular, $\ddot{\nu}$ tends to take its lowest values after glitches with exponential recoveries, corresponding to inter-glitch braking indices $n_{\text{ig}} \sim 6\text{--}9$. On the other hand, following glitches without exponential recoveries – even large ones – $\ddot{\nu}$ tends to be higher, with varied values that typically lead to braking indices above 9 (inferred n_{ig} mostly clustered between 10 and 35).

Key words. methods: observational – pulsars: general – pulsars: individual: PSR J0537-6910

1. Introduction

Pulsars are rapidly rotating neutron stars that emit radiation across multiple wavelengths, from radio to gamma rays. They have extremely high and stable moments of inertia, which give them remarkably steady rotations. This places them among the most accurate clocks in the Universe (Hobbs et al. 2012).

However, the smooth rotation of pulsars can sometimes be interrupted by sudden increases in rotational frequency ν , called glitches (Espinoza et al. 2011; Manchester 2018; Zubieta et al. 2023). In addition, these spin-up episodes often exhibit a sudden change in the pulsar spin-down rate, which then recovers, sometimes partially in an exponential-like manner (Yu et al. 2013; Liu et al. 2024). Currently, around 700 glitches have been

detected in more than 200 pulsars¹, with a wide range of amplitudes: observed frequency changes $\Delta\nu$ vary from $\sim 10^{-4} \mu\text{Hz}$ to nearly $100 \mu\text{Hz}$ (Basu et al. 2022).

Glitches are believed to be the manifestation of a sudden angular momentum transfer between a loosely coupled superfluid interior and the rest of the star (Baym et al. 1969; Anderson & Itoh 1975; Haskell & Melatos 2015; Antonopoulou et al. 2022; Zhou et al. 2022). The neutrons in the inner crust and core of neutron stars are in a superfluid state, which is characterised by irrotational flow with vanishing friction. Inside fast-rotating pulsars, however, it is energetically favourable for the superfluid to follow the rotation of its normal surroundings whilst remaining irrotational in its bulk. It does so by forming many microscopic vortices with non-superfluid

* Corresponding author: ezubieta@iar.unlp.edu.ar

¹ <http://www.jb.man.ac.uk/pulsar/glitches.html>

cores, which carry all circulation and whose density determines a local, average rotation rate. As the pulsar spins down, vortices move outwards and eventually vanish at the superfluid boundary, so that the superfluid decelerates as well. If in some stellar layer this vortex motion is prevented by forces that ‘pin’ vortices to the normal matter, excess vorticity builds up because the pulsar spins down whilst the local superfluid maintains a higher rotation rate. Eventually, either due to the increasing velocity difference between the components or aided by some external disturbance, a catastrophic unpinning event takes place: vortices move, quickly transferring the excess angular momentum to the crust, which spins-up at a glitch.

Weakly coupled superfluid regions lag behind the abrupt glitch-induced increase in the crust’s rotation rate and temporarily decouple, reducing the effective moment of inertia upon which the external spin-down torque acts. This causes a transient increase in the spin-down rate $|\dot{\nu}|$, commonly observed to accompany glitches, with typical amplitudes up to $\Delta\dot{\nu} \sim 0.1\% \dot{\nu}$. In a considerable minority, mostly consisting of large glitches in young pulsars, part of this change exponentially decays over characteristic relaxation timescales of days to several weeks: a signature of superfluid components slowly recoupling to the crust (Haskell & Melatos 2015; Haskell et al. 2020; Antonopoulou et al. 2022). For well-monitored neutron stars with regular large glitches, such as the Vela pulsar, detected decaying timescales range from minutes (Dodson et al. 2002; Palfreyman et al. 2018; Ashton et al. 2019; Zubieta et al. 2025b) to ~ 500 d (Zubieta et al. 2024). The study of glitches, and particularly their relaxation process, is key to understanding the internal physics of neutron stars.

PSR J0537–6910 stands out as the fastest spinning ($\nu \sim 62$ Hz), non-recycled, rotation-powered pulsar with the highest known spin-down energy loss rate ($\dot{E} \sim 4.88 \times 10^{38}$ erg s $^{-1}$). It is a young object, associated with the 1–5 kyr old supernova remnant N157B in the Large Magellanic Cloud (Wang & Gotthelf 1998; Marshall et al. 1998), and it is by far the most frequently glitching pulsar with a rate of approximately three events per year and a total of 60 reported glitches in nearly 18 years of cumulated data (Middleitch et al. 2006; Antonopoulou et al. 2018; Ferdman et al. 2018; Ho et al. 2022). About 70% of these glitches are large, with $\Delta\nu > 10$ μ Hz (90% above 1 μ Hz), which contribute to the peak at large values observed in the bimodal distribution of glitch sizes for all glitching pulsars (Fuentes et al. 2017).

Glitches in PSR J0537–6910 are accompanied by significant increases in the spin-down rate $|\dot{\nu}|$, which dominate $\dot{\nu}$ -time evolution (Fig. C.1), a common feature of glitching pulsars of similar age, such as Vela (Lyne et al. 2015; Espinoza et al. 2017). On shorter timescales, these pulsars exhibit inter-glitch $\dot{\nu}$ values that are considerably larger than the predictions of any model for pulsar braking. For example, the dipole braking model, in which the neutron star loses energy exclusively by magnetic dipole radiation, assuming constant moment of inertia and magnetic field, predicts a $\dot{\nu}$ rate that is 20 times lower than the observed values between glitches in PSR J0537–6910. Contrary to similar (glitching) pulsars, however, PSR J0537–6910 does not feature prominent exponential relaxations, and their presence has been somewhat under-explored (Ho et al. 2020). The first glitch observed was the largest and the only one for which an exponential recovery was clearly discernible and well-studied (Antonopoulou et al. 2018).

In this work, we present six new glitches detected in recent NICER (*Neutron star Interior Composition Explorer*) observations of PSR J0537–6910 and analyse all known glitches in

the entire dataset for the presence of exponential recoveries. In Sect. 2, we summarise the reduction process of the X-ray data from NICER and RXTE (*Rossi X-ray Timing Explorer*). In Sect. 3, we develop a systematic methodology to identify and characterise post-glitch exponential recoveries. In Sect. 4, we provide a report on the six new glitches and the results of our search for exponential recoveries on the 66 total glitches observed so far. Finally, in Sect. 5, we discuss the implications of our findings.

2. Observations

PSR J0537–6910 was discovered in the X-ray band in 1998 (Marshall et al. 1998), and only recently some weak and dispersed radio pulses were detected (Crawford et al. 2005; Crawford 2024). Monitoring its rotation with RXTE began on 1999 January 19 and lasted until 2011 December 31. Regular observations resumed 5.45 years later with NICER (Ho et al. 2020). In this work, we used all available RXTE and NICER data for this pulsar and calculated pulse times of arrival (TOAs) by cross-correlating pulse profiles with a pulse template. The RXTE TOAs used in our analysis are the same as those in Antonopoulou et al. (2018), derived using a method similar to that presented by Kuiper & Hermsen (2009) for PSR J1846–0258.

However, the template used for correlating with the folded profiles to obtain a TOA measurement comes from fitting a 60-bin asymmetric Lorentzian to a high signal-to-noise profile. In some cases, consecutive observations from RXTE were combined to achieve sufficient exposure for confidently measuring the spin frequency of the pulsar. This occurred mainly with observations performed after 2006, where generally only one or two of the five proportional counter units (PCUs) were operational during a standard observation.

For the NICER data, the methodology was laid out in Ho et al. (2020): HEASOFT 6.22–6.35.1 (Nasa High Energy Astrophysics Science Archive Research Center (Heasarc) 2014) and NICERDAS 2018-03-01_V003-2025- 03-11_V013A were used to process and filter the observations. We only kept events in the 1–7 keV range, where the pulse is easily detected (Kuiper & Hermsen 2015). As for RXTE data, there were some cases for which sets of close individual observations had to be merged to significantly detect the pulsar period. The typical total exposure of merged observations ranged between 4 and 9 ks. From each merged observation, one TOA was obtained. In this case, TOAs were measured by correlating the folded profiles with a template obtained from fitting a Gaussian to a series of NICER high signal-to-noise pulse profiles.

3. Methods

In the following, we discuss the 45 glitches published by Antonopoulou et al. (2018) for the complete RXTE observing period (MJD 51197–55926), plus 15 reported glitches from NICER observations, which we enumerate from 46 onward. Timing solutions for glitches 46–53 can be found in Ho et al. (2020) (glitches 1–7 in their work); glitches 54–56 were presented in Abbott et al. (2021) (glitches 8–10 in their work); and glitches 57–60 in Ho et al. (2022) (11–15 in their work).

We processed the latest NICER data (after MJD 59627) and report here on six new glitches that we identified and measured. Therefore, a total of 66 PSR J0537–6910 glitches were included in our post-glitch recovery analysis.

3.1. Pulsar timing

The pulsar rotation is monitored by recording TOAs of detected pulses and comparing them to the predictions of a spin-down model for the rotation phase ϕ , which can be approximated by a truncated Taylor series:

$$\phi(t) = \phi + \nu_0(t - t_0) + \frac{1}{2}\dot{\nu}_0(t - t_0)^2 + \frac{1}{6}\ddot{\nu}_0(t - t_0)^3, \quad (1)$$

where ν_0 , $\dot{\nu}_0$, and $\ddot{\nu}_0$ are the rotation frequency of the pulsar and its first and second derivatives, respectively, at the reference epoch t_0 . The differences between the observed and predicted arrival times are called timing residuals. These can be remarkably small, thanks to the accuracy of TOA measurements and the stability of the pulsar's rotation. Glitches introduce well-known signals in the residuals and, unless they are too small, are generally straightforward to identify.

To describe the changes in the rotation of the pulsar due to a glitch at time t_g and its post-glitch recovery, an additional phase ϕ_g must be added to the model of Eq. (1) for $t > t_g$, of the form (e.g. Yu et al. 2013; Zubieta et al. 2025a):

$$\begin{aligned} \phi_g(t) = & \Delta\phi + \Delta\nu_p(t - t_g) + \frac{1}{2}\Delta\dot{\nu}_p(t - t_g)^2 \\ & + \frac{1}{6}\Delta\ddot{\nu}_p(t - t_g)^3 + \left[1 - \exp\left(-\frac{t - t_g}{\tau_d}\right) \right] \Delta\nu_d \tau_d. \end{aligned} \quad (2)$$

The quantities $\Delta\nu_p$, $\Delta\dot{\nu}_p$, and $\Delta\ddot{\nu}_p$ describe persisting changes in the rotational parameters with respect to the pre-glitch model, whilst a non-zero $\Delta\phi$ might be required due to the uncertainty of the precise glitch epoch t_g . The last term describes a transient component of the spin-up, $\Delta\nu_d$, which exponentially relaxes on a timescale τ_d . Glitches can have one or more such exponential components in their recoveries; however, in this study of PSR J0537–6910 we examine glitch models with a maximum of one exponential term. We computed the total, unresolved initial frequency step as $\Delta\nu = \Delta\nu_p + \Delta\nu_d$ and its recovery factor as $Q = \Delta\nu_d/\Delta\nu$. The equivalent change in frequency derivative is $\Delta\dot{\nu} = \Delta\dot{\nu}_p - (\Delta\nu_d/\tau_d)$, and it is useful to also define the spin-down rate recovery factor as $Q' = \Delta\nu_d/(\Delta\nu_d - \tau_d\Delta\dot{\nu}_p)$ (Weltevrede et al. 2011).

3.2. Post-glitch recovery searches

We analysed each glitch (n) using TOAs between the previous glitch ($n - 1$) and the following one ($n + 1$). Within this data span, we used TEMPO2 to fit Eq. (1) to the pre-glitch TOAs. If an exponential recovery was detected for glitch $n - 1$, we excluded any TOAs clearly affected by it from the pre-glitch data span used to characterise the following glitch n . This ensured that the pre-glitch- n data could be accurately described solely by Eq. (1). In general, PSR J0537–6910 has a high inter-glitch $\dot{\nu}$, which can be determined even for short time intervals. We therefore included this term in the fits except in cases of exceptionally short (four or fewer TOAs) pre-glitch datasets, for which we set $\dot{\nu} = 1 \times 10^{-20} \text{ Hz s}^{-2}$ (Ho et al. 2020).

Indications of exponential signatures in the post-glitch recovery were searched for automatically in each inter-glitch dataset comprising of at least six TOAs. The first step of our methodology was to make an initial estimate of the glitch size by fitting a glitch model with only $\Delta\phi$, $\Delta\nu_p$, and $\Delta\dot{\nu}_p$ (maintaining $\Delta\nu_d = \Delta\ddot{\nu}_p = 0$). Glitch parameters from the literature were used as a starting solution, when available. In the next step, the full glitch model of Eq. (2) was fitted for, with $\Delta\nu_d$

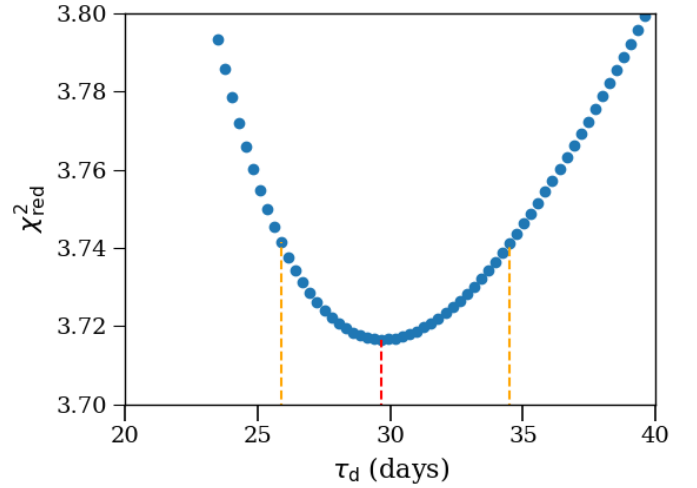


Fig. 1. Example of the exploratory process to determine possible values for an exponential relaxation timescale. The blue dots corresponds to the χ_{red}^2 values of glitch 64, over a subset of the total range of trial τ_d values focused around the timescale that returns the minimum χ_{red}^2 (marked by the red line). The dashed orange lines denote the 1σ uncertainty. These rather high χ_{red}^2 values are used only to identify a possible timescale. Further analyses are performed for the final measurements (Table 2).

and all other parameters allowed to vary except for τ_d , which was set at a trial value. For each post-glitch dataset, we performed such fits for a range of possible relaxation timescales: all integer values between one day and four times the total length of the post-glitch data span were used as trials for τ_d . Exponential recoveries on shorter timescales could not be accurately determined due to the timing resolution of the dataset. Similarly, the exponential signature was hard to discern if the relaxation timescale was considerably longer than the post-glitch fitted interval.

The variations in the χ_{red}^2 metric for each glitch fit with respect to τ_d were analysed. When a minimum was present, the search was refined around it, as exemplified in Fig. 1.

This was not a final determination of glitch recovery but rather a preliminary exploration to identify whether the data favour a possible recovery timescale. Glitches for which the minimum χ_{red}^2 occurred at either the beginning or the end of the explored τ_d range were not further explored. We assume in such cases that the absolute minimum might occur outside our trial τ_d range and therefore, as discussed above, we could not securely detect and characterise an exponential recovery with the available data.

For the final step, we considered two different glitch models to compare via marginal likelihood:

- Model 1 (without recovery) consists of $\Delta\phi$, $\Delta\nu_p$, and $\Delta\dot{\nu}_p$. We considered $\Delta\dot{\nu}_p$ when its inclusion led to a significant decrease in the residuals' RMS.
- Model 2 (with one exponential recovery) consists of $\Delta\phi$, $\Delta\nu_p$, $\Delta\dot{\nu}_p$, $\Delta\nu_d$, and τ_d . We incorporated $\Delta\dot{\nu}_p$ when its inclusion led to a significant decrease in the residuals' RMS. We used the best-fit recovery parameters from the previous step as initial values for these fits.

A nested sampling code² (Skilling 2004) based on the PINT software (Luo et al. 2021) was used to fit the models and infer the

² <https://github.com/kbarbary/nestle>

most likely values for the glitch parameters. We explored a wide range for each parameter to ensure that a global solution was reached.

Finally, to decide the appropriate model for each glitch, we calculated their Bayesian evidences, Z_1 and Z_2 , for Models 1 and 2, respectively, and assessed their logarithmic ratio:

$$\ln B_{21} = \ln Z_2 - \ln Z_1, \quad (3)$$

following the criterion used in Liu et al. (2024): if $\ln B_{21} < 2.5$, Model 2 is considered to have insufficient evidence over Model 1. In the case of $2.5 < \ln B_{21} < 5$, Model 2 is considered to have moderate evidence over Model 1. For $5 < \ln B_{21} < 10$, strong evidence for Model 2 over Model 1 is considered. However, if $\ln B_{21} > 10$, evidence for Model 2 is considered as decisive.

We treated all glitches for which the evidence for Model 2 was moderate, strong, or decisive as detections of exponential relaxations. In some cases, the posterior distributions for certain parameters are considerably asymmetric, and therefore we used the posterior mode – instead of the mean value – as the representative value for each parameter. From these, we constructed and used a timing solution for each glitch as a starting solution for a final fit, obtained using PINT. In most cases, the final fit only marginally reduced the RMS of the residuals compared to the starting solution. The difference between the modes and the values obtained after this fit was consistently smaller than the range defined by the 68% of the most likely values in the posterior distributions, even for the most asymmetric posteriors. We assigned final parameter uncertainties calculated as the distance from the mode to the 16th and 84th percentiles of its posterior samples.

4. Results

4.1. Detection of six new glitches with NICER

In the timing residuals of the most recent NICER data used in this analysis (MJD 59626–60601), the presence of six new glitches is evident – and in accordance with the known glitching rate of this pulsar. We present the parameters that describe these new glitches in Table 1.

To be consistent with earlier measuring techniques and to ensure the results are directly comparable to previously published glitch parameters, we followed Ho et al. (2020) and used TEMPO2 (Hobbs et al. 2006) to fit Eqs. (1) and (2) without the exponential relaxation term to the TOAs of each pre-glitch and post-glitch interval (described in Table A.1). A seventh new glitch (67) occurred around MJD 60592. Therefore, we include glitch 66 in the analysis but not glitch 67, since the post-glitch data were incomplete at the time of this analysis (see Ho et al. 2026).

Five out of the six new glitches have inferred amplitudes $\Delta\nu_p > 15 \mu\text{Hz}$, and all had detectable changes in $\dot{\nu}$. As we discuss next, two of these glitches present evidence for exponential relaxation in their post-glitch recovery. Refined solutions that include this term are presented in the following subsection.

4.2. Detection of exponential post-glitch relaxations

We analysed a total of 66 PSR J0537–6910 glitches following the methods outlined in Sect. 3.2, starting from initial solutions presented in Antonopoulou et al. (2018), Ho et al. (2020), Abbott et al. (2021), Ho et al. (2022) and Table 1 of this work. We found 12 glitches for which the Bayesian criterion points to the presence of an exponential recovery, with moderate evidence

in two cases and strong or decisive evidence in the other ten. These glitches are presented in Table A.2, together with information on the chosen solution for Models 1 and 2, as well as their Bayesian evidences.

Some glitches yielded a clear τ_d value during the initial systematic search, but their Bayesian evidence was not high enough to favour Model 2 over Model 1. This is likely due to the sparsity of post-glitch TOAs in these particular cases. Nonetheless, in these cases where Model 2 was favoured, the accuracy of the inferred exponential solutions does not appear strongly sensitive to variations in data quality. No clear correlation exists between the uncertainties of the inferred exponential parameters ($\Delta\nu_d$ and τ_d) and the TOA cadence or TOA errors of the respective post-glitch interval.

As noted in Table A.2, for some glitches Model 1 required a $\Delta\dot{\nu}_p \neq 0$ in order to produce reasonably flat residuals; however, when the recovery term (Model 2) was added, the change $\Delta\dot{\nu}_p$ became negligible. There were also two cases (glitches 43 and 45), for which including a $\Delta\dot{\nu}_p$ term did not significantly flatten the residuals for Model 1; nonetheless, the exponential recovery term in Model 2 did. Finally, there were four cases for which both $\Delta\dot{\nu}_p$ and a recovery term were required to significantly reduce the RMS of the residuals. In the model of the first glitch, we omitted $\Delta\dot{\nu}_p$, as we obtained only a few pre-glitch TOAs and therefore could not reliably determine $\dot{\nu}$ before the glitch. Similarly, $\Delta\dot{\nu}_p$ was not included for glitch 47, as only three pre-glitch TOAs were available. In this case, $\dot{\nu}$ was fitted to both pre- and post-glitch data.

The final solutions of Model 2 for the 12 glitches with evidence of an exponential term in their post-glitch recovery are shown in Table 2. The description of the observations used for each fit is presented in Table A.3. It is worth mentioning that for glitches 12 and 47, a fraction of the posterior samples of $\Delta\nu_d$ is consistent with 0, as indicated by their respective uncertainties. Nevertheless, we consider them significant detections because the modes of the posteriors are far from zero, and the Bayesian criterion favours an exponential term. These two glitches have short inferred τ_d values, so the lack of TOAs immediately after the glitch hampers their accurate measurement, as reflected by their uncertainties. The shorter (~ 3 d) recovery timescales seen in Table 2 emphasise the need to monitor PSR J0537–6910 with high resolution immediately after its glitches, to detect and resolve such fast transients. Such observations will also improve our detection capability for exponentials of any timescale.

4.3. The exponential post-glitch relaxations

In Fig. 2, we show two of the clearest detections of exponential relaxations in PSR J0537–6910. For completeness, in Fig. B.1 we also show two large glitches for which we found no evidence of exponential signatures. Plots for all 12 glitches are available in a GitHub repository³.

The exponentially relaxing component represents only a small fraction of the instantaneous change in spin frequency at the glitch. The inferred decaying amplitudes $\Delta\nu_d$ range between 0.05 and 0.8 μHz , while the persisting changes $\Delta\nu_p$ lie between 20 and 43 μHz . Therefore, the percentage of the initial frequency step that decayed exponentially, Q , is $< 4\%$ for all glitches. Most often, the calculated Q is close to or below 1%, which is typical of giant glitches in other pulsars (see, for instance, Zubieta et al. 2023).

³ <https://github.com/ezubieta/PSR-J0537-6910-recoveries>

Table 1. Timing solutions (without exponential recovery terms) for the six new glitches detected with NICER.

Glitch	t_0 (MJD)	t_g (MJD)	ν (Hz)	$\dot{\nu}$ (10^{-10} Hz s $^{-1}$)	$\ddot{\nu}$ (10^{-20} Hz s $^{-2}$)	$\Delta\phi$ (cycle)	$\Delta\nu_p$ (μ Hz)	$\Delta\dot{\nu}_p$ (10^{-13} Hz)	$\Delta\ddot{\nu}_p$ (10^{-20} Hz s $^{-2}$)
61	59607	59692(6)	61.8972423813(5)	-1.997915(3)	0.80(8)	-0.18(2)	5.83(1)	-0.86(6)	-
62	59716	59741(6)	61.895366836(4)	-1.99802(6)	1 \dagger	-0.42(2)	15.40(2)	-1.06(7)	0.56(4)
63	59813	59886(5)	61.893707533(1)	-1.997897(5)	1.56(2)	0.10(1)	24.375(9)	-1.77(4)	-0.71(5)
64	59958	60028(2)	61.891228926(7)	-1.998156(5)	0.84(2)	1.89(3)	34.22(2)	-1.43(6)	-2(6)
65	60150	60222(2)	61.887948098(4)	-1.998211(2)	0.58(2)	0.223(7)	21.105(9)	-7.22(9)	-
66	60290	60379(22)	61.885552197(1)	-1.998404(5)	0.90(2)	0.39(6)	31.49(3)	-1.22(9)	-1.3(8)

Notes. Solutions and their uncertainties were obtained by fitting the TOAs using TEMPO2. Glitch numbers follow the counting of RXTE glitches (Antonopoulou et al. 2018); to recover the NICER count, subtract 45. The dagger symbol (\dagger) indicates that $\dot{\nu}$ was fixed at that value. The numbers in parenthesis are the uncertainties in the last digit.

Table 2. Solutions for glitches exhibiting exponential recovery.

Glitch	t_g (MJD)	$\Delta\phi$ (cycle)	$\Delta\nu_p$ (μ Hz)	$\Delta\dot{\nu}_p$ (10^{-13} Hz s $^{-1}$)	$\Delta\ddot{\nu}_p$ (10^{-20} Hz s $^{-2}$)	τ_d (days)	$\Delta\nu_d$ (μ Hz)	Q (%)	Q' (%)
1	51278(16)	$0.5^{+0.1}_{-0.1}$	$42.27^{+0.02}_{-0.02}$	$-0.70^{+0.04}_{-0.04}$	-	21^{+5}_{-5}	$0.3^{+0.1}_{-0.1}$	$0.8^{+0.2}_{-0.2}$	71^{+8}_{-8}
2	51562(15)	$0.6^{+0.1}_{-0.1}$	$27.8^{+0.1}_{-0.1}$	$-1.2^{+0.1}_{-0.1}$	-	31^{+28}_{-12}	$0.3^{+0.1}_{-0.1}$	$1.0^{+0.5}_{-0.4}$	46^{+16}_{-24}
6	51960(5)	$0.0^{+0.1}_{-0.2}$	$28.14^{+0.03}_{-0.02}$	$-1.5^{+0.1}_{-0.1}$	$-2.2^{+0.3}_{-0.3}$	8^{+2}_{-5}	$0.2^{+0.6}_{-0.1}$	$0.8^{+2.2}_{-0.4}$	69^{+61}_{-12}
12	52545(6)	$0.57^{+0.06}_{-0.06}$	$26.09^{+0.01}_{-0.01}$	$-0.73^{+0.03}_{-0.04}$	-	7^{+7}_{-5}	$0.05^{+0.07}_{-0.06}$	$0.2^{+0.2}_{-0.2}$	97^{+2}_{-7}
38	55280(4)	$0.5^{+0.1}_{-0.2}$	$34.02^{+0.02}_{-0.02}$	$-0.4^{+0.1}_{-0.1}$	-	4^{+4}_{-1}	$0.8^{+1.1}_{-0.5}$	2^{+3}_{-1}	98^{+2}_{-1}
43	55615(4)	$0.1^{+0.1}_{-0.1}$	$28.0^{+0.1}_{-0.1}$	-	-	18^{+14}_{-8}	$0.2^{+0.1}_{-0.1}$	$0.8^{+0.4}_{-0.3}$	100^{+7}_{-12}
45	55819(2)	$-0.2^{+0.1}_{-0.1}$	$20.7^{+0.3}_{-0.5}$	-	-	37^{+30}_{-10}	$0.8^{+0.5}_{-0.2}$	4^{+2}_{-1}	100^{+7}_{-21}
47	58152(11)	$0.2^{+0.4}_{-0.2}$	$36.02^{+0.07}_{-0.09}$	$-1.5^{+0.3}_{-0.3}$	-	6^{+7}_{-2}	$0.5^{+0.6}_{-0.7}$	1^{+2}_{-2}	87^{+26}_{-24}
53	58868(5)	$-0.0^{+0.2}_{-0.1}$	$23.96^{+0.07}_{-0.09}$	$-2.1^{+0.4}_{-0.3}$	$-5.0^{+0.6}_{-0.7}$	5^{+5}_{-1}	$0.5^{+0.1}_{-0.4}$	$2.0^{+0.8}_{-0.8}$	84^{+36}_{-23}
56	59103(6)	$0.4^{+0.1}_{-0.1}$	$33.81^{+0.02}_{-0.01}$	$-1.3^{+0.2}_{-0.2}$	$-3.1^{+0.2}_{-0.2}$	9^{+5}_{-3}	$0.2^{+0.4}_{-0.1}$	$0.6^{+1.1}_{-0.3}$	37^{+34}_{-20}
64	60028(1)	$0.83^{+0.07}_{-0.04}$	$34.0^{+0.1}_{-0.1}$	$-1.0^{+0.2}_{-0.1}$	$-0.4^{+0.1}_{-0.1}$	30^{+13}_{-8}	$0.4^{+0.1}_{-0.1}$	$1.1^{+0.4}_{-0.3}$	61^{+11}_{-12}
65	60223(3)	$0.22^{+0.05}_{-0.04}$	$21.1^{+0.1}_{-0.1}$	$-0.7^{+0.1}_{-0.1}$	-	33^{+18}_{-11}	$0.3^{+0.1}_{-0.1}$	$1.3^{+0.6}_{-0.4}$	58^{+14}_{-15}

Notes. The posteriors samples for these glitches can be found in a GitHub repository (<https://github.com/ezubieta/PSR-J0537-6910-recoveries>). The uncertainties define 68% of the most likely values in the posteriors. Glitch numbers follow the counting of RXTE glitches (Antonopoulou et al. 2018). Subtract 45 to recover the NICER count.

It is remarkable that among the glitches with an exponential recovery, two (43 and 45) did not exhibit persistent changes $\Delta\dot{\nu}_p$ nor $\Delta\ddot{\nu}_p$. For these glitches, the value of Q' reaches 100%, meaning $\dot{\nu}$ returned to its pre-glitch value once the exponential decay was completed. Similar high Q' values are obtained for glitches 12 and 38. This is somewhat unusual for Vela-like glitches but has been observed in other pulsars, such as the Crab.

Although the persistent $\Delta\nu_p$ value does not change significantly with the addition of the exponential recovery term, the overall glitch and timing solution is altered. To investigate how including the exponentially recovering terms affects the inferred persisting $\Delta\dot{\nu}_p$ and $\ddot{\nu}$, we compare their values as calculated by Model 1 (no exponential) and Model 2 (one exponential term) in Fig. 3. As shown, the persistent spin-down change $\Delta\dot{\nu}_p$ considerably decreases when exponential relaxation is included in the model. Particularly interesting are the cases of glitches 43 and 45, discussed above. For these, Model 1 required a $\Delta\dot{\nu}_p \neq 0$ term, which becomes zero for Model 2. This shows that, in some cases, exponential signatures may be misinterpreted as changes $\Delta\dot{\nu}_p$ and $\Delta\ddot{\nu}_p$.

The nominal timescales of the retrieved solutions for Model 2 span from 4 to 37 d, but are often highly uncertain (as seen in Table 2). To better capture the inferred range of typical recov-

ery timescales of PSR J0537–6910, in Fig. 4 we show the distribution of τ_d values drawn from the posteriors of glitches with detected exponential recovery. There is some bimodality in the likely timescales, with a high peak around $\tau_d \sim 3$ d and a secondary one around $\tau_d \sim 20$ d. We note here that for those glitches where the inferred timescale was relatively long (glitches 1, 2, 43, 45, 64, and 65), we also tested a recovery model with two exponential terms but found no evidence for a second exponential. These results, which in general suggest short decaying timescales, intensify the need for dense observations as soon as possible after large glitches of PSR J0537–6910, in order to capture their short-term recovery behaviour.

Higher observing cadence close to a glitch will also allow us to place better limits on its size and improve the overall accuracy of all glitch parameters and timing solutions. Despite the small amplitude $\Delta\nu_d$ of the exponential relaxations compared to the total glitch change $\Delta\nu$, detecting and including these terms, when present, can affect several other results. This can be seen, for example, in Fig. 5, which displays the size $\Delta\nu$ of a glitch versus the length of the time interval ΔT_f until the following glitch. For PSR J0537–6910 these parameters are known to exhibit a strong, rather linear correlation (at least for intermediate to large glitches) (Middleditch et al. 2006; Antonopoulou et al. 2018;




Fig. 2. Examples of exponential glitch recovery detections, for glitches 1 (left) and 56 (right). From top to bottom: Phase residuals relative to Model 1, with the residuals of the most likely exponential recovery model (Model 2 relative to Model 1) superimposed (blue curve); phase residuals relative to Model 2; frequency residuals relative to Eq. (1) fitted to TOAs up to the glitch epoch, with the post-glitch data all lowered by a certain amount (the mean post-glitch frequency residual) for better visualisation; and $\dot{\nu}$ evolution with Model 2 shown as the blue curve.

Ho et al. 2020, 2022). Here, we study the correlation between these two parameters through the Spearman rank-order correlation coefficient (SPEARMANR from the SCIPY library). We first consider all glitch amplitudes obtained via Model 1 (no exponential term) and find a correlation coefficient $\rho_s = 0.89$ with a p -value $p = 1.5 \times 10^{-22}$. Interestingly, when we instead use the sizes $\Delta\nu$ retrieved by Model 2 for the 12 glitches with evidence of exponentials, the correlation measures improve to $\rho_s = 0.94$ and p -value of $p = 2.0 \times 10^{-32}$.

5. Discussion

Exponential recoveries were identified and measured following 12 glitches of PSR J0537–6910, all of which rank among the largest glitches observed. No evidence of exponential recovery was found for glitches with $\Delta\nu < 20 \mu\text{Hz}$. For greater glitch sizes, some – but not all – events show an exponential component, and only the largest glitches ($\Delta\nu \gtrsim 31.3 \mu\text{Hz}$) consistently exhibit detectable exponential terms. This behaviour can be seen in Fig. 5. By visually inspecting the timing residuals relative to a model – as in Eq. (2) – but without the exponential term, we confirm that no additional glitches of sizes between 20 and 30 μHz show obvious signs of exponential recovery. Nevertheless, not all residuals are adequately flat. For instance, the timing residuals for glitches 66 and 28, which are the first and second largest events without a detected recovery (with $\Delta\nu = 31.3 \mu\text{Hz}$ and $\Delta\nu = 30.3 \mu\text{Hz}$, respectively), are sufficiently flat, similar to what is observed for glitch 9 (the fourth largest glitch without a recovery, with $\Delta\nu = 26.4 \mu\text{Hz}$). However, glitch 51 (26.9 μHz), which is the third largest glitch without an exponential detec-

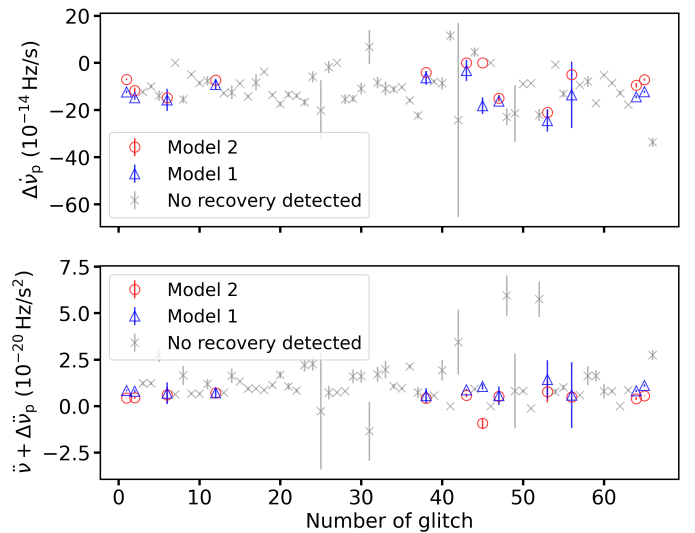


Fig. 3. Difference in inferred glitch parameters between Model 1 (blue points) and Model 2 (red points), for glitches with detected exponential recovery. Grey points represent values for glitches without detectable exponential terms (Model 1).

tion, exhibits noisier residuals relative to the same model. Other glitches within this size range behave similarly to these largest events, two of which are shown in Fig. B.1. We could not find any obvious reason – such as observational biases or failure to retrieve equally accurate solutions – for why exponentials were




Fig. 4. Distribution of τ_d values for the 12 glitches with a detected recovery, as represented by a histogram of 10 000 samples (weighted) from the posterior distributions. The red ticks correspond to the values yielded by the final solutions, as reported in Table 2.

not detected after these large glitches. Furthermore, the presence of a detectable exponential relaxation seems associated with lower inter-glitch $\dot{\nu}$ values, as discussed below. It is thus possible that the differences in post-glitch behaviour are intrinsic to the pulsar. Future observations can reveal whether that is the case and what it depends upon.

5.1. Inter-glitch spin-down rate evolution

When the favoured glitch model includes an exponential term, the post-glitch rotation typically evolves with a lower $\dot{\nu}$. This is not surprising at first, as all detections concern large glitches. These are followed by long (at least 100 d) inter-glitch intervals due to the aforementioned $\Delta\nu\text{-}\Delta T_f$ correlation, and it has been shown that the inter-glitch $\dot{\nu}$ tends to decrease over time from the glitch epoch (Andersson et al. 2018). To test this with our updated glitch sample, we fitted Eq. (1) to 90-day TOA intervals, starting after each glitch and moving forward by 20 days until the next glitch was reached. The resulting $\dot{\nu}$ values as a function of time since the preceding glitch are shown in Fig. 6. With all glitches collectively considered, $\dot{\nu}(t)$ indeed appears to decrease over time. However, beyond ~ 70 d the glitches without a detectable exponential have higher $\dot{\nu}$ (blue points) compared to glitches with exponential relaxation (red circles). We emphasise that the modelled exponential contributions were removed before calculating these $\dot{\nu}(t)$ values. If that had been the case, the red markers in Fig. 6 (circles and stars) would have been even lower, especially at early times. Therefore, we do not expect the blue points to be higher at late times due to unidentified exponentials of similar timescales as the detected recoveries. Although we cannot exclude contributions from undetected exponential relaxations on much longer timescales, given the current analysis and the broad range of trial timescales, this explanation seems considerably less likely than a different recovering regime.

This genuine difference between glitches with and without detected exponential terms can also be expressed via the braking index. Assuming a power law $\dot{\nu} \propto -\nu^n$ for the spin-down evolution, the braking index n can be calculated from observations as $n = \nu\ddot{\nu}/\dot{\nu}^2$. Typical inter-glitch braking indices n_{ig} in Vela-like pulsars are $n_{\text{ig}} = 10\text{--}40$, or higher. This is arguably the result of internal superfluid torques (Espinoza et al. 2017; Antonopoulou et al. 2018) rather than external magnetospheric torques, which are expected to give much lower values (e.g.

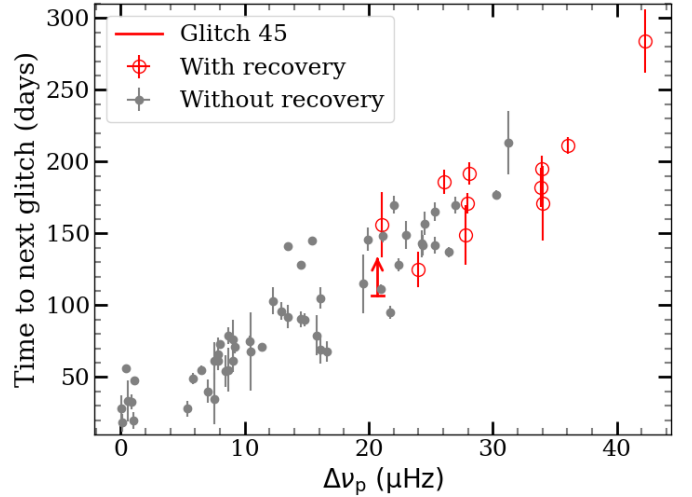


Fig. 5. Time to the next glitch as a function of $\Delta\nu_p$ for 66 glitches of PSR J0537–6910. Red markers indicate glitches for which an exponential recovery was detected. Using Model 2 solutions in these cases improves the size-waiting time correlation. Glitch 45, with a detected recovery, is the last glitch in the RXTE dataset, which ends 107 d after. Thus, the time to the next glitch is unknown, and it is plotted with a vertical arrow indicating the lower limit. The latest glitch reported in this work is glitch 66, which we are able to include in this plot as we know that glitch 67 occurred on MJD 60592.

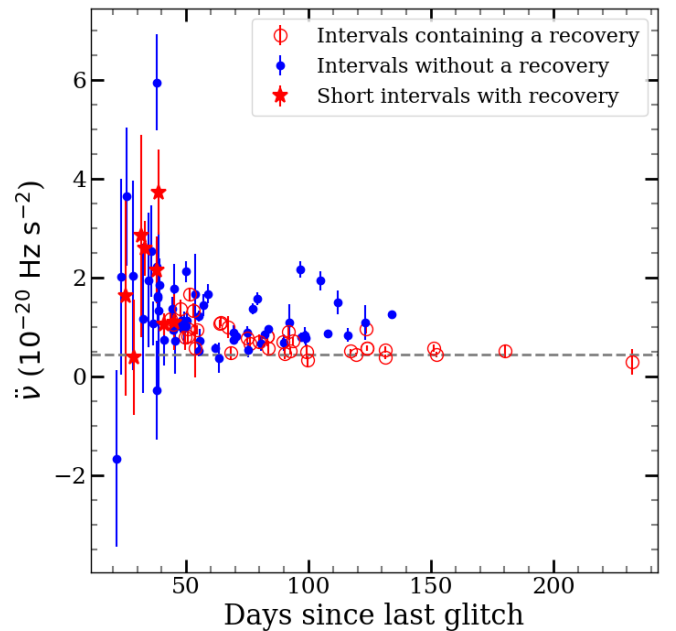


Fig. 6. $\dot{\nu}$ vs. time since the preceding glitch. The horizontal dashed line indicates $\dot{\nu}_{\text{ref}} = 0.44 \cdot 10^{-20} \text{ s}^{-3}$. The data points were obtained by fitting Eq. (1) over a 90-day sliding window, which moved forward 20 days at each step. The full inter-glitch data span was fitted if the inter-glitch interval was shorter than 90 days. The red stars were obtained by fitting smaller windows of at least 30 d and at least eight TOAs immediately after the glitches with a recovery were detected.

$n \sim 3$ in the standard magnetic dipole braking model). Whilst for most glitches of PSR J0537–6910, the inter-glitch n_{ig} is generally greater than 9 and as high as 35; after the glitches with detected exponential recoveries, it settles to lower values, between 6.2 and 9.4.




Fig. 7. Evolution of $\dot{\nu}$ corrected for the changes induced by the glitches. Vertical dotted lines indicate the epoch of the glitches. Note that the scales on the horizontal axes have different lengths.

To further investigate the $\dot{\nu}$ values between glitches and the overall rotational evolution, we studied $\dot{\nu}$ as a function of time. A $\dot{\nu}(t)$ dataset was constructed by fitting Eq. (1) to overlapping subsets of data containing at least five TOAs and spanning at least 30 d, but not through glitches. For these fits we kept $\dot{\nu}$ fixed at the mean value of each inter-glitch interval. The fitting window moved one TOA each time; hence, the data are heavily smoothed for visualisation purposes. In Appendix C we present the resulting $\dot{\nu}$ evolution, resembling a sawtooth curve typical of pulsars with repeating large glitches (e.g. Espinoza et al. 2017). In the case of PSR J0537–6910, $\dot{\nu}$ decreases over the years, with a corresponding effective negative slope and long-term negative effective braking index, a potential product of the high glitch activity (Antonopoulou et al. 2018; Ho et al. 2020). Between glitches, $\dot{\nu}$ increases rapidly: the inter-glitch evolution has high, positive $\ddot{\nu}$ values and $n_{\text{ig}} > 3$, which likely involves contributions from (at least) the exponentially and non-exponentially recovering superfluid, as well as the external electromagnetic and magnetospheric torques.

The overall stability of the inter-glitch $\dot{\nu}$ can be observed in Fig. 7, where the modelled exponential and persisting glitch components are ‘corrected’ for, by subtracting $[\Delta\dot{\nu}_p - \exp(-(t - t_g)/\tau_d)\Delta\nu_d/\tau_d]$ from $\dot{\nu}(t)$ after each glitch. This results in the disappearance of the sawtooth curve and the long-term negative slope of $\dot{\nu}$ observed in Fig. C.1, leaving a predominantly linear, smooth evolution under a remarkably stable positive slope ($\ddot{\nu} > 0$). This high inter-glitch $\dot{\nu}$ and some additional wandering features (above random scatter) seen in Fig. 7 remain unaccounted for by the current glitch model, which has only one exponential term to describe internal torques and does not for example include the response of non-linearly coupled superfluid components (Haskell et al. 2020). We performed separate linear fits for the RXTE and NICER data and found $\dot{\nu} = (1.059 \pm 0.003) \cdot 10^{-20} \text{ Hz s}^{-2}$ for RXTE and $\dot{\nu} = (1.182 \pm 0.013) \cdot 10^{-20} \text{ Hz s}^{-2}$ for NICER. The corresponding braking indices are high ($n_{\text{ig}} =$

16.49 ± 0.05 and $n_{\text{ig}} = 18.4 \pm 0.2$, respectively) as in other similar glitching pulsars.

To better understand the variations in $\dot{\nu}$, we subtracted (additionally to the corrections for the glitch components) a linear model with a fixed, reference value of $\dot{\nu}_{\text{ref}} = 0.44 \times 10^{-20} \text{ Hz s}^{-2}$, which corresponds to one of the lowest inter-glitch measurements of $\dot{\nu}$ (observed after the first and largest glitch; $n_{\text{ig}} = 6.8$). These new residuals are shown in Fig. 8. As expected, the $\dot{\nu}$ residuals are flat in the first inter-glitch interval; however, they also remain relatively flat after the second glitch (large, followed by an exponential relaxation), indicating that $\dot{\nu}$ remains close to $\dot{\nu}_{\text{ref}}$. However, around the time of the third glitch, which is large but has no detected exponential, $\dot{\nu}$ increases and remains high until the next exponentially recovering glitch occurs (glitch 6), after which it decreases again and remains slightly above $\dot{\nu}_{\text{ref}}$ throughout several glitches. Similar patterns can be seen in the rest of the data: $\dot{\nu}$ falls close to $\dot{\nu}_{\text{ref}}$ following glitches with detected exponentials but overall spends long periods around a higher, rather constant, value. This is clearest in the first half of the middle panel, when many glitches occurred but none had detectable exponentials: $\dot{\nu}$ remains high and relatively constant for about 2000 days. Remarkably, some of the sharper increases in $\dot{\nu}$ in Fig. 8 appear to occur around small glitches instead of intermediate or large ones (and never after an exponentially recovering glitch). To explore this possible relation between the inter-glitch $\dot{\nu}$ and the size of the previous glitch, we plot the value of $\dot{\nu}$ after each glitch, i.e. $\dot{\nu} + \Delta\dot{\nu}_p$, against $\Delta\nu_p$ of the preceding glitch. The result is shown in Fig. 9, which shows a weak anti-correlation between $\dot{\nu} + \Delta\dot{\nu}_p$ and $\Delta\nu_p$ with the Spearman-rank correlation coefficient $\rho_s = -0.314$ and a p -value of $p = 0.012$.

5.2. $\Delta\dot{\nu}_p$ and spin-down evolution

As the frequent, large persisting changes of the spin-down rate $\Delta\dot{\nu}_p$ are the most probable cause of the observed negative long-term trend of the spin-down evolution (Fig. C.1), we explored their correlations with the change of $\dot{\nu}$ during the time intervals preceding and following the glitch. The results can be seen in Fig. 10. We calculated the change in the spin-down rate during the inter-glitch intervals as $\dot{\nu} \cdot t^-$ for the pre-glitch data span and $(\dot{\nu} + \Delta\dot{\nu}_p) \cdot t^+$ for the post-glitch data span. Here, t^- refers to the time since the previous glitch, and t^+ refers to the time to the next glitch. Following Liu et al. (2024), the smallest glitches ($\Delta\nu/\nu < 10^{-7}$) were excluded from the analysis, and the waiting times were recalculated between the remaining glitches. We find a significant correlation between the increase in $\dot{\nu}$ during an inter-glitch interval and its change $|\Delta\dot{\nu}_p|$ at the subsequent glitch, with $\rho_s = 0.79$ and p -value = $9.3 \cdot 10^{-13}$. This is in accordance with previous results for this particular pulsar (Middleditch et al. 2006; Antonopoulou et al. 2018). On the other hand, a correlation of $|\Delta\dot{\nu}_p|$ with the change of $\dot{\nu}$ in the time until the next glitch is not as tight, with $\rho_s = 0.29$ and p -value = 0.03, although such a correlation has been observed for several other pulsars (Lower et al. 2021; Liu et al. 2024). It was already noted in Liu et al. (2024) that both correlations are seen for some glitching pulsars, and that for PSR J0537–6910 $|\Delta\dot{\nu}|$ is more strongly correlated with the preceding (rather than following) increase in $\dot{\nu}$. Our results, which use the persisting changes (instead of the total change, which includes the exponentially relaxing components) further strengthen this conclusion. Because $\dot{\nu} \cdot t^-$ is a proxy of the amount of superfluid that recoupled since the previous glitch and is therefore available to decouple again at the next glitch causing a bigger $\Delta\dot{\nu}$, this

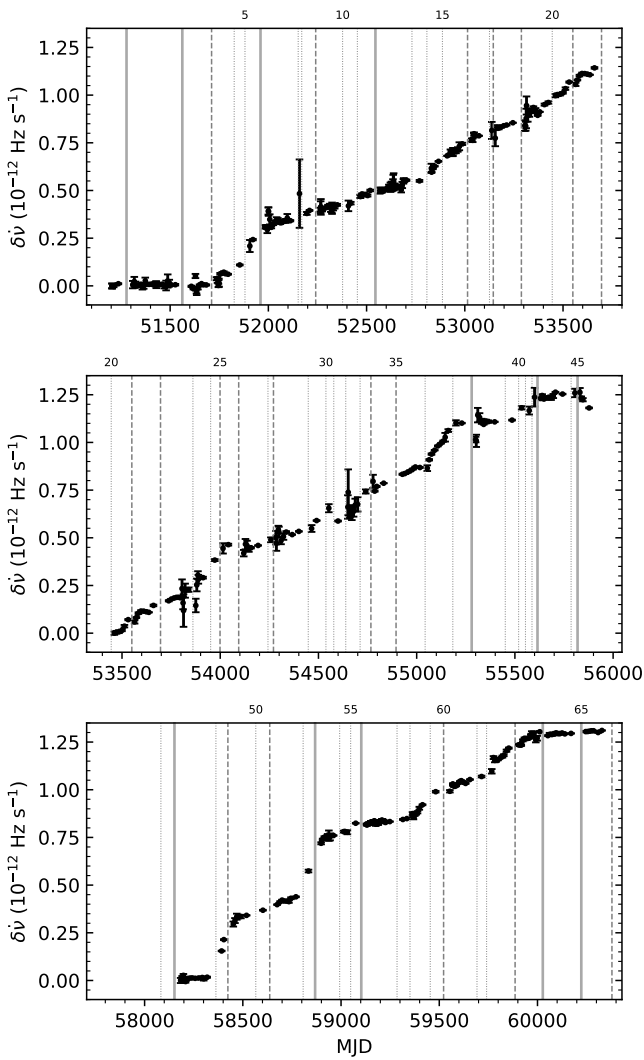


Fig. 8. Residuals of $\dot{\nu}$ evolution (as calculated for Fig. 7, with the modelled glitch persistent change and exponentially recovering components removed), relative to a linear model with slope $\dot{\nu} = 0.44 \times 10^{-20} \text{ Hz s}^{-2}$. The full dataset was divided into three sections of similar length, shown in the three panels (RXTE data in the two superior panels and NICER at the bottom). All plots have equal vertical and horizontal scales. The solid grey lines mark the epochs of the glitches with a detected exponential recovery; dashed lines mark the epochs of glitches with $\Delta\nu > 19 \mu\text{Hz}$; and thin dotted lines represent the epochs of smaller glitches. Glitch numbers are indicated at top of each panel.

correlation implies a ‘reservoir’ effect for the glitch $\dot{\nu}$ changes in PSR J0537–6910 (Antonopoulou et al. 2018; Liu et al. 2024).

5.3. Distributions of $\Delta\nu_p$, $\Delta\dot{\nu}_p$, and waiting times

The number of glitches studied in this work (66) might be small by statistical standards, but it is the largest glitch sample from a single pulsar ever investigated. We therefore explored the distributions of basic glitch parameters ($\Delta\nu_p$, $\Delta\dot{\nu}_p$, and inter-glitch waiting times) using our updated measurements. The unbinned data were modelled using the STATS package of the SCIPY library. For visualisation purposes, we binned the data using the minimum histogram bin width between the Sturges rule and the Freedman-Diaconis estimator. Histograms of the glitch param-

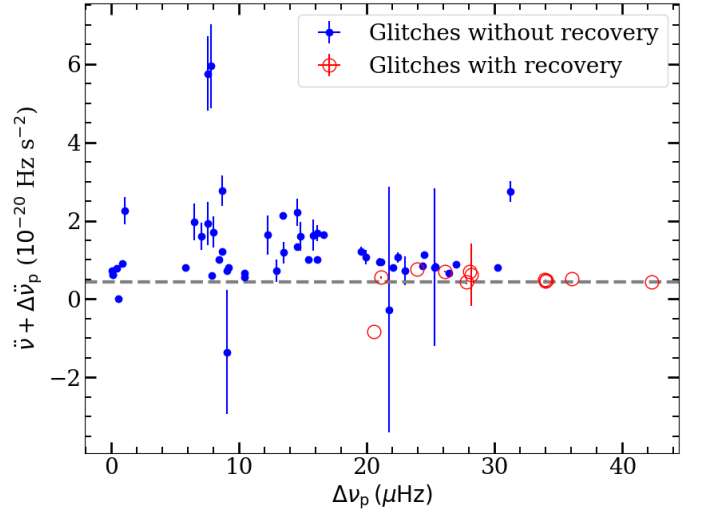


Fig. 9. $\dot{\nu} + \Delta\dot{\nu}_p$, as a function of $\Delta\nu_p$ (that is, inter-glitch $\dot{\nu}$ vs. the size of the previous glitch) for all 66 glitches. The horizontal dashed line indicates $\dot{\nu} = 0.44 \cdot 10^{-20} \text{ Hz s}^{-2}$.

Table 3. Description of the normal distributions of $\Delta\nu$, $\Delta\dot{\nu}$, and the waiting times.

	$\Delta\nu_p$	$\Delta\dot{\nu}_p$	Waiting time
KS	0.10	0.09	0.12
p -value	0.48	0.54	0.22
μ	16(1)	$-11.3(9)$	108(7)
σ	10.2(9)	7.3(7)	57(5)

Notes. KS is the result of the Kolmogorov-Smirnov test, μ is the centre of the distribution, and σ is its width. For $\Delta\dot{\nu}_p$, μ and σ are in units of ten to the minus fourteen hertz per second. For $\Delta\nu_p$, they are expressed in microhertz, and, for the waiting time, they are expressed in days. For μ and σ , the numbers in parenthesis correspond to their uncertainties in the last digit.

eters are shown in Fig. 11, together with the fitted models for Gaussian distributions, which are an acceptable fit to the data.

We find that $\Delta\nu_p$ and $\Delta\dot{\nu}_p$ are consistent with normal distributions centred at $16 \mu\text{Hz}$ and $-11.3 \times 10^{-14} \text{ Hz s}^{-1}$, respectively, with p -value of 0.48 and 0.54. With respect to the inter-glitch waiting times, the p -value of the normal distribution is 0.22, and it is centred at $\mu = 108$ days – in accordance to that in Fig. 5, which shows that a waiting time of ≈ 108 d corresponds to a $\Delta\nu_p \approx 16 \mu\text{Hz}$. The parameters of the fitted distributions are presented in Table 3.

6. Conclusion

PSR J0537–6910 is the most frequently glitching pulsar, exhibiting an average of three glitches per year. Analysis of some recent NICER observations delivered the detection of six new glitches with properties typical for this pulsar (Table 1). With these new detections, the total number of published glitches for this pulsar is now 66. The vast majority of the glitches in PSR J0537–6910 are large compared to the majority of glitches detected in other pulsars, with 70% of them having $\Delta\nu > 10 \mu\text{Hz}$. In other young pulsars, large glitches such as these are often followed by exponential-like recoveries. However, in the case of PSR J0537–6910, only one exponential recovery had been reported despite the detection of more than 50 large glitches.




Fig. 10. Comparison between $\Delta\dot{\nu}_p$ and the waiting time (τ^\pm) multiplied by the value of the second derivative of the frequency ($\dot{\nu}$ before the glitch, $\dot{\nu} + \Delta\dot{\nu}_p$ after the glitch). The red points represent the glitches for which recoveries were detected, and the blue points represent the glitches for which we did not detect a recovery. The blue square represents a glitch for which we measured a positive value of $\Delta\dot{\nu}_p$. Top panel: Time (τ^-) since the previous glitch. The vertical arrows indicate the lower limits for glitches 1 (red arrow) and 46 (blue arrow). Bottom panel: Waiting time (τ^+) until the following glitch. The vertical arrow indicates the lower limit for glitch 45. As in Fig. 5, we included glitch 66 in this plot.

We performed a systematic search for exponential recoveries following all major glitches of PSR J0537–6910 and defined detection criteria based on model comparison and Bayesian evidence. Using this method, 11 new exponential recoveries were identified and characterised.

The presence of detectable exponential recoveries seems to depend on glitch size. Every glitch with sizes $\Delta\nu$ above $31.3 \mu\text{Hz}$ exhibits an exponential recovery, but only some glitches of amplitudes $20 \mu\text{Hz} < \Delta\nu < 31.3 \mu\text{Hz}$ do, and no exponentials are detected for any glitch with sizes below $20 \mu\text{s}$. PSR J0537–6910 has a uniquely strong correlation between glitch size and the length of its following inter-glitch interval. This could introduce an observational bias as, for example, glitches with sizes below $10 \mu\text{s}$ are followed by only a ~ 50 d interval, which might be too short to securely detect the exponential recovery. It cannot, however, account for the lack of detections for glitches with $\Delta\nu \gtrsim 20 \mu\text{s}$, which we believe may be an intrinsic behaviour of the pulsar.

For the detected recoveries, the measured $\Delta\nu_d$ values are between $0.05 \mu\text{Hz}$ and $0.8 \mu\text{Hz}$, while τ_d ranges between 3 d and 37 d, which evidences the need for a high-cadence monitoring of this pulsar to thoroughly characterise its glitches. All the values of Q are below 4%, while Q' consistently lie between 46% and 100%. Including exponential recovery in the timing model




Fig. 11. Distribution of $\Delta\nu_p$ (top panel), $\Delta\dot{\nu}_p$ (middle panel), and the waiting times (bottom panel) of the glitches studied in this work. $\Delta\nu_p$ and $\Delta\dot{\nu}_p$ are fitted with a normal distribution, and the waiting times with a normal and a bimodal distribution.

affects the inferred parameters, especially $\Delta\nu_p$ and $\dot{\nu}$. We re-evaluated known correlations of this pulsar using the new solutions and find a very strong correlation between $\Delta\nu_p$ and the time until the following glitch, as well as a strong correlation between the change of $\dot{\nu}$ over a pre-glitch interval t^- , calculated as $(\dot{\nu} \cdot t^-)$, and $|\Delta\dot{\nu}_p|$ at the following glitch. We also find a mild anti-correlation between $\Delta\nu_p$ and its post-glitch value of $\dot{\nu}$. Importantly, it was demonstrated that the spin-down rate $\dot{\nu}$

evolves at a lower $\dot{\nu}$ following glitches with exponential recovery (once the exponential relaxation is over) compared to the generally stable evolution with higher $\dot{\nu}$ following glitches without detected exponential decay.

Our models do not include a red noise component. While timing noise can certainly influence some of the inferred recovery parameters or even mimic an exponential relaxation, the combination of theoretical expectations, consistent trends across glitches, and the distinct residual structures strengthens our confidence that the majority of the detected exponential components represent real post-glitch behaviour.

These findings can help us understand the mechanism behind glitches and the superfluid recoupling process that follows them, as well as shed light on the interplay of various internal and external torques for this pulsar. Continuous high-cadence observations of glitching pulsars are essential for studying neutron star dynamics, and PSR J0537–6910 is an excellent target for such monitoring due to its frequent, large glitches and unique characteristics. As we showed in this work, aspects of its rotational behaviour remain under-explored, and fast post-glitch transients require good time coverage to achieve further progress.

Acknowledgements. We thank Zaven Arzoumanian and Keith Gendreau for scheduling extra observations with NICER, which allowed the detection of some of the post-glitch recoveries reported here. CME acknowledges support from ANID/FONDECYT, grant 1211964. DA acknowledges support from an EPSRC/STFC fellowship (EP/T017325/1). WCGH acknowledges support through grant 80NSSC22K1305 from NASA. FG is a CONICET researcher. FG acknowledges support from PIBAA 1275 and PIP 0113 (CONICET). We used computing resources made available by ANID Chile SIA, grant SA77210112 (PI: Pinto). This work is supported by NASA through the NICER mission and the Astrophysics Explorers Program and uses data and software provided by the High Energy Astrophysics Science Archive Research Center (HEASARC), which is a service of the Astrophysics Science Division at NASA/GSFC.

References

- Abbott, R., Abbott, T. D., Abraham, S., et al. 2021, *ApJ*, 913, L27
- Anderson, P. W., & Itoh, N. 1975, *Nature*, 256, 25
- Andersson, N., Antonopoulou, D., Espinoza, C. M., Haskell, B., & Ho, W. C. G. 2018, *ApJ*, 864, 137
- Antonopoulou, D., Espinoza, C. M., Kuiper, L., & Andersson, N. 2018, *MNRAS*, 473, 1644
- Antonopoulou, D., Haskell, B., & Espinoza, C. M. 2022, *Rep. Prog. Phys.*, 85, 126901
- Ashton, G., Lasky, P. D., Graber, V., & Palfreyman, J. 2019, *Nat. Astron.*, 3, 1143
- Basu, A., Shaw, B., Antonopoulou, D., et al. 2022, *MNRAS*, 510, 4049
- Baym, G., Pethick, C., & Pines, D. 1969, *Nature*, 224, 673
- Crawford, F. 2024, *American Astronomical Society Meeting Abstracts*, 243, 132.04
- Crawford, F., McLaughlin, M., Johnston, S., Romani, R., & Sorrelgreen, E. 2005, *Adv. Space Res.*, 35, 1181
- Dodson, R. G., McCulloch, P. M., & Lewis, D. R. 2002, *ApJ*, 564, L85
- Espinoza, C. M., Lyne, A. G., Stappers, B. W., & Kramer, M. 2011, *MNRAS*, 414, 1679
- Espinoza, C. M., Lyne, A. G., & Stappers, B. W. 2017, *MNRAS*, 466, 147
- Ferdman, R. D., Archibald, R. F., Gourgouliatos, K. N., & Kaspi, V. M. 2018, *ApJ*, 852, 123
- Fuentes, J. R., Espinoza, C. M., Reisenegger, A., et al. 2017, *A&A*, 608, A131
- Haskell, B., & Melatos, A. 2015, *Int. J. Mod. Phys. D*, 24, 1530008
- Haskell, B., Antonopoulou, D., & Barenghi, C. 2020, *MNRAS*, 499, 161
- Ho, W. C. G., Espinoza, C. M., Arzoumanian, Z., et al. 2020, *MNRAS*, 498, 4605
- Ho, W. C. G., Kuiper, L., Espinoza, C. M., et al. 2022, *ApJ*, 939, 7
- Ho, W. C. G., Kuiper, L., Espinoza, C. M., et al. 2026, *MNRAS*, accepted
- Hobbs, G. B., Edwards, R. T., & Manchester, R. N. 2006, *MNRAS*, 369, 655
- Hobbs, G., Coles, W., Manchester, R. N., et al. 2012, *MNRAS*, 427, 2780
- Kuiper, L., & Hermsen, W. 2009, *A&A*, 501, 1031
- Kuiper, L., & Hermsen, W. 2015, *MNRAS*, 449, 3827
- Liu, Y., Keith, M. J., Antonopoulou, D., et al. 2024, *MNRAS*, 532, 859
- Lower, M. E., Johnston, S., Dunn, L., et al. 2021, *MNRAS*, 508, 3251
- Luo, J., Ransom, S., Demorest, P., et al. 2021, *ApJ*, 911, 45
- Lyne, A. G., Jordan, C. A., Graham-Smith, F., et al. 2015, *MNRAS*, 446, 857
- Manchester, R. N. 2018, in *Pulsar Astrophysics the Next Fifty Years*, eds. P. Weltevrede, B. B. P. Perera, L. L. Preston, & S. Sanidas, 337, 197
- Marshall, F. E., Gotthelf, E. V., Zhang, W., Middleditch, J., & Wang, Q. D. 1998, *ApJ*, 499, L179
- Middleditch, J., Marshall, F. E., Wang, Q. D., Gotthelf, E. V., & Zhang, W. 2006, *ApJ*, 652, 1531
- Nasa High Energy Astrophysics Science Archive Research Center (Heasarc) 2014, *Astrophysics Source Code Library* [record ascl:1408.004]
- Palfreyman, J., Dickey, J. M., Hotan, A., Ellingsen, S., & van Straten, W. 2018, *Nature*, 556, 219
- Skilling, J. 2004, in *Bayesian Inference and Maximum Entropy Methods in Science and Engineering: 24th International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, eds. R. Fischer, R. Preuss, & U. V. Toussaint (AIP), *AIP Conf. Ser.*, 735, 395
- Wang, Q. D., & Gotthelf, E. V. 1998, *ApJ*, 494, 623
- Weltevrede, P., Johnston, S., & Espinoza, C. M. 2011, *MNRAS*, 411, 1917
- Yu, M., Manchester, R. N., Hobbs, G., et al. 2013, *MNRAS*, 429, 688
- Zhou, S., Gügercinoğlu, E., Yuan, J., Ge, M., & Yu, C. 2022, *Universe*, 8, 641
- Zubieta, E., Missel, R., Sosa Fiscella, V., et al. 2023, *MNRAS*, 521, 4504
- Zubieta, E., García, F., del Palacio, S., et al. 2024, *A&A*, 689, A191
- Zubieta, E., García, F., del Palacio, S., et al. 2025a, *A&A*, 694, A124
- Zubieta, E., Missel, R., Araujo Furlan, S. B., et al. 2025b, *A&A*, 698, A72

Appendix A: Data intervals used to measure new glitches and exponential recoveries

Table A.1 gives details on the fits of six new glitches (Table 2). Table A.2 shows the Bayesian evidence of Models 1 and 2 for each glitch where model 2 was favoured. Table A.3 presents details of the fits of glitch exponential recoveries to 12 glitches (Table 2).

Table A.1. Data used to model six new glitches (Table 1).

Glitch	START (MJD)	FINISH (MJD)	Pre-glitch TOAs	Post-glitch TOAs	Residual RMS (μs)	χ^2_{red}
61	59529.5	59735.4	27	12	138.7	4.4
62	59695.9	59882.6	12	22	142.5	4.8
63	59745.7	60026.9	22	26	148.5	6.0
64	59888.5	60221.3	26	21	275.1	20.4
65	60071.8	60357.4	16	15	124.4	4.1
66	60287.7	60601.4	10	7	104.5	4.1

Notes. Glitches 64 and 65 exhibit clear post-glitch recovery signatures. We removed the TOAs affected by these recoveries from the pre-glitch data span of glitches 65 and 66 to obtain an accurate pre-glitch solution before characterising the respective glitches.

Table A.2. Bayesian evidence for the models corresponding to the 12 positive detections.

Glitch	t_g (MJD)	$\Delta\dot{v}_p$ in Model 1	$\Delta\dot{v}_p$ in Model 2	Evidence Model 1	Evidence Model 2	$\ln B_{21}$	Evidence criteria
1	51278(16)	No	No	318.96	355.63	16	Decisive
2	51562(15)	Yes	No	285.94	292.45	2.82	Moderate
6	51960(5)	Yes	Yes	248.69	265.81	7.48	Strong
12	52545(6)	Yes	No	340.81	352.68	5.16	Strong
38	55280(4)	Yes	No	161.71	181.80	8.72	Strong
43	55615(4)	No	No	141.72	164.92	10.07	Decisive
45	55819(2)	No	No	118.74	127.78	3.92	Moderate
47	58152(11)	No	No	174.70	186.76	5.23	Strong
53	58868(5)	Yes	Yes	124.58	166.76	18.31	Decisive
56	59103(6)	Yes	No	240.66	328.32	38.07	Decisive
64	60027(1)	Yes	Yes	29.60	348.15	138.34	Decisive
65	60223(3)	Yes	No	222.24	246.29	10.44	Decisive

Table A.3. Description of the data used to measure exponential glitch recoveries (Table 2).

Glitch	START (MJD)	FINISH (MJD)	Pre-glitch TOAs	Post-glitch TOAs	Residual RMS (μs)	χ^2_{red}	Time to next glitch (days)
1	51197.1	51546.7	11	32	104.6	0.9	284
2	51423.0	51705.2	18	18	123.9	1.32	149
6	51886.9	52144.1	10	26	147.4	1.42	192
12	52460.0	52717.4	14	28	124.9	1.23	186
38	55185.4	55444.8	9	15	200.0	2.27	171
43	55589.3	55786.1	6	15	157.2	1.6	171
45	55794.7	55926.9	5	11	122.9	1.2	≥ 107
47	58108.1	58348.8	3	21	89.6	2.9	211
53	58810.0	58990.2	6	17	93.5	3.8	125
56	59051.4	59282.4	10	31	82.2	2.0	182
64	59888.5	60221.3	26	21	120.3	4.1	195
65	60071.8	60357.4	16	15	77.4	1.6	156

Appendix B: Examples of glitches without a detected recovery

We present two examples of large glitches for which we did not detect exponential recovery. In these cases, the algorithm described in Sect. 3.2 either did not identify a minimum of χ_{red}^2 within 1 day to 4 times the length of the data span, or, if it did, the Bayesian evidence for Model 2 was not significant compared to Model 1.




Fig. B.1. Examples of glitches 28 (left) and 51 (right) without a detected exponential recovery. From top to bottom: Phase residuals relative to Model 1, with the residuals (blue curve) of the most likely exponential recovery model (Model 2 relative to Model 1), corresponding here to zero, superimposed; phase residuals relative to Model 2; frequency residuals relative to Eq. 1 fitted to TOAs up to the glitch epoch, with the post-glitch data all lowered by a certain amount (the mean post-glitch frequency residual) for better visualisation; and $\dot{\nu}$ evolution with Model 2, when used as the favoured model, shown as the blue curve.

Appendix C: Evolution of $\dot{\nu}$

We present the evolution of $\dot{\nu}$ in Fig. C.1. The evolution shows a global negative slope, which is caused by repeated large glitches. The slope becomes positive when the effects of the glitches are removed.




Fig. C.1. Evolution of the spin-down rate. For visual purposes, the quantity $\delta\dot{\nu} = \dot{\nu} + 199 \times 10^{-12}$ was plotted. The full dataset was divided as in Fig. 8. The dashed vertical lines mark the epochs of the glitches.