Pulsar glitches: observations and physical interpretation

Danai Antonopoulou[1], Brynmor Haskell[2] and Cristóbal M. Espinoza[3,4]

[1]Jodrell Bank Centre for Astrophysics, School of Physics and Astronomy, University of Manchester, Manchester UK, M13 9PL
[2]Nicolaus Copernicus Astronomical Center of the Polish Academy of Sciences, Bartycka 18, 00-716 Warsaw, Poland
[3]Departamento de Física, Universidad de Santiago de Chile (USACH), Av. Victor Jara 3493, Estación Central, Chile
[4]Center for Interdisciplinary Research in Astrophysics and Space Sciences (CIRAS), Universidad de Santiago de Chile, Chile

E-mail: danai.antonopoulou@manchester.ac.uk

September 2022

Abstract.

The interpretation of pulsar rotational glitches, the sudden increase in spin frequency of neutron stars, is a half-century-old challenge. The common view is that glitches are driven by the dynamics of the stellar interior, and connect in particular to the interactions between a large-scale neutron superfluid and the other stellar components. This thesis is corroborated by observational data of glitches and the post-glitch response seen in pulsars' rotation, which often involves very long timescales, from months to years. As such, glitch observables combined with consistent models incorporating the rich physics of neutron stars – from the lattice structure of their crust to the equation of state for matter beyond nuclear densities – can be very powerful at placing limits on, and reduce uncertainties of, the internal properties. This review summarises glitch observations, current data, and recent analyses, and connects them to the underlying mechanisms and microphysical parameters in the context of the most advanced theoretical glitch models to date.

Keywords: neutron star, pulsar, glitch

1. Introduction

1.1. Historical overview

Straight after the neutron particle was discovered in 1932, theoretical speculations about neutron stars began [1, 2, 3]. The confirmation of their existence did not, however, come until 1967 when the first pulsar was discovered, followed promptly by the detection of three more such sources [4]. Thereafter the population of observed pulsars quickly increased; today we know of over 3000 neutron stars detected at various wavelengths, from the radio band up to high energies such as TeV γ-rays [5]. Pulsar emission appears modulated by the star's rotational frequency $\nu = \Omega/(2\pi)$ (where Ω is the angular velocity), which is remarkably stable and allows for microsecond precision in the measurement of the pulsed signal [6].

When, in November 1968, the source PSR J0534+2200 was discovered in the Crab Nebula, its period was pinned down to be $P = 1/\nu = 33.09$ ms. Such a short spin period is difficult to explain if the source were a pulsating white dwarf, which gave confidence that pulsars are rotating, magnetised, neutron stars; as proposed by Gold [7]. A few weeks later the period of the Crab pulsar was found to be slowly but steadily increasing [8], providing enough energy to power its surrounding nebula. This secular spin-down was observed in other pulsars [9] and fitted in with the theoretical interpretation of these sources.

But then, between February 24 and March 3, 1969, something unexpected happened. The source PSR B0833−45, associated with the Vela supernova remnant, sped up, its period decreasing by 134 ns, and then resumed its spindown at a higher rate than before [10, 11]. Pulse intensity did not seem to vary before or after the spin-up [10]. It was immediately appreciated that this event could be connected to the internal physics of neutron stars and driven by the relatively high spindown rate of the Vela pulsar, meaning the Crab pulsar would also be a good candidate to display such features [11]. Indeed, in September of the same year, the Crab pulsar also spun-up, although by a much smaller amount [12].

The intriguing phenomenon, nowadays called a pulsar *glitch*, caught the attention of many researchers who soon investigated various hypotheses for its origin. A natural explanation that was originally considered as the cause of either the spin-up or the change in spindown rate was a change ΔI in stellar moment of inertia I. Under the standard model for pulsar spindown, the relative change $|\Delta I|/I$ would lead to fractional changes in frequency $|\Delta\nu|/\nu$ and its rate of change $|\Delta\dot{\nu}/\dot{\nu}|$ of similar order of magnitude – this is not observed, therefore a decrease of moment of inertia ΔI cannot explain both. Several mechanisms were proposed. For example, mass ejection could lead to the spin-up [10]; alternatively the spin-up itself could be due to some small contraction [11] together with an increase in the braking torque [10]. B. Durney explored the possibility of mass accretion that then leads to contraction (and, interestingly, this is likely the first model that suggests neutron star masses could be inferred from pulsar glitches) [13]. Other early ideas included collision with a smaller body [14], expulsion of plasmoids from the magnetosphere [15], differential rotation of a fluid star and transfer (by mixing) of angular momentum from the core to the outer layers, due to instability-driven changes in composition gradients [16].

A promising idea was put forward by M. Ruderman in 1969. He realised that the crust of a neutron star will become solid soon after its birth [17] and so the star will be 'frozen' in an oblate shape according to its, then higher, rotation rate; in this state the crust is relaxed (unstressed). As the star spins down, the deviation between the crust's relaxed state and the equilibrium one corresponding to its present-time rotation rate will grow, resulting in an increasing stress to the crust lattice. Eventually, a starquake will take place that relaxes the stress and decreases the moment of inertia, giving rise to an observed spin-up of the pulsar [18]. The early calculations suggested this could lead to a glitch size close to that observed in Vela pulsar. However, as Ruderman recognised, the fact that a glitch was already seen within the first year of observations could be problematic, as such starquakes should be rare events [18]. Indeed, subsequent observations of more glitches of the Vela pulsar revealed that a standalone starquake model cannot explain their frequency and size.

Following the starquake conception, Baym, Pethick, Pines and Ruderman presented a very insightful model that offered an explanation for the glitch change in spindown rate, which leads to a recovery of the glitch spin-up on long timescales. The idea built on the realisation by Migdal [19], and Ginzburg and Kirzhnits [20], that the neutrons inside the star will be in a superfluid state, having essentially zero viscosity. The star can thus be viewed as consisting of two components: one comprising the crust and charged particles which can be assumed uniformly rotating, the other being the neutron superfluid that can have a different rotational velocity. The two interact via a weak mutual friction force that couples them together on long timescales [21]. At a glitch, whilst the charged component quickly follows the spin-up of the crust (be it due to a starquake or other mechanism), the neutron superfluid responds slowly, on the characteristic timescale of the mutual friction, and temporarily disengages from the rest of the star. This gives rise to

the glitch change in spindown rate, since the external torque (assumed constant) now acts on a reduced effective moment of inertia. A slow relaxation towards the pre-glitch rotational state is expected as the superfluid couples back, and is indeed commonly detected. The observed relaxation timescales range from days to months, which is hard to reconcile with normal viscous processes, but will be the natural result of the presence of a weakly coupled superfluid. As the authors noted, "the event [the Vela glitch] can be taken as evidence that the interior of the pulsar is a superfluid" [21].

The above mechanism remains the prevailing model for the observable relaxation following pulsar glitches. In addition to the post-glitch response, the superfluid can also be invoked as the origin of the spin-up itself. Central to this picture is the presence and dynamics of microscopic neutron vortices (detailed in Section 2). Such vortices can form in otherwise irrotational superfluids and each of them carries a discrete amount of circulation. A superfluid can then sustain a rotation rate that scales with the density of vortices permeating it. In a decelerating system such as the spinning-down neutron star, vortices must move out of the superfluid for its angular velocity to reduce. However, this might not be possible in certain regions where continuous vortex flow is impeded by their interaction with normal matter. For example, in the inner crust of neutron stars the co-existence of the neutron superfluid with the solid lattice provides a way to "pin" quantised vortices in place [22]. As the crust spins down, the superfluid – unable to expel pinned vortices – can maintain a faster rotation rate which will decrease in a non-continuous fashion via vortex creep, in analogy to what is seen in superconductors [23, 24]. Glitches will arise when a large group of vortices collectively unpins and migrates outwards, rapidly transferring stored angular momentum from the superfluid to the crust which then spins up [25]. A yet open question within this picture is what drives the great number of vortices, $N_v \sim 10^{13}$, required to produce a large Vela glitch, to unpin simultaneously and move over many pinning sites, to annihilate at the edge of the superfluid. Several possibilities have been suggested, including crustquake-triggered unpinning, fluid instabilities and, more recently, vortex avalanches which will allow a small scale perturbation to grow to a large-scale unpinning event [26]. As will be covered in the following, this process of vortex avalanches appears sufficient to explain both the size and frequency of observed glitches, but could also operate alongside other glitch triggers.

1.2. Present-day overview

As more pulsars are being discovered and observed frequently enough to follow their rotation accurately over long timescales, the number of detected glitches has grown rapidly (figure 1). By early 2022, the number of reported glitches approached 700, identified in over 200 different pulsars. Although a simple estimate implies only about $\sim 10\%$ of neutron stars display glitches, it needs to be stressed that the real fraction of "glitching" pulsars could be significantly higher as glitches are generally infrequent events and only a small number of sources have been monitored closely over the years. In fact, despite their perceived rarity, glitches have been observed in many different classes of neutron stars such as in slowly rotating magnetars (young neutron stars with characteristic magnetic fields $B \gtrsim 10^{13} - 10^{15}$G, whose activity is powered by magnetic, rather than rotational, energy), rapidly-rotating millisecond pulsars, accreting sources in binary systems, slowly-rotating old pulsars, and are a very common feature among young (age $< 10^5$ yr), rotationally-powered, pulsars. Ongoing and forthcoming pulsar timing campaigns, such as monitoring at the Jodrell Bank Observatory, the CSIRO Parkes Observatory, and the CHIME, MeerKAT, and UTMOST projects, which aim to closely follow the rotation of many hundreds of pulsars, will improve the estimates of glitch prevalence in the pulsar population.

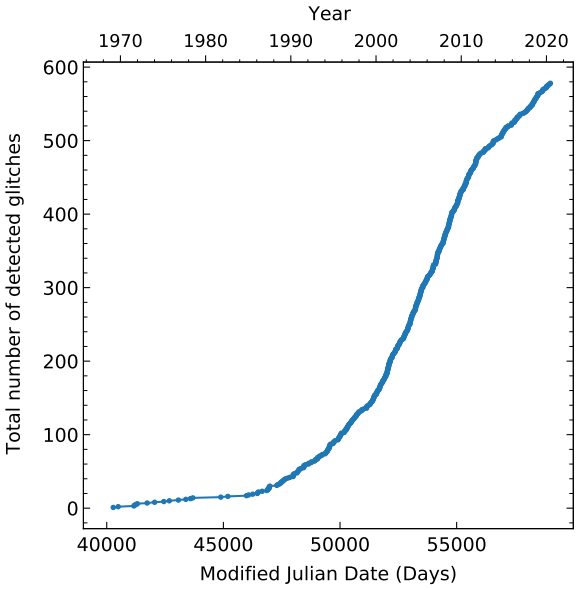

Figure 1. The total number of reported glitches as a function of their occurrence epoch, from 1969 until early 2021. The effect of targeted glitch searches such as the ones carried out using the Jodrell Bank Observatory [27] and Parkes Observatory [28] pulsar timing datasets is evident between Modified Julian Dates (MJD) $\sim 50000 - 56000$ and highlights the importance of such comprehensive analyses.

The typical properties of a glitch are first and foremost the size of the observed spin-up, $\Delta\nu$, and the change $\Delta\dot{\nu}$ in spindown rate that accompanies

it in most cases. The majority of inferred glitch sizes fall between 10^{-4} and 50 μHz. When a change $\Delta\dot{\nu}$ is observed, it usually has a relative size $\Delta\dot{\nu}/\dot{\nu}$ that is orders of magnitude greater than the relative change in spin frequency $\Delta\nu/\nu$. For a large Vela-like glitch, typical values would be $\Delta\nu/\nu \sim 10^{-6}$ whilst $\Delta\dot{\nu}/\dot{\nu} \sim 10^{-3}$. For most pulsars these changes appear instantaneous within the accuracy of the data, although there are a few instances for which the glitch rise has been partially resolved in time (Section 5.3).

Post-glitch response is highly diverse between pulsars as well as for different glitches of the same object [29, 30, 31]. In some cases there appears to be no evolution towards the pre-glitch state, with the glitch changes persisting over time. Most often, partial recovery is observed which can be characterised either by one or multiple exponential terms, and/or a change in the second frequency derivative $\ddot{\nu}$ (Section 5.4). The recovery can be on very long timescales (months to several years) and might be interrupted by future glitches. An assortment of glitches exemplifying the variety of the post-glitch behaviour can be found in figure 6.

Whilst a lack of detectable emission changes seems to characterise the majority of glitches, and is considered affirmation of their internal origin, the situation is different in pulsars with high inferred magnetic fields and most notably in magnetars. Glitches in such neutron stars sometimes coincide with radiative outbursts or more subtle magnetospheric changes. The latest observations however provide indications that the pulsar magnetosphere might also be involved even in pulsars with moderate magnetic field strengths (Section 6).

This review focuses on the observational attributes of glitches and highlights examples of the information they carry about the underlying mechanism and the physical properties of neutron stars. It does not aim to be exhaustive in the discussion of theoretical models (a recent review of glitch models can be found in [26]). Instead, Section 2 contains some general ideas common to most models and introduces the key parameters that enter in the description of the internal superfluid dynamics, which are frequently used to tie observable properties to the microphysics. Theoretical details of specific aspects are discussed along their observable consequences in subsequent sections.

Pulsar observations, glitch detection and characterisation are presented in Section 3. Section 4 is about the properties of the most basic glitch parameters, such as their size and temporal distribution, and their possible interpretation. The different stages of the glitch phenomenon, from the fast rise to the long post-glitch recoveries, are considered in Section 5 together with the information they carry about the underlying microphysical mechanisms and internal stellar structure. Possible connections between glitches and radiative changes in ordinary pulsars as well as magnetars are summarised in Section 6. Finally, the impact of glitches on the long-term evolution of neutron stars is discussed in Section 7, followed by a brief summary and outlook in Section 8.

2. The multicomponent neutron star model

The expected structure of the interior of mature neutron stars is illustrated in figure 2. Non-accreting pulsars will be considered, under the assumption that matter is in its ground state. Underneath a very thin envelope of light elements such as hydrogen and helium, we can identify the outer crust region in the density range $10^6\,\mathrm{g\,cm^{-3}} - 10^{11}\,\mathrm{g\,cm^{-3}}$. It is composed of relativistic electrons and fully ionised, heavy, nuclei. Electron capture by the nuclei becomes possible for densities $> 10^7\,\mathrm{g\,cm^{-3}}$, resulting in an increasing neutron fraction inside the nuclei as density grows. Neutrons drip out of the nuclei beyond a density $\sim 4 \times 10^{11}\,\mathrm{g\,cm^{-3}}$. This neutron drip density denotes the transition to the inner crust region, in which a fluid of unbound neutrons coexists with the nuclei and ultrarelativistic electrons. For temperatures lower than $T \sim 10^9 - 10^{10}\,\mathrm{K}$, the nuclei in these two outer layers will form a solid. Microscopically, the equilibrium state of this solid is expected to be a body centred cubic lattice, but history-dependent deformations can appear and the exact properties of the lattice are an open research question (see [32] for details on the outer and inner crust layers).

As densities approach the saturation density of symmetrical nuclear matter $\rho_0 = 2.8 \times 10^{14}\,\mathrm{g\,cm^{-3}}$, nuclei dissolve. The outer core of the star begins, at a density around $0.5\rho_0$. This stellar part consists of a fluid mixture of neutrons, and a much smaller fraction of protons and electrons (as well as muons in higher densities). At even higher densities ($\gtrsim 2\rho_0$), extrapolations from the physics of the states of matter familiar to us become more uncertain. This inner core section, of relevance to the most massive neutron stars whose central densities can reach $5 - 10\rho_0$, possibly contains hyperons and other hadrons (e.g. mesons), or even deconfined quarks.

In the inner crust and core areas, and at low enough temperatures, the long-range attractive component of the nucleon-nucleon interaction will result in the formation of bound neutron pairs. At densities pertinent to the inner crust, neutrons will form pairs of zero total spin (singlet state, $^1S_0$). In the outer core, close to the crust-core boundary, this state is expected to give place to pairing in the triplet state (total pair spin of 1). A characteristic of the

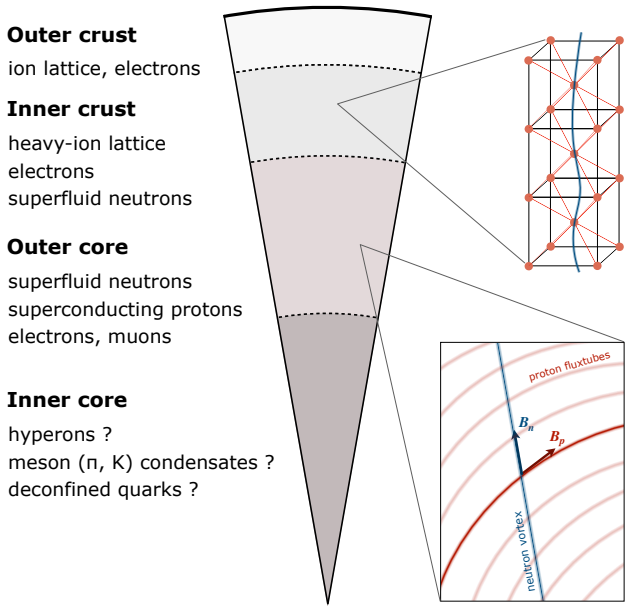

Figure 2. Neutron star composition and possible sites of vortex pinning. The outer crust and inner crust are each expected to be a few hundred meters thick. Superfluid neutron vortices in the inner crust interact with, and can "pin" to, the lattice of nuclei. This is shown in schematic form at the top right corner of the figure, for a vortex (in blue colour) immersed in the lattice. The extent of the outer core region is typically $\sim$ 4 km, but may be the entire core for some equations of state and neutron star masses. In this region neutron vortices will interact with magnetic fluxtubes if the protons form a type II superconductor (bottom right insert). The inner core is estimated to be $\sim$ 5 km in radius and reaches densities of the order $10^{15}\,\mathrm{g\,cm^{-3}}$. The composition and state of matter in this central part is uncertain.

paired system is the existence of an energy gap: a minimum energy is required to create a fermionic excitation from the ground state, since a pair has to be broken. Because of this non-zero pairing gap, below a critical temperature the neutrons are superfluid: they can flow without dissipation as scatterings which lead to ordinary viscosity are being suppressed. This transition to the superfluid state will start very early in the life of a neutron star, as the critical temperature can be as high as $\sim 10^{10}\,\mathrm{K}$ in some densities. For pulsar ages $\sim 10\,\mathrm{kyr}$ the bulk of the inner crust and core neutrons will be superfluid. Protons in the stellar core will also bind in pairs and form a superconductor, characterised by non-dissipative current flows.

The neutron superfluid lies at the heart of the most successful efforts to date at explaining glitches, which assume a star comprising of three basic components: the "normal" matter, which includes the solid crust and any elements coupled to it on short dynamical timescales (with respect to typical glitch timescales), a superfluid region that acts as an angular momentum reservoir and drives the spin-up, and the superfluid neutrons that respond to the glitch (the last two components do not necessarily correspond to two distinct physical regions).

The normal component typically contains the crustal ions and electrons, charged particles of the core (e.g. electrons, muons, protons – even if the latter are superconducting), and any non-superfluid neutrons. On the dynamical timescales of interest for glitches the normal component rotates as a rigid body, with its constituents strongly coupled (mainly electromagnetically) and its rotation rate Ω_c can be identified from the observed spin frequency ν of the crust as $\Omega_\mathrm{c} = 2\pi\nu$. On the other hand, the superfluid can sustain differential rotation dependent on its local coupling to the other components, which – as will be detailed hereinafter – can be very weak in some layers. The total angular momentum of the neutron star can thus be written as

$$L = I_\mathrm{c}\Omega_\mathrm{c} + \int \Omega_\mathrm{s}(\boldsymbol{r})\mathrm{d}I_\mathrm{s}\ , \tag{1}$$

where I_c and $\mathrm{d}I_\mathrm{s}(\boldsymbol{r})$ are the moments of inertia of the normal component and of a superfluid element at location $\boldsymbol{r}$, respectively, and $\Omega_\mathrm{s}(\boldsymbol{r})$ the superfluid angular velocity. The total angular momentum loss rate $\dot{L}$ is dictated by the external torque N_sd:

$$\dot{L} = \dot{L}_\mathrm{c} + \dot{L}_\mathrm{s} = N_\mathrm{sd}\ . \tag{2}$$

In essence, glitch models assume an exchange of angular momentum from part of the hidden superfluid to the observed normal component (that is, from the second to the first term in the right-hand side of (1)), whilst the total L is conserved because the effect of N_sd is negligible on timescales of seconds or hours as those involved in the glitch spin-up. The induced change in relative angular velocities causes a drop in $|\dot{L}_\mathrm{s}|$, which then gradually grows back towards its pre-glitch levels. This affects the evolution of $|\dot{L}_\mathrm{c}| = I_\mathrm{c}|\dot{\Omega}_\mathrm{c}|$ via equation (2) and is observed as the glitch recovery. This behaviour is explained by the way the quantity Ω_s, and its time evolution that defines $\dot{L}_\mathrm{s}$, connect to the distribution and moving of neutron vortices which is the focus of the upcoming section.

2.1. Superfluid vorticity and rotational dynamics

Rotation of the superfluid is realised by the appearance of a large number of vortices which carry a quantised amount of circulation whilst the bulk of the superfluid remains irrotational, as depicted in figure 3. The vortex core – of radius comparable to the superfluid coherence length ($\sim 10^2$ fm) – comprises normal matter. Outside the vortex core, the superfluid circulates around each vortex line with a velocity that falls as $1/r$ with distance r from the vortex core axis. In neutron stars, where superfluidity arises

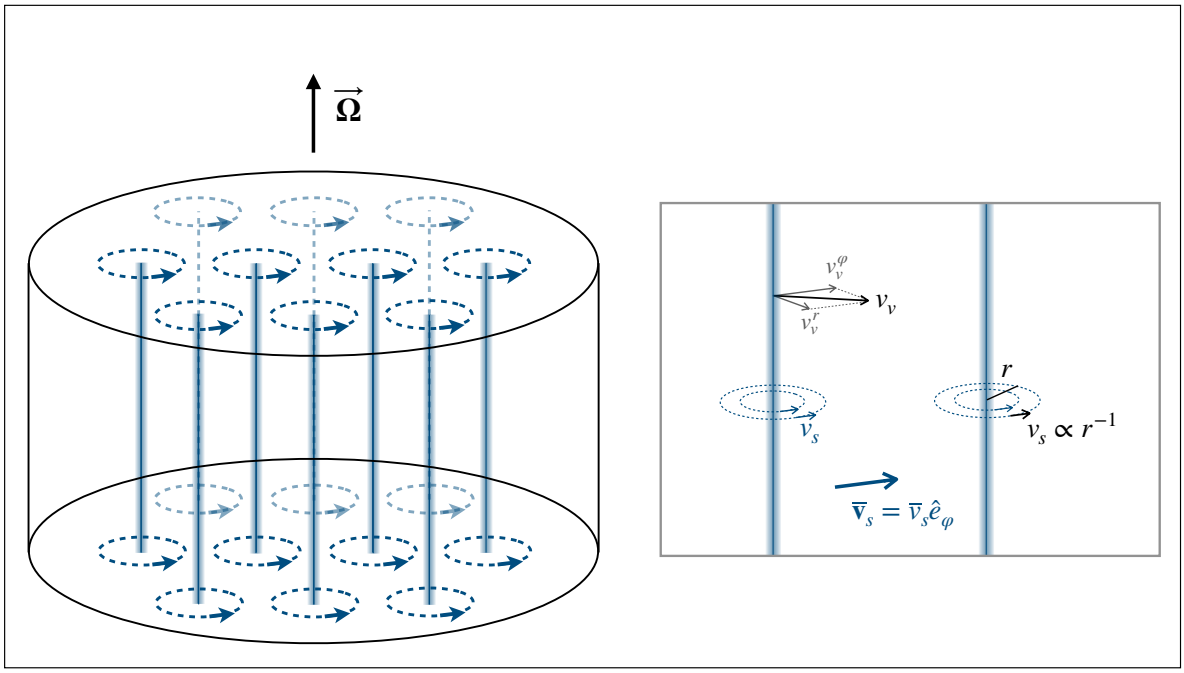

Figure 3. Left: Vortices will form and support the rotation of a superfluid when the angular velocity Ω of its container exceeds a critical value. Right: zooming in on a pair of vortices. Superfluid velocity v_s around a vortex line scales inversely proportional to the distance from the vortex axis. The presence of many vortices (assumed straight in this illustration) gives rise to an average $\overline{v}_s$ (which is in the φ direction in this case where $\mathbf{\Omega} = \Omega\hat{e}_z$). Under deceleration of the container, the vortex velocity v_v (otherwise following the superflow) acquires a non-zero radial component v_v^r. The induced outwards vortex motion decreases the total angular momentum of the superfluid, as described in the text.

because of neutron pairs, the quantum of circulation is $\kappa = h/(2m_n) \approx 1.978 \times 10^{-3}$ cm^2 s^{-1} with h Planck's constant and m_n the bare neutron mass. For simplicity, we will hereafter always assume singly-quantised vortices, which are generally expected to be the lowest energy state in the crust and probably the entire core of neutron stars. Multiply quantized vortices are possible and could lead to observable differences, but are potentially unstable.

The local rotation rate Ω_s of the superfluid is determined by averaging over the microscopic scale circulation induced by individual vortices. Ω_s therefore depends on the density of vortices n_v; were the superfluid to follow rigid rotation, the areal density of vortices would come from the Onsager-Feynman relationship (by equating the fluid circulation to that carried by N_v total vortices):

$$\kappa n_v = 2\Omega_\mathrm{s} \, . \tag{3}$$

Indicatively, taking $\Omega_\mathrm{s} = \Omega_\mathrm{c}$ the total number of vortices should be $N_v \sim 10^{17} - 10^{18}$ in a fast pulsar such as the Crab.

Any changes in the superfluid rotation rate must come from adjustments to its vortex density. Because the normal component of the star is slowing down, the superfluid momentarily finds itself with a higher rotation rate than the normal component. A lag in rotational velocities develops, $\omega(\boldsymbol{r}) = \Omega_\mathrm{s}(\boldsymbol{r}) - \Omega_\mathrm{c}$, between the normal component and the superfluid – and consequently, between the normal component and vortices, which would normally follow the superfluid flow. Unlike the bulk of the superfluid, however, vortices can interact with the normal component (details on the nature of these interactions can be found in Section 5.1), and so the lag between the two gives rise to forces that act on the vortices. Due to the effect of these forces a velocity difference develops between vortices and the superfluid flow, which in turn results in the appearance of a 'lift' (Magnus) force which makes vortices move outwards. This dissipative vortex motion mediates the exchange of angular momentum between the superfluid and the normal component. It is in this way that a non-zero $\dot{\Omega}_\mathrm{s}$ is induced and the superfluid can follow, at least partially, the spin-down of the normal component.

In regions where excess vorticity is removed in an unrestricted manner, the normal and superfluid components essentially spin down together ($\dot{\Omega}_\mathrm{s} = \dot{\Omega}_\mathrm{c}$), maintaining an equilibrium lag ω_eq. The superfluid responds (in the sense that ω_eq is restored) to changes in the rotation of the normal component Ω_c on short timescales if the two are strongly coupled, or longer

timescales if coupling is weaker.

There are regions, though, where vortex motion is counteracted by *pinning* forces. Pinning arises because there are energetically favourable positions for a neutron vortex to be in, as will be detailed in Section 2.3. For example, in the inner crust it might be preferable for the vortex core to coincide with (or exclude) the proton clusters (nuclei) of the lattice. In the neutron star core, pinning can occur as a result of the interaction between a magnetised neutron vortex and the superconducting proton vortices (see figure 2). Whichever the exact pinning mechanism, if strong it poses an impediment to the vortex flow which limits $|\dot{\Omega}|_{\rm s} < |\dot{\Omega}|_{\rm c}$. In these regions then, as $\Omega_{\rm c}$ decreases due to the spin-down, the lag $\omega = \Omega_{\rm s} - \Omega_{\rm c}$ grows over time and so does the relative velocity between the superfluid and pinned vortices (forced to co-rotate with their pinning sites). This results in an increasing Magnus force exerted on them:

$$\boldsymbol{f}_{\rm M} = \rho_{\rm s}\boldsymbol{\kappa} \times (\boldsymbol{v}_{\rm s} - \boldsymbol{v}_L) \tag{4}$$

per unit vortex length, where $\boldsymbol{v}_{\rm s}$ and $\boldsymbol{v}_L$ are the velocities of superfluid and vortex line respectively, $\rho_{\rm s}$ the superfluid mass density and $\boldsymbol{\kappa}$ has modulus κ and points along the vortex axis. When the lag reaches a threshold $\omega_{\rm cr}$, calculated by balancing the Magnus and pinning forces per unit vortex length, vortices unpin. Therefore $\omega_{\rm cr}$ places a maximum on the excess angular momentum that can be stored in a given superfluid region. Fast unpinning on a large scale, i.e. a vortex avalanche, will rapidly transfer some of the excess angular momentum to the crust, causing a glitch.

Whilst $\omega < \omega_{\rm cr}$, vortex lines in pinning regions can still move between pinning potentials because of thermal fluctuations (when vortex thermal energy exceeds the pinning barrier) and quantum tunnelling. At finite temperatures this creep process is presumed to depend on temperature via a Boltzmann factor $\exp(-E_{\rm p}^{\rm f}/k_{\rm B}T)$, where the activation energy $E_{\rm p}^{\rm f}$ is the free energy associated with overcoming the potential, k_B the Boltzmann constant, and T the temperature, in equivalence with the flux creep theory developed for superconductors [33, 34]. Although creep is stochastic and consists of discrete vortex motion, collectively it contributes to a continuous $\dot{\Omega}_{\rm s}$. In the absence of a lag ω, vortices are equally likely to move to any direction – and thus there is no net vortex current – but as the lag increases the activation energy becomes lower for vortices moving radially outwards than inwards due to the Magnus force. This renders $\dot{\Omega}_{\rm s} < 0$ possible even if pinning is strong. Creep rate is expected to be very slow for small ω but plays a greater role as the rotational lag gets close to its critical value $\omega_{\rm cr}$. We will consider vortex creep in more detail in Section 5.2, where we show how it can be included in hydrodynamical models as a non-linear form of the mutual friction interaction, the basics of which we discuss next.

2.2. Vortex-mediated mutual friction

In mature, cold, neutron stars, excitations of the neutron superfluid can, at first approximation, be ignored. The main coupling between the normal and superfluid components is then via their respective interactions with the vortices. Thus the resulting force between the two components (which defines $\dot{\Omega}_{\rm s}$) will depend on the microphysical mechanisms on a vortex level, as well as spatial and temporal characteristics. A convenient description of this coupling borrows from the theory of terrestrial superfluidity and has the form of a mutual friction that acts on the two components [35, 36]. This is a long-wavelength approach that does not fully capture the physics at the microscopic vortex scale, but rather uses phenomenological parameters which must be determined from experiments or simulations of vortex motion. Nevertheless it offers a powerful tool to carry out coarse-grained hydrodynamical studies, on scales on which resolving individual vortices would be unfeasible.

In general, one assumes that dissipative forces on the vortices can phenomenologically be expressed as a drag force that is linear in the difference of velocities between the vortex and the normal component:

$$\boldsymbol{f}_{\rm d} = \mathcal{R}(\boldsymbol{v}_{\rm c} - \boldsymbol{v}_L) \tag{5}$$

with $\boldsymbol{v}_{\rm c}$ the velocity of the normal component and $\mathcal{R}$ the drag coefficient. The drag force leads to vortex motion with respect to the superfluid and to the aforementioned Magnus effect. Neglecting the minuscule vortex inertia, the forces in equations (4) and (5) must be balanced, which allows us to eliminate $\boldsymbol{v}_L$ in (5) and arrive at the mutual friction force per unit length on a vortex:

$$\begin{aligned}\boldsymbol{f}_{\rm MF}^{L} = {} & \rho_{\rm s}\frac{D^2}{1+D^2}\boldsymbol{\kappa} \times (\boldsymbol{v}_{\rm s} - \boldsymbol{v}_{\rm c}) \\ & + \frac{\rho_{\rm s}}{\kappa}\frac{D}{1+D^2}\boldsymbol{\kappa} \times [\boldsymbol{\kappa} \times (\boldsymbol{v}_{\rm s} - \boldsymbol{v}_{\rm c})]\end{aligned} \tag{6}$$

where we have defined a dimensionless parameter $D = \mathcal{R}/(\rho\kappa)$. In terms of vortex motion, the drag force introduces a component of the vortex velocity perpendicular to the superfluid flow, and the trajectories can be parametrized in terms of a dissipation angle [37] $\theta_D = \tan^{-1} D$, such that for $D = 0$ one has $\theta_D = 0$ and the vortex flows with the superfluid.

The same force as in (6) should be acting on the neutron superfluid (and an equal but opposite force on

the normal fluid) per unit volume, which can be found by multiplying $\boldsymbol{f}_{\rm MF}^{L}$ with the total vortex line length per unit volume.

The above description holds in the absence of pinning. For completely pinned vortices $\boldsymbol{f}_{\rm d}=0$, since $\boldsymbol{v}_{\rm c}=\boldsymbol{v}_L$, and the Magnus force should instead be balanced by the pinning force. If pinning is weaker, so that at any given time there is some 'free' vorticity which moves under the drag (e.g. creep current), equation (6) may still be a valid approximation as long as the vortex segment is not well within a pinning potential. In the opposite case the force balance equation should also take into account the pinning force at the location of the vortex and the mutual friction form will be different from (6). In the partial pinning regime, care should be taken when translating $\boldsymbol{f}_{\rm MF}^{L}$ to fluid force per unit volume as now only the 'free' total vortex length experiences the drag force. This can, for example, be introduced by re-scaling the mutual friction coefficient by the fraction of free over total vortex length $\gamma_{\rm fv}(\boldsymbol{r},t)$ [38, 39].

To connect to macroscopic quantities, such as a local angular velocity $\Omega_{\rm s}$ for the superfluid and the coupling between the two components, and hence to the observable parameters, appropriate averaging over length-scales that involve many vortices must be performed (the typical inter-vortex spacing is $d_v\sim\left[10^3(\Omega_{\rm s}/1\,{\rm Hz})\right]^{-1/2}$ cm). We present here, as a starting point, a special case where Newtonian treatment and a collection of assumptions lead to simple dynamical equations; other situations, for example the case of superfluid turbulence [40], will be considered later in connection to the observations.

Relaxing the rigid-body rotation constraint for the superfluid, but maintaining the view that vortices will be straight and parallel over the length-scale of interest, equation (3) can be replaced with a more realistic solution for the areal vortex density in the context of neutron stars:

$$\kappa n_v = 2\left[\Omega_{\rm s}+\varepsilon_n(\Omega_{\rm c}-\Omega_{\rm s})\right]+\varpi\frac{\partial}{\partial\varpi}\left[\Omega_{\rm s}+\varepsilon_n(\Omega_{\rm c}-\Omega_{\rm s})\right] \tag{7}$$

where ϖ is the cylindrical radius (taking Ω of all components to be along the z-axis). Whilst the last term accounts for potential differential rotation in the superfluid, another significant difference from equation (3) is allowing non-zero entrainment (Andreev-Bashkin effect) between the neutrons and the normal component. Although the nature of entrainment in the crust and core of neutron stars is different (as will be discussed in the following), the result is that currents of both constituents feed into their respective momenta, as reflected in the terms with the neutron entrainment coefficient ε_n in (7).

Under the same assumption of straight, aligned vortices, the mutual friction per unit fluid volume $\boldsymbol{F}_{\rm MF}$ is simply $n_v\boldsymbol{f}_{\rm MF}^{L}$:

$$\begin{aligned}\boldsymbol{F}_{\rm MF} &= \kappa n_v\rho_{\rm s}\frac{D^2}{1+D^2}\hat{\kappa}\times(\boldsymbol{v}_{\rm s}-\boldsymbol{v}_{\rm c}) \\ &+\kappa n_v\rho_{\rm s}\frac{D}{1+D^2}\hat{\kappa}\times[\hat{\kappa}\times(\boldsymbol{v}_{\rm s}-\boldsymbol{v}_{\rm c})]\end{aligned} \tag{8}$$

with $\hat{\kappa}=\boldsymbol{\kappa}/\kappa$ and n_v as in (7) assuming no pinning. For laminar flows of the two fluids, orthogonal to the vortex axes, we first note that $\hat{\kappa}\times[\hat{\kappa}\times(\boldsymbol{v}_{\rm s}-\boldsymbol{v}_{\rm c})]=-(\boldsymbol{v}_{\rm s}-\boldsymbol{v}_{\rm c})$, and so, in a sense, we can identify the first, non-dissipative, term of (8) with an equivalent Magnus effect, whilst the dissipative second term takes the form of a drag linear in the velocity difference. Habitually, the dimensionless, dissipative, mutual friction coefficient is defined in the literature as:

$$\mathcal{B}\equiv D/(1+D^2)\,. \tag{9}$$

Using the above form for the mutual friction we can arrive (e.g. by taking the curl of the Euler equations for the two fluids, see also [41]), at the equations of motion:

$$(1-\varepsilon_n)\dot{\Omega}_{\rm s}+\varepsilon_n\dot{\Omega}_{\rm c}=\frac{\kappa n_v\rho_{\rm s}\mathcal{B}}{\rho_{\rm s}}\left(\Omega_{\rm c}-\Omega_{\rm s}\right) \tag{10}$$

$$\dot{\Omega}_{\rm c}=-\frac{1}{(1-\varepsilon_p-\varepsilon_n)}\frac{\kappa n_v\rho_{\rm s}\mathcal{B}}{\rho_{\rm c}}\left(\Omega_{\rm c}-\Omega_{\rm s}\right)+F_{\rm sd}^{'}/\rho_{\rm c}$$

where we have used the fact that the entrainment coefficients ε_n and ε_p obey $\rho_{\rm s}\varepsilon_n=\rho_{\rm c}\varepsilon_p$. The above furthermore assumes – besides zero pinning – that the only other force is the external spindown $F_{\rm sd}^{'}$ that acts on the normal component; we ignore, for example, elastic forces in the solid crust.

With the more stringent condition of rigid-body rotation for the normal component, we identify its moment of inertia

$$I_{\rm c}=\int\rho_{\rm c}(\boldsymbol{r})\varpi^2{\rm d}V$$

and assuming constant moments of inertia in (1), equations (10) can be used to express the angular momentum evolution of the whole system, which now reads

$$I_{\rm c}\dot{\Omega}_{\rm c}+\int\frac{1}{(1-\varepsilon_p-\varepsilon_n)}\,\kappa n_v\,\mathcal{B}\,(\Omega_{\rm c}-\Omega_{\rm s})\rho_{\rm s}\varpi^2{\rm d}V=N_{\rm sd}. \tag{11}$$

In the above, ε_p, ε_n, n_v, $\mathcal{B}$, $\Omega_{\rm s}$ and of course $\rho_{\rm s}$ are all functions of $\boldsymbol{r}$ inside the star.

In general, the coupling between the two components will not be linear in their rotational lag $\omega=\Omega_{\rm s}-\Omega_{\rm c}$. For instance, its form will differ if vortices form a tangle, and even in (11) both n_v and

$\mathcal{B}$ can be functions of ω. The replacement $\mathcal{B} \to \gamma_{\rm fv}\mathcal{B}$ discussed for the case of partial pinning means that now the coupling strength depends also on the lag, as $\gamma_{\rm fv}$ increases with increasing ω. Furthermore, the drag coefficient itself can vary depending on the vortex velocity [42].

The full non-linear form of the mutual friction will modulate the superfluid response and we examine it in the context of its effect on the observed timescales and form of post-glitch recovery in Section 5. It is useful, however, to first consider the simpler case where all other factors in (11) are not lag-dependent: then (11) is linear in $\omega = -(\Omega_{\rm c} - \Omega_{\rm s})$ and it becomes transparent how it can lead to quasi-exponential relaxation as a response to the lag change that happens at a glitch. In return, for glitches where such exponential signatures are observed, the associated timescales can be used to constrain internal properties such as $\mathcal{B}(\boldsymbol{r})$. Taking a single, rigidly rotating, superfluid component $I_{\rm s}$, zero entrainment and ignoring the external torque, we can use the equations (10) to write an evolution equation for the lag ω:

$$\dot{\omega} = -2\Omega_{\rm s}\mathcal{B}\left(\frac{I_{\rm s}+I_{\rm c}}{I_{\rm c}}\right)\omega\,. \tag{12}$$

Making the reasonable assumption that $\omega \ll \Omega_{\rm s}$, so that we can ignore variations of $\Omega_{\rm s}$ during the glitch, we see that given an initial lag ω_0, the solution to (12) is an exponential relaxation, of the form $\omega = \omega_0 \exp\left(-t/\tau_c\right)$, with a characteristic timescale:

$$\tau_c \sim \frac{I_{\rm c}}{I_{\rm s}+I_{\rm c}}\frac{1}{2\Omega_{\rm s}\mathcal{B}}\quad . \tag{13}$$

Note that this is the dynamical timescale on which the two macroscopic components of the star recouple, and thus depends on the corresponding moments of inertia. It should not be confused with the (microscopic) timescale on which the vortex array couples to the normal components (e.g. in the lattice, or flux tubes and electrons in the core), which, as we shall see in the following, can be obtained directly from microphysical calculations and is used to determine the mutual friction parameter $\mathcal{B}$ [43].

It is also easy to see how the simplest three-component model would work, schematically illustrated in figure 4, in which glitches are driven by a region of the superfluid (of moment of inertia $I_{\rm gl}$) which is decoupled, or at least very weakly coupled, due to strong pinning. Taking the extreme case where $\gamma_{\rm fv} = 0$ for the $I_{\rm gl}$ component (and again ignoring entrainment), we have from (10) that $\dot{\Omega}_{\rm s}^{\rm gl} = 0$. At the other extreme, we assume that the rest of the superfluid, $I_{\rm s}^{\rm r}$, is more strongly coupled and has reached its equilibrium lag so that $\dot{\Omega}_{\rm s}^{\rm r} = \dot{\Omega}_{\rm c}^{\rm pre}$ before the glitch. While these components spin down together as $(I_{\rm c} + I_{\rm s}^{\rm r})\dot{\Omega}_{\rm c}^{\rm pre} = N_{\rm sd}$, the lag builds up in the $I_{\rm gl}$ region until pinning fails under the increasing Magnus force. Then catastrophic unpinning of vortices takes place and $I_{\rm gl}$ couples quickly to $I_{\rm c}$. On such a short timescale the effect of the external torque can be neglected in (10) and thus conservation of total angular momentum means $\Omega_{\rm c}$ increases. Once the spin-up phase is over, vortices have re-pinned and $I_{\rm gl}$ is again decoupled. The regions of $I_{\rm s}^{\rm r}$ that couple to $I_{\rm c}$ on a longer timescale than the glitch rise end up with a local $\omega < \omega_{\rm eq}$, and accordingly in these regions post-glitch $|\dot{\Omega}_{\rm s}^{\rm r}| < |\dot{\Omega}_{\rm c}|$. Thus, assuming $N_{\rm sd}$ constant, it follows that $|\dot{\Omega}_{\rm c}^{\rm post}| > |\dot{\Omega}_{\rm c}^{\rm pre}|$ (this is the observed glitch spindown rate change, $\Delta\dot{\Omega}_{\rm c} < 0$) which then recovers according to (11).

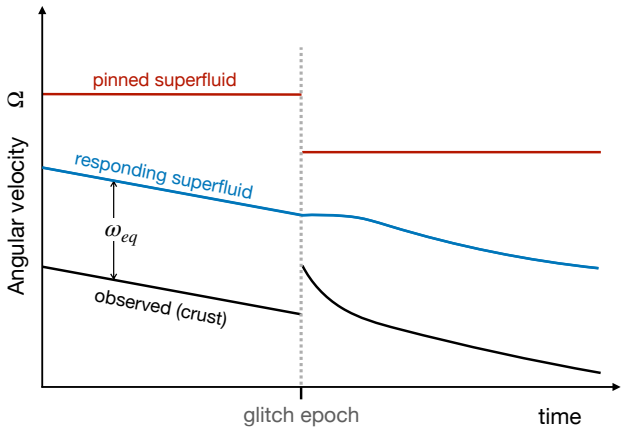

Figure 4. A sketch of a rudimentary three-component glitch model (see text for details). Amplitudes not in scale; for clarity we use arbitrary vertical offsets so that curves appear separate. The pinned superfluid region $I_{\rm gl}$ (red, top, line) stores angular momentum between glitches, that is then transferred to the normal component (black, bottom, line) abruptly during the glitch. Co-rotation of the two once the spin-up is completed is possible, but not necessarily achieved. The lag in the superfluid component $I_{\rm s}^{\rm r}$, which was spinning down together with the crust before the glitch (blue, middle, line), reduces from its equilibrium value. Consequently $I_{\rm s}^{\rm r}$ decouples at the glitch; its re-coupling drives the observed post-glitch relaxation.

2.3. Vortex pinning

Superfluid-driven glitch models rely on the existence of one or multiple regions within the star where vortex pinning occurs and bears a growing rotational lag until the glitch is triggered. Studying the exact pinning properties is important in order to determine the possible sizes and frequency of glitches, as well as to investigate vortex motion and the evolution of $\dot{\Omega}_{\rm s}$ between glitches.

In the inner crust, pinning can result from the interaction between a vortex and a proton cluster (ion) of the lattice. This interaction has two main contributions: the one arising from the velocity field of

the vortex line, and that of the change in condensation energy (i.e. the energy gain due to pairing, which is density dependent). Therefore the pinning energy per pinning site (i.e. the energy difference between configurations where a single ion is inside or outside the vortex core) depends on the difference in kinetic energy between the two configurations, and the cost involved with destroying superfluidity in the vortex core, related to the pairing gap of the nuclei and that of the neutron superfluid. The latter also defines the coherence length of the superfluid and thus the size ratio between a vortex core and nucleus. Another factor to consider in pinning energy is the deformation of the lattice, if nuclei can be slightly displaced due to their interaction with the vortex. Vortices are, however, macroscopic objects and so to understand the effect of pinning on the dynamical evolution the relevant quantity is not the force per pinning site but force per vortex length. Even if the pinning energy is large, the overall pinning force on the vortex can be very weak when averaged over many pinning sites along the vortex length and considering the lattice periodicities and orientation [44].

There are two calculations of importance to glitches. The first is the maximum pinning force on a vortex segment, which will define the maximum lag $(v_{\rm s} - v_{\rm c})$ that can develop before the vortex unpins due to the Magnus force (analogous in a sense to static friction). The second is the short-range force acting on a vortex as it moves through the pinning potentials of the lattice, which enters in the vortex equation of motion via the mutual friction term.

As far as the maximum pinning force in the crust is concerned, there are still significant computational uncertainties regarding even its sign (i.e. whether the force is attractive or repulsive). Recent static calculations, based on energy minimisation arguments and averaging over vortex configurations, obtain an attractive force throughout most of the crust [45]. On the other hand, three dimensional dynamical simulations (in the time-dependent superfluid local density approximation) recover a repulsive interaction when simulating the interaction between a vortex and a pinning centre [46], as do quantum mechanical calculations of the interaction, which also take into account shell-effects in the nuclei [47, 48]. Although there is general agreement on the overall magnitude of the maximum pinning force (expected to be of the order of 10^{18} dyn/cm), which is independent of the sign, and sufficient to explain large glitches, more work is clearly needed to understand vortex pinning in a realistic lattice.

Simulations of vortex motion in a lattice have been performed in two dimensions, to model mutual friction [49, 50]. While in these calculations the interaction between a single vortex and a nucleus is needed as an input, they allow one to extract robust prescriptions to model the fraction of free vortices close to the unpinning threshold, and recover standard mutual friction expressions [50]. The behaviour close to unpinning is of particular interest, as in this case a significant number of vortices are free. Depending on the physical conditions this can lead to either a steady outward creep of vorticity and $\dot{\Omega}_{\rm s} < 0$, or to vortices knocking each other on, leading to avalanches and a spasmodic spindown. However, to fully understand the impact of the additional degrees of freedom and vortex rigidity, three dimensional simulations will be needed.

Vortex pinning can also arise in the core of neutron stars, where the proton fluid is expected to be in a type-II superconducting state (although regions of type I superconductivity may also exist at high densities [43]). In a type-II superconductor the magnetic flux is quantised into proton vortices called fluxtubes. The neutron vortices also become magnetised due to entrainment between neutrons and protons (see Section 5.1). In this case it is energetically favourable for vortices and fluxtubes to overlap [51, 52], reducing the condensation and magnetic interaction energies (see [53] for a review). If vortices pin to fluxtubes, this can significantly increase the angular momentum available to exchange in a glitch, as one is also tapping the moment of inertia of the core [54]. Pinning to fluxtubes can strongly affect the mutual friction force and the response of the star to a glitch, as we shall see in the following [55, 56, 57, 58, 59]. Interactions between vortices and fluxtubes also contribute to the magnetic field evolution, as magnetic flux is 'pushed' out by the vortices [60, 61], or if fluxtubes force vortices to form bundles [62].

2.4. Glitch triggers

In most current models the transfer of angular momentum at a glitch is explained by previously pinned vortices being set free, moving, and their density reducing. This in turn assumes the existence of a trigger mechanism that causes unpinning of a large number of vortices, $N_v \gtrsim 10^{11}$ [63], on the short timescales of days or even seconds associated with the glitch spin-up (Section 5.3). The maximum lag $\omega_{\rm cr}$ that vortex pinning can sustain sets an upper limit for the size of glitches that can be powered this way. In practice however, the size of the glitch will depend on the lag that has been built up by the time the glitch is triggered, which need not be $\omega_{\rm cr}$. There is still no consensus on the exact trigger mechanism for glitches, and in fact it has been suggested that several may be acting, even in the same pulsar.

Crustquakes were one of the first trigger mechanisms to be proposed [18], and in some early models

the glitch was entirely attributed to the release of stress in the crust and subsequent change in moment of inertia. Whilst the latter hypothesis cannot explain the glitch activity of many pulsars [26, 64], crustquakes may play a role in triggering glitches by leading to vortex unpinning during the crust failing event, and have been suggested to be occurring in the Crab pulsar.

Another idea is that large scale unpinning is caused by vortex knock on, that is, few vortices unpin due to a local fluctuation, and as they move they trigger further unpinning, like falling dominoes. Quantum dynamical simulations of vortices in a rotating trap show that it is indeed possible for vortices to knock-on each other, essentially "colliding" with their nearest neighbours, forcing them to unpin due to the increased velocity field in their vicinity, and causing "avalanches" [65, 66]. This means that global stresses caused by the spin-down are released locally by nearest neighbour interactions, leading to a self-organised critical system and scale invariance, a point discussed further in Section 4.4 in relation to the statistical properties of glitches, such as their size distribution. It is important to point out that such simulations are carried out in a regime which is appropriate for laboratory experiments with Bose Einstein Condensates, but not necessarily for the neutron star problem, and the number of vortices simulated is several orders of magnitude less than what would be needed to simulate a glitch. Furthermore the calculations are not carried out in full 3D, thus not allowing for additional degrees of freedom which may allow vortices to creep out more freely by bending around pinning sites, rather than being forced to obey the kind of stick-slip dynamics observed in current models. Nevertheless, despite the caveats that must be addressed by future work, the mechanism is promising, and scaling up the results of current simulations allows to make predictions for glitch size and waiting time distributions to compare with observations [67, 68], which are presented in Section 4.

Avalanches are not, however, the only way that large-scale vortex unpinning can occur in a superfluid. The extra degrees of freedom allow for a number of additional instabilities with respect to a normal fluid, which can lead to both classical (hydrodynamical) and quantum turbulence over a vast range of scales [69, 70]. Transitions between a laminar and turbulent regime may trigger a glitch [71, 72], and also impact on the post-glitch recovery [73].

Even from this rather brief general summary, it will already be clear that glitch physics is a complex problem that ranges over the microphysics of transport properties, through the mesoscopic dynamics of vortex lines, and up to macroscopic hydrodynamics. With a multitude of inherent uncertainties, it is not feasible to build reliable glitch models on purely theoretical grounds. Instead, the subject has developed and continues to develop in concert with careful analysis of the growing observational data for glitches, and the layout of this review aims to reflect this. Laying the groundwork for this, the next section begins with a detailed description of pulsar timing and the detection of glitches.

3. Pulsar timing and glitch observations

The majority of pulsars are discovered and followed in the radio wavelengths, typically using single large parabolic antennae (of diameters 30 m or greater) or arrays of smaller antennae. However, neutron stars emit in many wavelengths and an increasing number of them are detected at higher frequencies, making monitoring by X-ray or gamma-ray instruments aboard satellites feasible. High-energy observations of pulsed emission are especially important for radio-quiet pulsars, or pulsars that are too faint in the radio and demand long observing times.

Observations can reveal the rotational evolution by a technique called pulsar timing, which tracks the neutron star revolution turn by turn. We outline below the basic data reduction process and the observational procedures used to identify and characterise glitches.

3.1. Pulse arrival times and pulsar monitoring programs

The rotation of a pulsar can be studied by recording the arrival times of its pulses at the telescope and – assuming the time between consecutive pulses is the pulsar spin period – connecting them to a model of the rotational phase [74, 75]. To start with, data from a relatively long observing session can be used to detect the precise periodicity. This is done either directly, or by Fourier transforming the data [76, 77]. The recorded time series is then folded at the rotation period to superimpose the light curves and create an average pulse profile. Regular monitoring sessions, which can be separated by days or weeks, are also designed to produce mean pulse profiles but are normally shorter – for example tens of minutes for typical radio telescopes or a few hours for X-ray timing observations, or even days for gamma rays [78, 79]. A time of arrival (ToA) of a fiducial point of rotational phase is calculated for each average pulse profile.

Times of arrival are found by cross-correlating the observed mean pulse profiles with a standard profile (usually a high signal-to-noise pulse profile obtained from a long observation) representative of the particular pulsar. The measured ToA will typically correspond to the arrival time of the pulse closest to

the middle of the observation stretch. For a given observing system, the observation length needed to achieve the desired accuracy for the ToA depends primarily on the pulsed flux density, but also on the pulse shape and its stability. Indicatively (but depending on specifications of the observing system and the pulsar considered), integrating over a few minutes of observations from a large, single-dish, radio telescope is enough to obtain precision at the $10^{-3}P$ level. Before these ToAs are standardized for use in a timing model, a number of effects must be taken into account. The ToAs are barycentred, that is, corrected for the Earth's motion (taking the location of the observing site also into account) by shifting the arrival time to what it would have been at the barycentre of the Solar System, which is used as an inertial frame of reference. Timing corrections include propagation dispersive effects, the Rømer delay to account for the difference in light travel time, and two general relativistic effects: the Shapiro delay (to account for spacetime curvature) and the Einstein delay (to account for time dilation).

There are several pulsar monitoring programs – with varying capabilities, scientific targets, duration, and scheduling – which produce regular ToAs and contribute to glitch studies. One of the first campaigns started in the early 70's at the Jodrell Bank Observatory (JBO) in the UK, and by today has provided nearly five decades of timing data for some pulsars. The JBO dataset includes almost continuous, daily, observations of the Crab pulsar [63, 80], and more than two decades of observations of over 500 pulsars [81, 82, 27]. Other observatories started some time later and can already offer 15-30 years of timing data; a prominent example is the Parkes radio telescope in Australia (e.g. [83, 84, 85]), which follows $\sim$ 200 pulsars and has also been used for important pulsar surveys. Decades of pulsar monitoring and glitch searches have also been carried out with the Nanshan 25-m radio telescope at Urumqi in China (e.g. [86, 87]) and at the Pushchino Observatory in Russia (e.g. [88, 89]). A decades-long campaign to observe the Vela pulsar has been carried out at the Mount Pleasant Radio Observatory (Hobart, Australia), which involves years of daily monitoring [90, 91, 92]. Most recent initiatives are also very important, as they offer not only an increase in the volume of data available but often also better precision observations. They include, amongst others, the timing program carried out in India with the Ooty Radio Telescope and the upgraded Giant Metrewave Radio Telescope (uGMRT) [93]; UTMOST in Australia [94]; the Thousand-Pulsar-Array project of MeerKAT in South Africa [95]; and CHIME in Canada [96, 97]. Finally, great contributions in long-term timing of X-ray pulsars and magnetars were achieved with the RXTE satellite (e.g. [98, 99, 100]) and, more recently, with NICER (e.g. [101]), as well as with the Fermi LAT (e.g. [78, 102]) for γ-ray pulsars. Most of the above programs produce weekly to monthly ToAs and only very few may commit to higher ToA cadences for some pulsars (like for example JBO, UTMOST and CHIME).

3.2. Timing model

A timing model describes the rotation of a pulsar by comparing the observed ToAs with their expected rotational phases for a given set of astrometric and rotational parameters. The evolution of the rotational phase can be calculated from a Taylor series expansion using reference values of the pulsar frequency $\nu = 1/P$ and its time derivatives at t_0:

$$\phi(t) = \phi_0 + \nu_0(t-t_0) + \frac{\dot{\nu}_0}{2}(t-t_0)^2 + \frac{\ddot{\nu}_0}{6}(t-t_0)^3 + \dots \,, \tag{14}$$

where ϕ_0 is an arbitrary phase reference. Other parameters, for example the pulsar proper motion or its orbital motion, if relevant, can also enter the timing model. $\phi(t)$ will increase by exactly one unit at every full rotation if the reference rotational parameters are correct, and should return an integer at the time of any observed ToA if t_0 is chosen as one particular ToA. Under the above definition, the angular velocity of the pulsar is $\Omega_c = 2\pi\,\mathrm{d}\phi/\mathrm{d}t$. Deviations from the assumed timing model are expressed as timing residuals, i.e. the difference between the observed and expected rotational phases, typically multiplied by the pulse period P to convert to time units. By minimising the timing residuals, updated rotational parameters can be obtained. Depending on the pulsar and the time span of the observations, it is typically enough to include only the first two (frequency ν and spindown rate $\dot{\nu}$) terms in equation (14), and sometimes $\ddot{\nu}$, in order to obtain small residuals scattered around zero.

However, when the effect of the secular spin-down described by the timing model is subtracted, the timing residuals often exhibit continuous, low-level fluctuations known as timing noise. Timing noise is more prominent in young pulsars [74, 83, 103] but is present for most, if not all, pulsars [82, 104], including millisecond pulsars which show the lowest levels of noise [105, 106]. Although timing noise mostly gives rise to smooth and somewhat slow irregularities in the residuals (in contrast with glitches which typically present a sharp change), it can have a faster component that presents as discontinuities of the frequency and/or the spindown rate. This kind of timing noise can be confused with glitches of small amplitude, especially when the cadence of observations is low [107, 104, 92]. Moreover, the presence of timing noise can compromise the accuracy by which glitch recoveries can be studied.

3.3. The glitch signature

When a given dataset contains a glitch, a timing model with constant reference ν_0 and $\dot{\nu}_0$, as in (14), will not provide an adequate description of all the ToAs. This is depicted in panel (a) of figure 5, where such a model is unable to describe well the rotation over the full time interval because of the presence of a glitch, which can be noticed as a sharp cusp. The positive and then negative slope of the residuals, before and after the glitch respectively, indicate that the ν_0 value obtained from the fitting procedure is higher, and then lower, than the actual spin frequency of the pulsar at those times. Similar signatures can be seen in figure 7, where a single, basic, timing model is used to produce timing residuals of a longer dataset containing three glitches.

If the timing model is optimised using only the ToAs before the glitch, the glitch signature will appear as a linear increase of the residuals $|\phi(t)|$ with time (panel (b) of figure 5). Because the residuals are defined as *data*−*model*, the trend is towards negative values, i.e. post-glitch ToAs arrive earlier than expected. Most glitches involve a spindown rate change, $\Delta\dot{\nu} < 0$, in addition to the change in frequency, $\Delta\nu > 0$. When that is the case, the residuals' evolution after the glitch will progressively resemble more a quadratic curve, as seen in panel (b) of figure 5. The actual changes in frequency and spindown rate for this glitch are shown in the last two panels of the same figure.

After a glitch, the evolution of the spindown rate (and consequently, of the spin frequency) often continues at a different rate $\ddot{\nu}$ than before the glitch. Commonly, some recovery towards the pre-glitch rotational parameters is observed, as revealed in the two lowest panels of figure 5 which present the frequency residuals (panel (c)) with respect to the pre-glitch timing model, and the evolution of the spin-down rate $\dot{\nu}$ (panel (d)). Similar plots for glitches, large and small, in other pulsars are shown in figure 6, which includes some representative examples of common, as well as more unusual, signatures of post-glitch recoveries. For many glitches the recovery process can be approximated by a combination of exponentially decaying terms and/or a linear decrease of $|\Delta\dot{\nu}|$ (see the top row of glitch examples in figure 6, and Section 5.4 for an in-depth discussion and interpretation of these features). It is thus practical to model the phase residuals attributed to the glitch as:

$$\phi_{\rm g} = \Delta\phi_{\rm g} + \Delta\nu_{\rm p}(t-t_{\rm g}) + \frac{1}{2}\Delta\dot{\nu}_{\rm p}(t-t_{\rm g})^2 + \\ + \frac{1}{6}\Delta\ddot{\nu}_{\rm p}(t-t_{\rm g})^3 + \sum_i \Delta\nu_{{\rm d}_i}\tau_{{\rm d}_i}(1-e^{-(t-t_{\rm g})/\tau_{{\rm d}_i}}) \quad , \tag{15}$$

which set in for $t \geq t_{\rm g}$, with $t_{\rm g}$ the inferred glitch epoch, in addition to the pre-glitch evolution of equation (14).

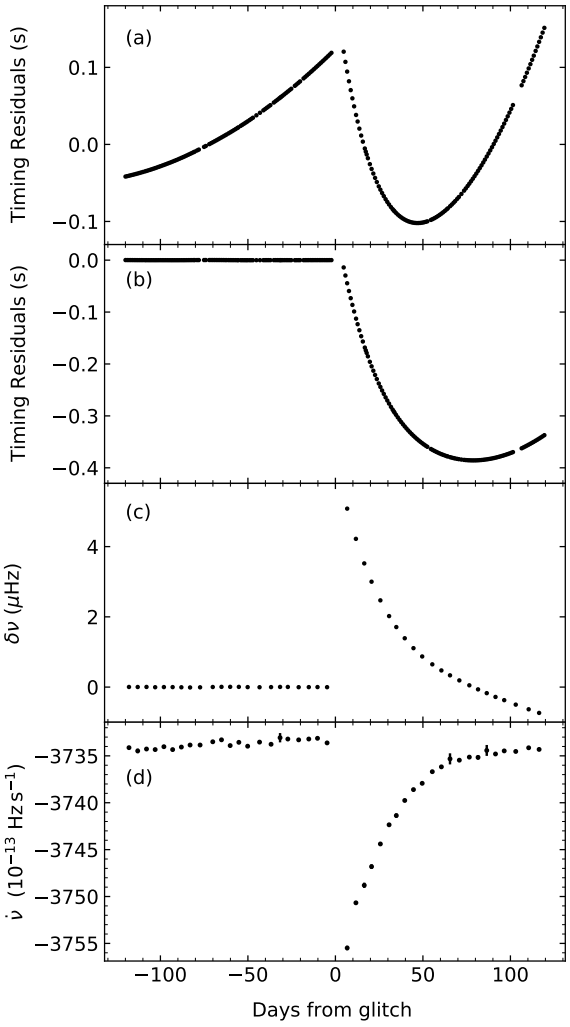

Figure 5. A glitch in the Crab pulsar ($t_{\rm g}$ = MJD 53067.1) with $\Delta\nu = 6.37\,\mu$Hz and $\Delta\dot{\nu} = -23.02\times10^{-13}\,{\rm Hz\,s^{-1}}$. The data was collected by the 42-ft telescope at Jodrell Bank observatory and presented in [27]. **(a)** Timing residuals relative to a single timing model fitted to the full interval. **(b)** Timing residuals relative to a timing model fitted to pre-glitch data only. **(c)** Frequency residuals with respect to the pre-glitch timing solution. **(d)** The evolution of the spin-down rate. In all cases the rotational model is that of equation (14) using two frequency derivatives. The gradual post-glitch relaxation is evident in the last two panels.

The terms with index p denote persisting changes that do not appear to recover in the duration of the dataset (e.g. as in the glitch of PSR B2334+61 shown in figure 6), whilst the last term captures any changes that show exponential decay on characteristic timescales $\tau_{{\rm d}_i}$ (visible after several glitches in figure 6, e.g. in PSR B0833-45, and also in the Crab glitch of figure 5).

For most pulsars, observations before and after the glitch can be days, even weeks, apart, which results in ambiguity in the exact epoch $t_{\rm g}$ when the glitch occurred. Typically this is constrained by demanding phase continuity through the spin-up ($\Delta\phi_{\rm g} = 0$),

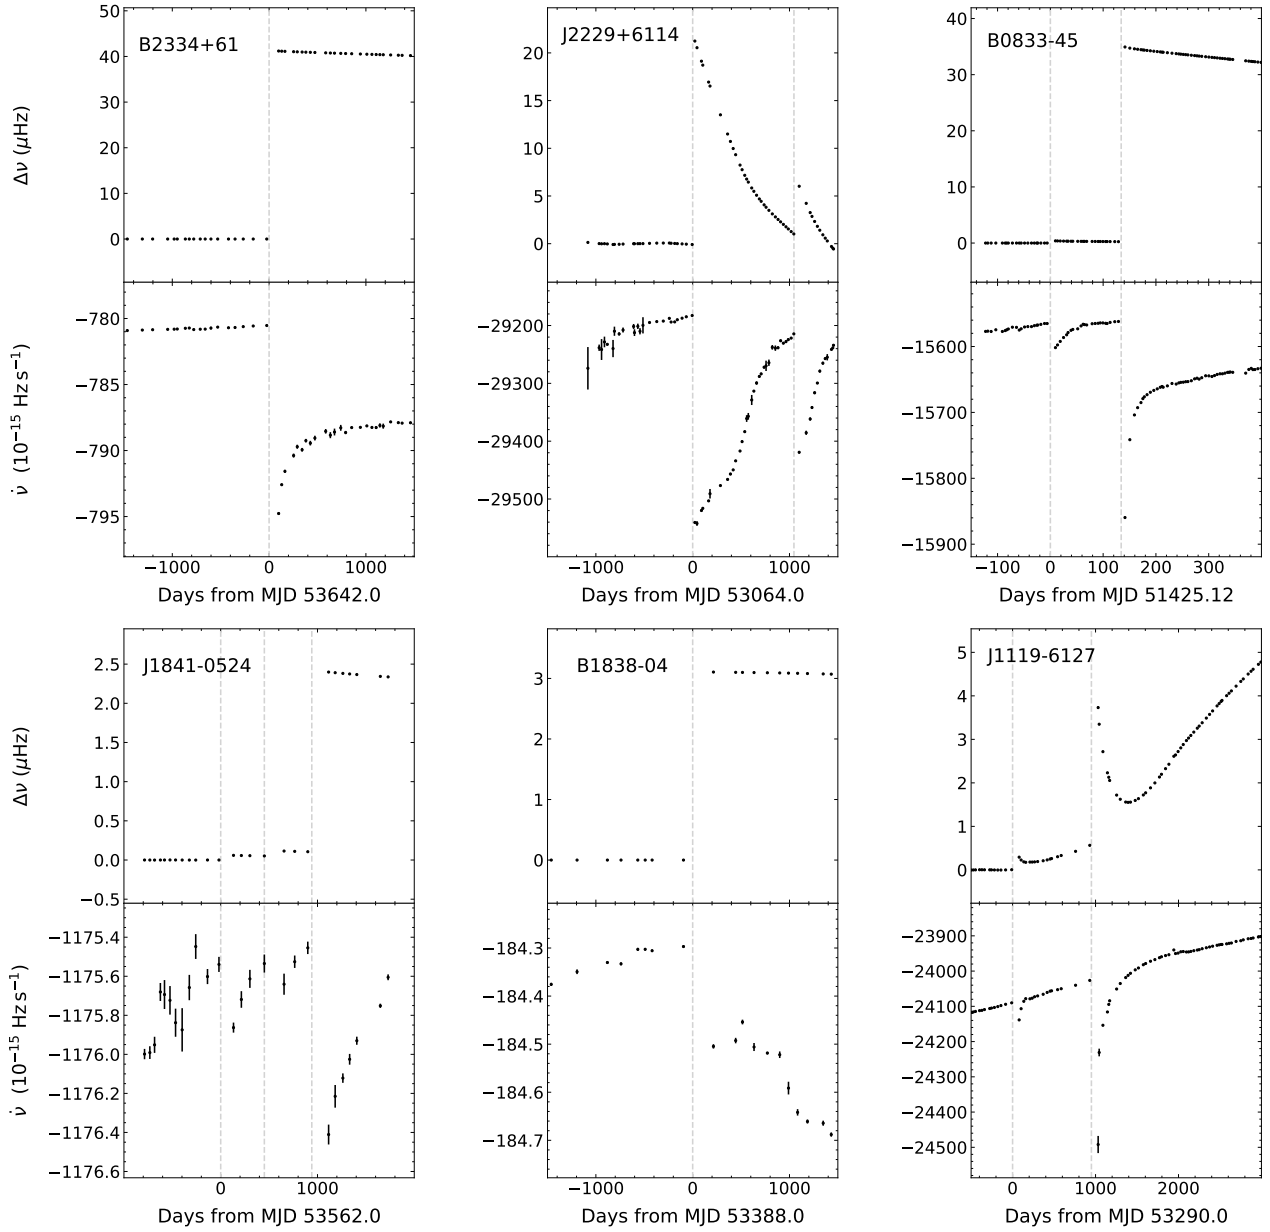

Figure 6. Examples of the diverse glitch phenomenology, from six different pulsars. Glitch epochs are denoted with dashed, grey, vertical lines. Top panels of each subplot: Frequency residuals with respect to a quadratic model fitted to frequency data between the start of the plot and the epoch of the first shown glitch. Bottom panels: The time evolution of $\dot{\nu}$. Both frequency and frequency derivative data were calculated via fits of a rotational model (where ν and $\dot{\nu}$ can vary, whilst $\ddot{\nu}$ is fixed) to small segments of data containing at least 10 ToAs and covering 50 to 500 d, depending on the pulsar. Note that while the time span for all windows is 3000 d, the one for PSR B0833−45 (the Vela pulsar) is only 550 d, and the one for PSR J1119−6127 is 3500 d. Three of the five glitches shown in the subplots at the top row are large, with $\Delta\nu$ well above 15μHz, and exhibit typical persistent steps $\Delta\dot{\nu}_{\rm p} < 0$. Note the particular recovery curve exhibited by PSR J2229+6114 in comparison to the more typical exponential-like recoveries seen in the other two pulsars. A late phase of linear (in $\dot{\nu}$) post-glitch recovery is clear after the second glitch shown for Vela. The glitches shown in the bottom subplots are smaller. The behaviour of PSR J1119−6127 is particularly interesting because of the atypical positive persistent steps $\Delta\dot{\nu}_{\rm p} > 0$, clearly visible after the second glitch shown. The recoveries of the other two pulsars are more obscured by timing noise. These glitches were published by [27, 108, 109, 92] and the data used here come in its majority from observations carried out at the Jodrell Bank Observatory in the UK, but also by the Parkes and Hobart radio telescopes in Australia (for PSR J1119−6127 and Vela, respectively).

however, especially for large glitches, pulse count can be lost and multiple solutions exist in the interval bracketing the glitch. Commonly, an epoch within the observational gap is used as the glitch epoch (typically the midpoint) and the term $\Delta\phi_{\rm g}$, which has no physical meaning, is added to emulate phase continuity through the glitch. For the opposite, much rarer, occasions when the glitch happens during an observing session, the uncertainty of the glitch epoch could be significantly reduced if the pulsar is bright, though it may still be larger than several pulsar rotations. Such observations might allow for the spin-up timescale to be resolved, or at least constrained, in which case the glitch model in (15) must include a term describing the glitch rise. This is usually also modelled as an exponential [109, 91, 110, 111].

Time differentiation of $\phi_{\rm g}$ provides the evolution of the spin frequency residuals after the glitch, and a second differentiation provides the evolution of the spindown rate residuals. These are residuals relative to the pre-glitch model in (14). The *instantaneous* changes at the glitch ($t = t_{\rm g}$) are

$$\Delta\nu = \Delta\nu_{\rm p} + \sum_i \Delta\nu_{{\rm d}_i} \tag{16}$$

$$\Delta\dot{\nu} = \Delta\dot{\nu}_{\rm p} - \sum_i \frac{\Delta\nu_{{\rm d}_i}}{\tau_{{\rm d}_i}} \quad . \tag{17}$$

The change in the second derivative at $t = t_{\rm g}$ can be calculated in the same way, but is rarely quoted. A persistent change $\Delta\ddot{\nu}_{\rm p}$ is sometimes included in the model and reported, usually for large glitches and models that include exponential recoveries. Positive and negative values for this term have been measured for different glitches in the same pulsar (e.g. [28]). This is contrary to what is usually observed for the other two persistent changes, which tend to be $\Delta\nu_{\rm p} > 0$ for standard large glitches and $\Delta\dot{\nu}_{\rm p} < 0$ for the vast majority of glitches [112, 113, 28].

3.4. Glitch detection

In typical timing residuals, what is usually visible is the continuous fluctuation caused by timing noise. If the data contain a glitch, a sharp feature like the ones in figures 5(a) and 7 might be striking out amongst the noise.

Generally speaking, the threshold for detecting a glitch will depend on its amplitude, especially relative to the intrinsic timing noise levels, the cadence of the observations, and the uncertainties on the derived ToAs [114, 27, 63, 115]. For example, short-term (say comparable to the observing cadence) timing noise can obscure the detection of small glitches; or low cadence observations can make it very hard to distinguish a small glitch from the noise, even if the noise levels are low and the fluctuations timescale long. On the other hand, large glitches (like those in the high end of the overall size distribution, see figure 10 in Section 4), are easy to detect because their signatures dominate the timing residuals for over hundreds of days. Even somewhat smaller glitches ($\Delta\nu/\nu > 10^{-9}$) are usually easy to detect visually in most datasets [114, 27].

At the largest glitch events, and unless cadence of ToAs is very high, the phase coherence implied by the modelling of ToAs prior to the glitch can be lost, thereby making the phases of the post-glitch ToAs ambiguous (see examples in [28]). This happens when the step in frequency, and therefore the drift of the residuals, is so large that one or more rotations could be unaccounted between the post-glitch observations. The higher the observing cadence is after the glitch, the larger the glitch has to be for this problem to arise. In some cases, finding the new frequency after the glitch can be difficult, and both low cadence ToAs or ToAs taken at too-regular time intervals can lead to wrong results [116].

For a given ToA cadence and precision, there is a glitch size below which the identification of glitches becomes uncertain, mainly because their signatures can be confused with noise [114, 63, 115, 92]. In fact, because of this, the detection of very small glitches has progressed slowly. Before Janssen & Stappers (2006) [114], who reported several very small glitches in regular radio pulsars, the only small glitch known was one in a millisecond pulsar, which was detected easily thanks to the pulsar's intrinsic rotational stability [117]. The incompleteness of the observed glitch sample at lower magnitudes poses a problem when trying to model the glitch rate, size distribution, and waiting times (i.e. the time intervals between consecutive glitches) distribution, thereby restricting the capacity of the observations to constrain physical models, for example about the trigger mechanism.

Espinoza et al. (2014) [63] explored the limiting factors in identifying the standard glitch signature presented in figure 5(b) and defined a minimum detectable size $\Delta\nu$ depending on timing noise, ToA accuracy and temporal density, as well as on the change $\Delta\dot{\nu} < 0$ that might accompany the glitch. The latter becomes important when the ratio $\Delta\nu/\Delta\dot{\nu}$ is comparable to or smaller than the typical time interval between ToAs, because then residuals grow to positive values in a short time compared to the observing cadence – the glitch can be mistaken for a change in $\Delta\dot{\nu}$ alone, or even a spin-down instead of a spin-up. They also developed an automated technique to search for glitches in large datasets, which has been applied to observations of the Crab and Vela pulsars, as will be detailed below [63, 92]. These studies focused on short time intervals (10-20 d) following a (potential) glitch and did not consider possible quasi-

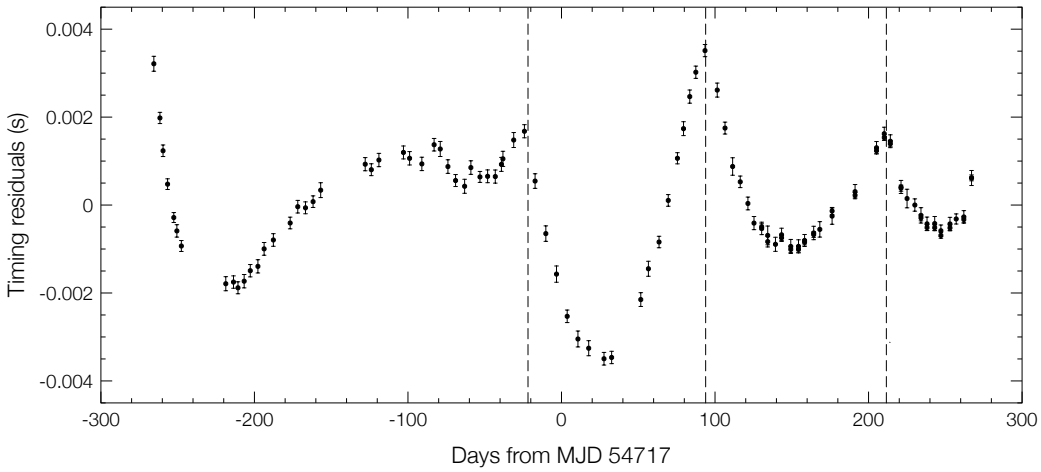

Figure 7. Three small glitches (epochs denoted by the dashed lines) in the rotation of PSR B1737−30, as seen in the timing residuals with respect to a model that uses ν, $\dot{\nu}$ and $\ddot{\nu}$ (fitted over the time interval shown). The more gradual variations of the residuals that can be seen before the epoch of the first glitch are likely caused by timing noise. The data was obtained by the Jodrell Bank Observatory and has been presented in [27].

exponential recoveries. This approximation is valid provided τ_{d_i} is long compared to the examined time window so that the effect of decaying terms can be absorbed in $\Delta\dot{\nu}_{\mathrm{p}}$. Whilst including decaying transients makes detectability limits less trivial, the characteristic exponential signature can actually help to distinguish small glitches from timing noise. The results confirmed (and quantified) the importance of monitoring pulsars frequently if the aim is to use observations to discover glitches.

Methods to detect glitches automatically face similar challenges as their visual counterparts but, once a set of rules has been defined, they have the benefit of being impartial. In the presence of noise, the identification of small glitches becomes a problem of probability assignment [115], hence automatic methods to find glitches become desirable. There are two, connected but distinguishable, kinds of automated methods. First, there are real-time glitch detection algorithms, as part of an observatory real-time data analysis (e.g. [118]). Their main objective is to infer the presence of a glitch from the smallest possible number of ToAs after the event, in order to make quick decisions with respect to additional (upcoming) observations such as multi-frequency follow up, increased monitoring etc. Secondly, there are algorithms intended for use on existing data with the aim of performing an as-exhaustive-as-possible search for glitches [63, 115]. So far, however, the vast majority of known glitches have been discovered either because they are visible in an inspection of the timing residuals or because trying to adjust a simple rotational model to the data fails.

The automated search technique which has been so far used on archival timing data of the Crab and Vela pulsars [63, 92] tests for the presence of a glitch after every available ToA by using standard fitting techniques and considering data 10-20 d around the epoch of the possible glitch. A glitch was defined as a step $\Delta\nu > 0$ that could occur simultaneously with a step $\Delta\dot{\nu} \leq 0$, without exponential recoveries (only the second and third terms in (15)). Other timing events such as "antiglitches" (that is $\Delta\nu < 0$ and $\Delta\dot{\nu} \geq 0$) as well as the case of $\Delta\nu = 0$ but $\Delta\dot{\nu} \neq 0$ were also included in the search [92]. The method has been applied to high cadence data (1 day between ToAs on average) for the Crab [63] and Vela [92] pulsars. The results included several glitches, both those known and some newly discovered, as well as a large population of smaller irregularities (glitch-like, antiglitch-like and $\dot{\nu}$-noise). These irregularities have different properties (e.g. size distribution [63]) and appear disconnected from the standard glitch population of these pulsars. This finding suggests that the truncation of the glitch size distribution towards small amplitudes might not only be an effect of observational biases, as previously thought, but a real feature of the glitch mechanism as well.

A different automated method [115] uses a hidden Markov model to infer the pulse frequency evolution on relevant time scales (such as those determined by timing noise and glitches), and aims at distinguishing glitches from timing noise by comparing the Bayes factor of models with and without a glitch. The method was tested on simulated ToA data, in which glitches were introduced as instantaneous changes in ν and $\dot{\nu}$, whilst the underlying rotational model involves one frequency derivative and a stochastic part to simulate timing noise. The detector has also been applied on real timing data and can be used to put

upper limits in the glitch size detection threshold [85], which is important for understanding the glitch sample completeness.

While these methods process ToAs to search for glitches, observations in gamma rays (and perhaps in X rays too) can be explored for glitches directly, without the need to create ToAs [119]. This is possible by exploiting the effectively continuous capture of photons provided by instruments like the Fermi-LAT, which survey the whole sky every few hours [120]. A rotational phase is assigned to each photon using a model like the one in equation (14). If the model is correct, the majority of the photons will align and distribute on a range of phases which is determined by the pulse shape. But if a glitch was present in the data, there will be a sudden misalignment of the photons, followed by a linear drift (just like figure 5(b), see [78]).

3.5. Glitch parameterization

The spin-up amplitude, that is the glitch size $\Delta\nu$, as well as any changes $\Delta\dot{\nu}$ and $\Delta\ddot{\nu}$ associated with the glitch, can be deduced by extrapolating the rotational parameters of a pre- and post-glitch timing model to the inferred glitch epoch $t_{\rm g}$ and taking their difference. This is a local approach for which use of relatively short pre- and post-glitch datasets is required, covering about 100-200 days or less (to avoid significant contamination from timing noise) [121, 27]. For longer time intervals the simple spin-down model might not suffice to describe the data adequately. If the glitch-induced changes show some recovery, the rotational model for the post-glitch observations will include decaying terms like those in (15) and the analysis must involve enough post-glitch data to faithfully capture those terms. In general, when trying to characterise the glitch recovery, long post-glitch intervals must be used – ideally the entire post-glitch data available.

This method of measuring glitch parameters is implemented with the widely used software package TEMPO2 [122], in which the timing model is defined as the sum of a pre-glitch model and a glitch contribution, as in equations (14) and (15) respectively. The software can fit for a model which includes several glitches in a long data set, and allows all, or a selected subset, of parameters to be varied to minimize the phase residuals. Traditionally, the glitch epoch $t_{\rm g}$ had to be pre-defined by the user.

As discussed in Section 3.3, it is often the case that $t_{\rm g}$ cannot be accurately determined from the observations; this introduces uncertainties in the inferred glitch parameters, including the main glitch parameter $\Delta\nu$ [100], and can become the dominant source of error for all measured parameters if the observational gap surrounding the glitch is large [28]. Such extra uncertainties are not taken into account by the standard TEMPO2 software and are often not reported in glitch measurements found in the literature; quoted $1\,\sigma$ errors generally underestimate the true uncertainties. When decaying terms are considered, several fitting algorithms require starting values for the timescales $\tau_{{\rm d}_i}$, which need to be carefully chosen in order to achieve convergence (e.g. [28]). When modelling a long dataset with a single model (that can contain one or more glitches) timing noise becomes important and can affect the results if it is not accounted for.

Glitch parameter inference has traditionally been done using weighted least-square minimization methods as in TEMPO2 [122], but other pulsar timing analysis software, or extensions, such as ENTERPRISE [123], which use more efficient methods to explore the parameter space and may offer more realistic uncertainties, are becoming increasingly popular (e.g. [111]). Bayesian approaches are gaining ground for estimation of the likelihood of the fitted models. As an example, the software package TEMPONEST uses Bayesian inference and nested sampling to explore the parameter space and assess the likelihood of tested solutions [124] (or see [107, 83, 104, 116] for the use of TEMPONEST on glitches). Both ENTERPRISE and TEMPONEST allow the timing model (which might include glitches and their recoveries) to be adjusted simultaneously with other parameters, such as those used to model timing noise or changes in the pulse shape, including ones describing stochastic processes [125, 111]. This allows the determination of a timing model that incorporates glitches, their recoveries, timing noise, and any other factors known to affect the rotation (e.g. [107, 111]), and can be used on very long datasets. An important advantage of packages like TEMPONEST is that they facilitate comparison of different models based on their Bayesian evidence.

Characterising the glitch recoveries presents various challenges and results can be ambiguous. First, whilst the standard fitted model of equation (15) might be adequately capturing the main behaviour after most glitches (i.e. returns 'flat' residuals scattered around zero), other times there are clearly remaining features in the residuals which could be unmodelled timing noise but could also be an unaccounted part of the post-glitch recovery. The elongated-S shape seen in the $\dot{\nu}$ curve after the first glitch of PSR J2229+6114 in figure 6 is potentially one such example. Secondly, data do not always allow the distinction between models that include different combinations of terms (in their nature and/or number) in (15). Scarcity of post-glitch data or the occurrence of a new glitch are two of the top limiting factors when it comes

to differentiating recovery models [126, 108]. Last, different functional forms for the recovery are rarely tested against equation (15), even though they might also be physically motivated and in some cases describe observations equally well [127, 73].

When fitting exponentially decaying terms, it is often necessary to first produce secondary data like $\nu(t)$ or $\dot{\nu}(t)$ which are calculated by fitting (14) to short subsections of the ToA data. Then these can be used to access the likely timescales τ_{d_i} involved in the recovery, which serve as starting values for the fitting process of the ToA dataset [28, 128]. Algorithms that are better suited to explore broad multidimensional parameter spaces and are designed to perform one overall fit to a long dataset might offer a better approach. Whichever the method, rapidly decaying terms are hard or often impossible to detect due to paucity of data. If a short-term recovery took place but was missed by observations, the inferred $\Delta\nu$ and $|\Delta\dot{\nu}|$ at the glitch epoch are underestimating the real glitch impact. On the other hand, recoveries on long timescales and persistent changes can be hard to quantify unless there is good ToA coverage for an extended period of time (which can be many years), without interruption from a subsequent glitch. The effect that the length of the post-glitch dataset has on the inferred recovery parameters means that earlier results, using restricted post-glitch data, should be considered preliminary and preferably be reviewed when more data accumulates [112, 27].

3.6. Ordinary glitches, anti-glitches and slow glitches

A rapid, positive, change in frequency is the minimum requirement for a timing feature to classify as a potential glitch. For $\sim$ 70% of detected glitches a change $\Delta\dot{\nu}$ has also been measured, which was $\Delta\dot{\nu} < 0$ in about 80% of the cases according to the Jodrell Bank Glitch catalogue (`http://www.jb.man.ac.uk/pulsar/glitches.html`, [27]). The glitches for which a $\Delta\dot{\nu} > 0$ is inferred are generally small ($\Delta\nu < 10\,\mu$Hz), and it is unclear if their properties are otherwise different from those of ordinary glitches. In the presence of timing noise, measurements of $\Delta\dot{\nu}$ carry more uncertainty, an effect which could explain at least part of the anomalous $\Delta\dot{\nu} > 0$ findings.

The vast majority of sharp frequency changes observed in pulsars are events for which $\Delta\nu > 0$. The lack of observed events with the opposite signature ("anti-glitches", where $\Delta\nu < 0$) is not a result of observational bias, as these would have been detected in the same way – either by visual inspection of timing residuals or automated methods. Only a few negative frequency changes have been reported, in two magnetars [129, 130, 131, 132, 133] and two accreting pulsars [134, 135].

Magnetars in general show a more complex timing behaviour compared to rotationally-powered pulsars (see Section 6.2), which is not restricted to the display of anti-glitches. One detected spin-down event was interpreted as a microglitch [129], which are small irregularities showing as either spin-up or spin-down, of comparable amplitudes and numbers, known to populate pulsar rotation – possibly as part of their timing noise, as mentioned already for the Crab and Vela pulsar [29, 63, 92]. Glitches and anti-glitches in the magnetar 1E 2259+586 have indeed very similar amplitudes [98]. Note that there may be more $\Delta\nu < 0$ events in other magnetars, but are harder to confirm [136, 137]. Magnetars exhibit in general stronger timing noise than ordinary pulsars, which can be confused for small glitches and anti-glitches [92], especially when the cadence of the observations is not very high, as is in fact the case for most magnetars. Hence it is possible that at least some of the reported magnetar anti-glitches are due to timing noise, and likely of magnetospheric origin. Radiative changes are sometimes, but not always, observed to coincide with such events (e.g. [130], see Section 6.2). Alternatively, anti-glitches could be driven by the evolution of the internal magnetic field [138, 139] or by processes related to the neutron superfluid similarly to ordinary glitches [140].

Glitches and anti-glitches in accreting pulsars might not be unexpected given their relatively high $|\dot{\nu}| \sim 10^{-11}\,\mathrm{Hz\,s^{-1}}$. A result of torques exerted by the accretion process, $\dot{\nu}$ may be either positive or negative and can remain stable for years [141]. There is a small number (< 10) of sudden ν changes reported (some with less certainty than others) for a handful of such neutron stars [142, 143, 134, 144, 135]. Observed sizes are $|\Delta\nu| > 1\,\mu$Hz but the relatively low cadence of timing observations could hinder detection of smaller events. No radiative counterparts have been found associated with the rotational events. The few anti-glitches among those events were seen in the rotation of two sources that are spinning up due to accretion ($\dot{\nu} > 0$) [134, 135], which led to the suggestion that they are caused by a similar mechanism to the one operating in standard spinning-down pulsars [141, 135]. Indeed, simulations of pinned vortices in a spinning-up container show that inward-travelling avalanches (leading to $\Delta\nu < 0$ jumps, i.e. anti-glitches) will occur by the same process of vortex unpinning and knock-on effects as for normal glitches [145].

While the glitches and anti-glitches discussed above involve steps $\Delta\nu$ that are unresolved to the accuracy of the data (constrained to occur in tens of seconds in one case), there is a type of rotational feature in which the spin frequency increases gradually over hundreds of days (compared to predictions of

the preceding timing solution). These are sometimes called slow glitches and follow a rapid, sometimes unresolved in time, decrease of the spin-down rate (hence the apparent frequency increase) [146, 28, 147]. Although slow glitches have been reported for at least 7 pulsars (e.g. [147]), in the three most studied pulsars with regards to this (PSRs B0919+06, B1642−03, and B1822−09, [148, 149, 146]), it was found that the events are part of a quasi periodic modulation of the phase residuals. Such a quasi-periodicity is not uncommon among pulsars [82, 150]. Furthermore, the timing modulation of PSR B1822−09 appears to be caused by switching between states of different spindown rate, and these $\dot{\nu}$ switches correlate with observed changes of the pulse profile [150]. Pulse profile changes have been reported for other slow-glitching pulsars too [147]. All this suggests that slow glitches are produced by $\dot{\nu}$ changes, possibly connected to the magnetospheric torque, and belong to the kind of irregularities that fall under the umbrella term timing noise instead of relating to the mechanism responsible for standard glitches [151, 150, 28]. They should also not be confused with glitches for which part of the rise has been resolved and present an extended period of spinning-up (Section 5.3) as both the timescales and, significantly, the accompanying $\dot{\nu}$ evolution, are very different.

4. Glitch magnitudes ($\Delta\nu$, $\Delta\dot{\nu}$) and interglitch waiting times

With nearly 600 glitches detected to date, and many pulsars showing multiple glitches, it is possible to study statistically their overall properties. Several important trends have emerged as the examined glitch sample grows [152, 153, 27, 154]. In this section we will present the properties of the general population, focusing on the glitch rate, the instantaneous changes $\Delta\nu$ and $\Delta\dot{\nu}$, interglitch waiting times, and glitch activity (which is defined as the accumulated effect of glitches in a given pulsar over the years), as well as some results for frequently glitching pulsars and those of particular interest. The relation between these basic observables and physical parameters of glitch models will be discussed, with insights gained on aspects such as the extent and location of the superfluid reservoir or the glitch trigger highlighted.

4.1. Glitch rate of occurrence

Most glitches have been seen in the rotation of ordinary, isolated pulsars (which is the largest population of observed neutron stars). Only two glitches have been detected in millisecond pulsars (MSPs, [117, 155]), despite extensive monitoring campaigns of 60 or more MSPs as part of the effort to use them to detect gravitational waves [156]. The reason for this could be that their typical spin-down rates are too low to drive frequent glitches [157], but could also be related to the fact that these objects are usually recycled pulsars hence a much older population compared to the standard glitching pulsars. Glitches appear common to magnetars [158, 98], and have also been reported in Rotating Radio Transients (RRATs, [159]) and Compact Central Objects (CCOs, [160]). The CCO 1E 1207.4−5209 has a very low spin-down rate but showed rotational irregularities attributable to two small glitches. Such activity is up to 7 times larger than what is expected from extrapolation of the trend defined by the general pulsar population [160]. CCO's have rather weak surface magnetic fields in general, but some of their properties could be explained by the emergence of a stronger field buried under the surface [161, 162]. On the other hand, the RRAT PSR 1819−1458 has a fairly large surface magnetic field ($\sim 5 \times 10^{13}$ G) and is one of the two pulsars that have exhibited a recovery with $\Delta\dot{\nu}_{\rm p} > 0$; the other one being the high magnetic field, magnetar-like PSR J1119−6127 whose two glitches with $\Delta\dot{\nu}_{\rm p} > 0$ are shown in figure 6 [108, 163]. Magnetars often show bursting activity and unusual rotational evolution associated with glitches. Therefore all these three families of neutron stars present uncommon glitch properties, likely associated with their magnetic fields (see also Section 6.2).

Glitches are generally rare events, though their average occurrence rate varies widely from pulsar to pulsar. In some glitching neutron stars consecutive spin-ups are sometimes just a few days to months apart, yet in others only a single glitch has been detected despite decades of observations. The great majority of pulsars have never been observed to glitch, even though they have a similar distribution of total observing time spans as for glitching pulsars. The fraction of glitching pulsars is smaller when calculated only over the population of neutron stars observed for less than ~ 10 years, as expected.

An average glitch rate for a given pulsar is defined as

$$r_{\rm g} = \frac{\text{number of detected glitches}}{\text{total observing time}} \qquad (18)$$

where the denominator is the total observing time over which ToAs are dense enough to allow detection of glitches beyond a certain size. Usually $r_{\rm g}$ can only be approximated because of the difficulty in assessing how complete the glitch sample for a given pulsar and observing period is. Furthermore, the glitch rate is variable over time for many sources, therefore $r_{\rm g}$ misrepresents the true nature of their glitch activity (see middle and bottom panels in figure 8 for two examples). Only for some pulsars are the time intervals between glitches rather regular, corresponding to a

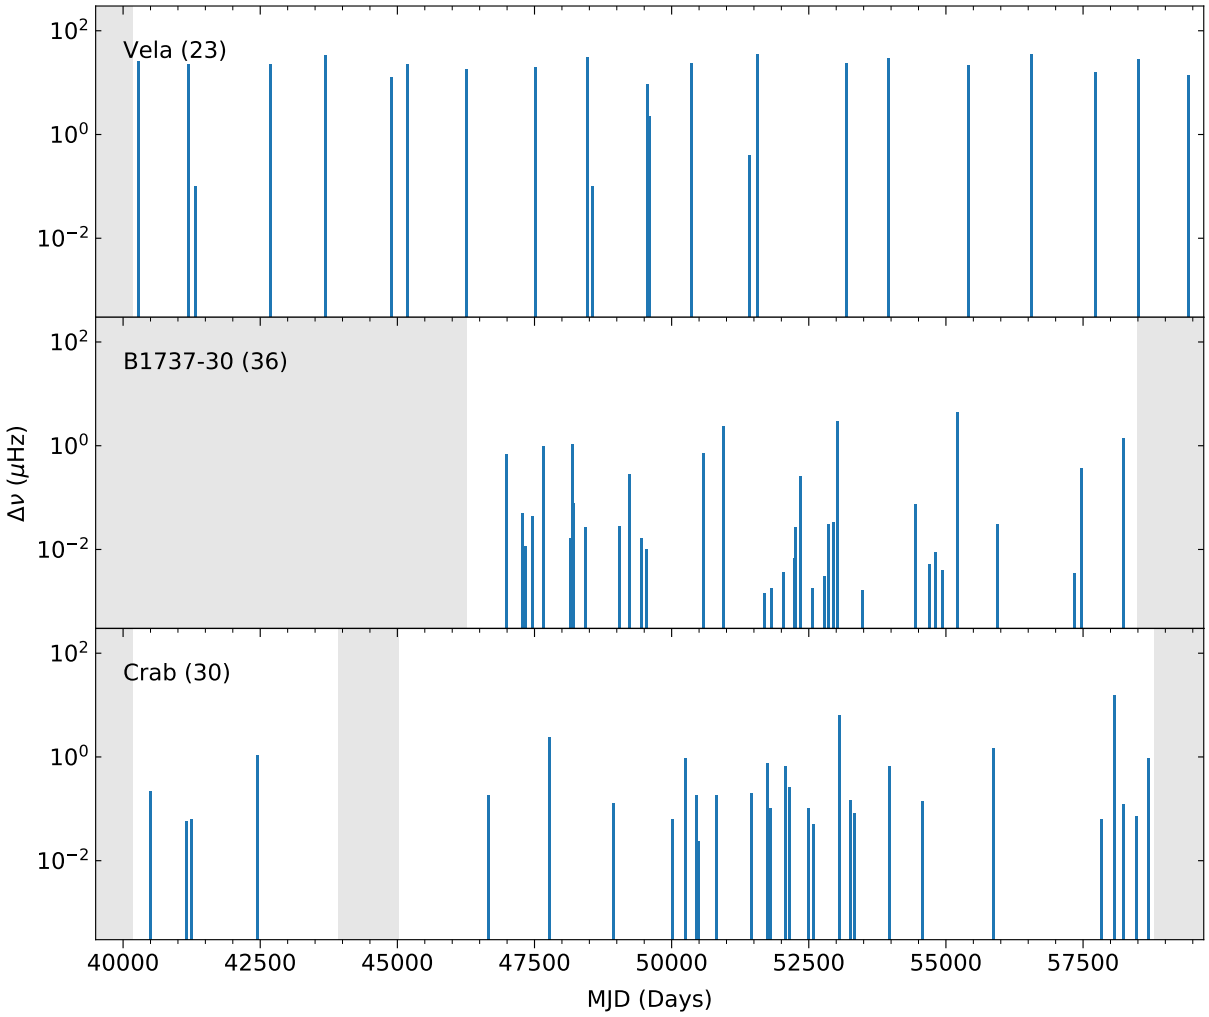

Figure 8. The temporal distribution (shown with vertical bars) and size of glitches (the length of the bars, in logarithmic scale) up to MJD 59000 in three pulsars that have frequent glitches. Grey shaded areas represent time intervals for which observations were not available. Top panel: the Vela pulsar, showing mostly large glitches at remarkably regular time intervals, maintaining an approximately constant glitch rate r_g for over five decades. Middle panel: PSR B1737−30 (J1740−3015) has a high glitch rate that varies on relatively short (order 10^3 days) timescales. The average r_g would have been under/overestimated if observations covered only a few years around a period of lower/higher activity. Bottom panel: The Crab pulsar, whose glitch rate also varies with time. The lack of glitches around MJD 43000-46000 is in part due to scarcer monitoring, but very frequent observations began at MJD ∼46300 and the sample is believed to be mostly complete from then onward. The clustering of glitches between MJD 50000–54000, which is preceded and followed by a period of less activity, is at odds with a random distribution of interglitch time intervals [80].

somewhat constant rate over the years of observations. The Vela pulsar is such an example, as shown in the upper panel of figure 8.

For some of the most frequently glitching pulsars, the distributions of interglitch waiting times can be modelled as an exponential, consistent with a random Poissonian process ([67], Section 4.3). If that is the case then the underlying constant Poisson rate (i.e. the average number of glitches per time) can be estimated for pulsars with multiple glitches after sufficiently large total observing times. The Crab pulsar however, whose rotation is being followed since its discovery and with particularly high observing cadence since mid-1980s, casts doubt on the homogeneous (constant-rate) Poisson process hypothesis. For about 11 years, between 1995 and 2006, it displayed a significantly increased activity (bottom panel of figure 8) [80]. This could indicate a time-varying Poisson rate if glitches are independent events, or clustering due to correlated events as in the case of earthquakes where the principal event(s) is surrounded by pre- and after-shocks [164]. Such clustering might also be expected if glitches are triggered by vortex avalanches as will be discussed in the following.

Keeping in mind these limitations, it is useful to examine how r_g varies in the pulsar population

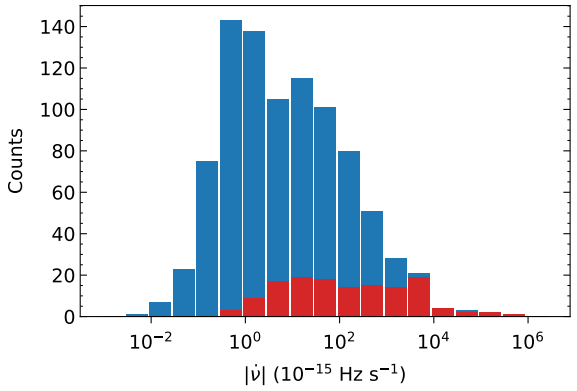

Figure 9. Histogram (in a logarithmic scale) of $|\dot{\nu}|$ values for 898 pulsars whose rotation has been searched for glitches [157]. Marked in red are the 137 pulsars in the sample for which glitches have been found. The percentage of pulsars seen to glitch in each bin decreases with decreasing $|\dot{\nu}|$. Data kindly provided by J.R. Fuentes.

as a whole. The trend that is revealed this way is a decrease of $r_{\rm g}$ with decreasing spin-down rate $|\dot{\nu}|$, as well as with increasing characteristic age $\tau_{\rm ch} = -\nu/(2\dot{\nu})$ [152, 27]. Currently, to the best of our knowledge, pulsar B0410+69 is the radio source with the lowest spin-down rate ($\dot{\nu} \simeq -5 \times 10^{-16}\,{\rm s}^{-2}$) ever seen to glitch. At the other extreme, the vast majority of very high $|\dot{\nu}|$ pulsars ($|\dot{\nu}| \gtrsim 10^{-11}\,{\rm s}^{-2}$) have displayed glitches, as seen in figure 9. The pulsar with the highest glitch rate, PSR J0537−6910 (with over 3 glitches per year) [100], has the second highest $|\dot{\nu}|$ known. Although $r_{\rm g}$ might be more strongly underestimated for older pulsars, with low spin-down rate, as these tend to show mostly small glitches which are harder to detect, this effect alone cannot explain the observed correlation between $r_{\rm g}$ and $|\dot{\nu}|$ (and anti-correlation with $\tau_{\rm ch}$). The existence of such a correlation is also supported by theoretical glitch models. The two commonly invoked mechanisms for glitch triggers, the unpinning/avalanches of vortices and (spin-down driven) crustquakes, will be occurring at a rate that depends on the external driver, in this case the magnetospheric torque that sets the average $|\dot{\nu}|$.

4.2. The changes in frequency and spindown rate

The measured changes in pulsar frequency at the moment of a glitch cover a very broad range of sizes, from about 10^{-5} to few tens of μHz, or, put in relative sizes, $10^{-12} \lesssim \Delta\nu/\nu \lesssim 6 \times 10^{-4}$. The distribution of glitch sizes, displayed in the first panel of figure 10, appears bi-modal (or multi-modal) [157, 165]. Using a sample of over 500 glitches, the resulting $\Delta\nu$ distribution in logarithmic space can be described as a mixture of two Gaussian components. Small and

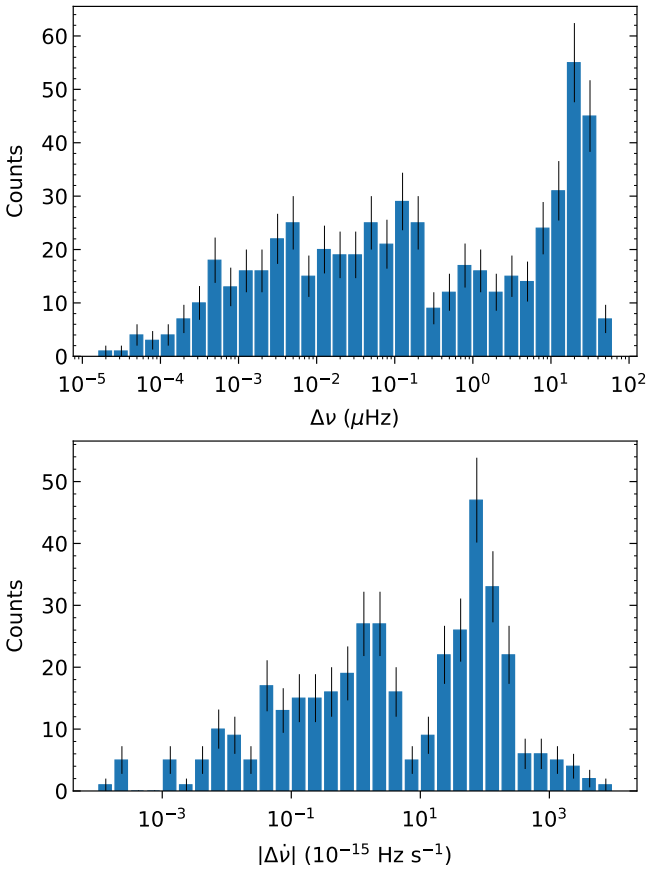

Figure 10. The distribution of glitch amplitudes $\Delta\nu$ (upper panel) and $|\Delta\dot{\nu}|$ (lower panel, data shown only for glitches where $\Delta\dot{\nu} < 0$) in logarithmic scales. Uncertainties are calculated as $\sqrt{N}$, where N is the number of elements in each bin. Glitch parameters were retrieved from the online Jodrell Bank Glitch Catalogue [27].

intermediate size glitches make up the first, wide, component, with mean $0.032\,\mu$Hz, whilst a second, narrow, component includes the largest events and has a mean of $18\,\mu$Hz [154]. The peak at large glitch sizes ($\gtrsim 10\,\mu$Hz), which are sometimes called "giant" glitches, is produced mainly by a modest subset of pulsars, typically young objects of characteristic age similar to the Vela pulsar ($\sim 10^4$ yr) [27, 28, 157]. Currently, about 20% of glitching pulsars have shown at least one large glitch. The bi-modality of the $\Delta\nu$ distribution is interesting from a theoretical point of view, as it suggests differences in the processes at work in "giant" glitches compared to smaller ones. It has been proposed that the glitch trigger might differ in those two cases; another possibility is that the response of the superfluid varies with glitch size, leading to an amplification of large glitches and the second peak in the distribution. We come back to this point in the following, when we discuss how large spin-ups might proceed faster and result in decoupling of a larger fraction of the star.

Observational biases and timing noise affect mostly the lower end of the $\Delta\nu$ distribution (Section 3.4) but pose greater limitations on the accuracy of $\Delta\dot{\nu}$ measurements. The $\Delta\dot{\nu}$ distribution also presents bi-modality (figure 10, lower panel), likely a result of an apparent correlation between frequency changes $\Delta\nu$ and spindown rate changes $\Delta\dot{\nu}$, as seen in figure 11.

The (physical) significance of the observed $\Delta\nu$-$\Delta\dot{\nu}$ correlation is hard to address because several systematic biases might be at work, especially for glitches of small amplitude where the correlation appears tighter. As pointed out in Section 3, glitches with small $\Delta\nu/\Delta\dot{\nu}$ ratio, and in the absence of clear exponential-like recoveries, can be confused with changes in $\Delta\dot{\nu}$ alone, or even spin-down rather than spin-up events, and not get classified as glitches. Low cadence of ToAs exaggerates this effect, as demonstrated by the long-dashed limiting line in figure 11, for which monthly ToAs are considered. Fortunately a very large percentage of pulsars known to glitch are monitored more closely (for example every 10-15 days), but there is probably still a contribution of this bias to the observed $\Delta\nu$-$\Delta\dot{\nu}$ relation. Perhaps of more relevance to the trend in figure 11, the ToA error or timing noise will place an upper bound to the resolved parameter space, which is exemplified by the dotted line for a ToA uncertainty or RMS of the timing residuals (whichever is greater) of 100 μs. Additionally, small changes in $\dot{\nu}$ are intrinsically hard to detect in the presence of timing noise, which most likely results in an under-population of the lower part of figure 11. Whilst this bias exists irrespectively of the $\Delta\nu$ amplitude, it likely affects mostly small glitches, as large events tend to have detectable spin-down changes. All these issues must be quantitatively factored in – which is not trivial as limits vary from pulsar to pulsar but also in time – to verify whether a real $\Delta\nu$-$\Delta\dot{\nu}$ correlation is present for the glitch population as a whole.

It is more straight forward to examine the $\Delta\nu$-$\Delta\dot{\nu}$ relation on a case-by-case basis. We present an example of four individual pulsars with multiple glitches in figure 12. Glitches where $\Delta\dot{\nu}$ has not been measured or was found to be $\Delta\dot{\nu} \geq 0$ have been excluded. For the Crab pulsar (bottom left panel of figure 12), which is generally monitored very frequently and whose ToA errors are relatively small (see figure 2 in [63]), the correlation between glitch parameters seems genuine and is also apparent when one uses the persistent changes in spindown rate $\Delta\dot{\nu}_{\rm p}$ instead of the instantaneous ones. Using ten such measurements (calculated after the short-term exponential recovery was over), an approximately linear relationship can be fitted as

$$|\Delta\dot{\nu}_{\rm p}| = 7 \times 10^{-8}\ \Delta\nu\ \mathrm{Hz/s} \qquad (19)$$

where $\Delta\nu$ is measured in Hz [80].

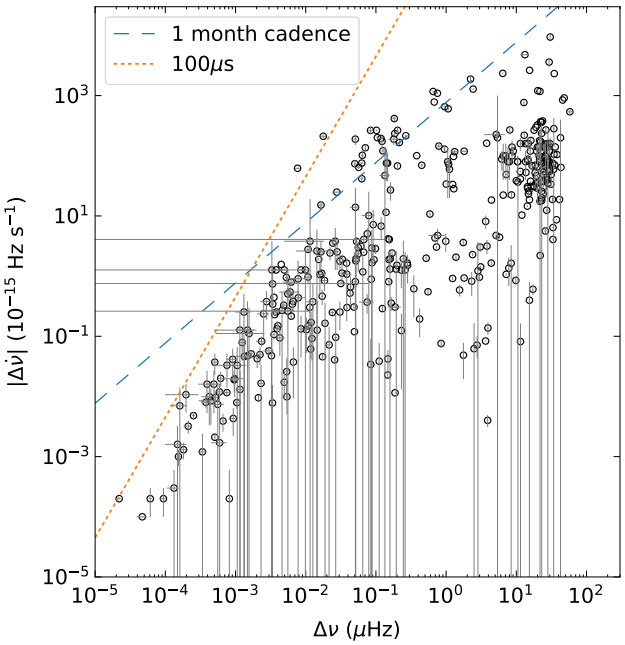

Figure 11. $|\Delta\dot{\nu}|$ versus $\Delta\nu$ in logarithmic scales for glitches where the former has been measured negative. There is an apparent positive correlation between the two parameters, to which contributions from observational biases should not be ignored. The indicative effect of two limiting factors is presented with the two straight lines (following [63]). The long-dashed line represents the repercussion of infrequent observations. When the ToA cadence is once per month or lower, glitches with parameters in the space above that line might not be picked up as such. The dotted line, which becomes important for smaller events, places an upper limit to the glitch parameters that can be probed when either the timing residuals RMS or ToAs' uncertainties are on average 100 μs or greater. See text for a detailed discussion.

In general, glitch parameters $\Delta\nu$ and $\Delta\dot{\nu}$ are calculated by extrapolating the pre- and post-glitch timing solutions to the inferred glitch epoch $t_{\rm g}$, since it is rare for a glitch to occur during observations. Measured parameters can thus deviate from their true value; whether they will be under- or over-estimates depends on several factors like the offset of $t_{\rm g}$ from the exact glitch moment and the accuracy of fitted timing models. Often though, especially when the true epoch is well constrained and the glitch model includes only steps in ν and $\dot{\nu}$, glitch parameters $\Delta\nu$ and $\Delta\dot{\nu}$ might be regarded as lower limits for their true counterparts. This is because, with current observing setups, sampling of ToAs is not frequent enough to allow the detection of fast-decaying changes most of the time, for most pulsars. Such changes have, however, been seen in some instances [153, 109], with timescales from minutes to few days, and might be a widespread feature of glitches. Nonetheless, the inferred amplitudes $\Delta\nu$ and $\Delta\dot{\nu}$ place constraints on the moment of inertia of the superfluid components

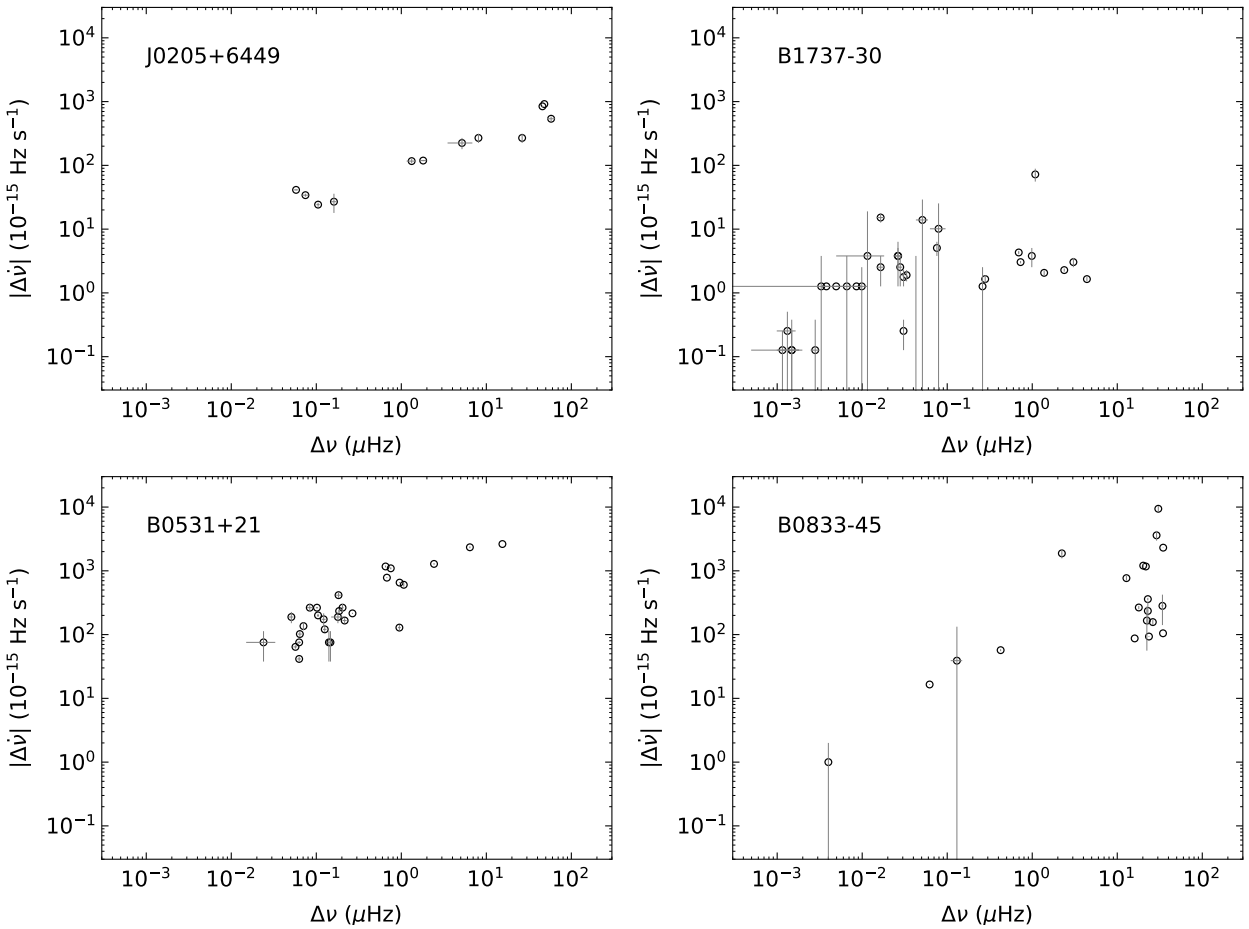

Figure 12. $|\Delta\dot{\nu}|$ versus $\Delta\nu$ in logarithmic scales, for those glitches where the former has been measured and $\Delta\dot{\nu} < 0$, in four selected pulsars including the Crab (bottom left) and Vela (bottom right) pulsars.

that contribute to the glitch, as well as the coupling timescales of the superfluid, and can be ultimately used to study the structure of neutron stars. One of the earlier examples of this comes from application of the vortex creep model to Vela glitches [166], leading to deductions for its mass and equation of state [167].

A first insight comes from the changes $\Delta\dot{\nu}$ in the spin-down rate. It can be assumed that in pulsars without observed emission changes close to the glitch, the external torque $N_{\rm sd}$ did not vary on short timescales, except for its dependence on the spin frequency. For example, in the standard form for a rotating magnetic dipole in vacuum:

$$N_{\rm sd} = -\frac{B_\perp^2 R_\star^6}{6c^3}\Omega_{\rm c}^3, \tag{20}$$

where $R_\star$ the stellar radius, $B_\perp = B_{\rm p}\sin\alpha_B$ with $B_{\rm p}$ the polar dipole magnetic field component, α_B the inclination angle, i.e. its angle to the rotational axis, and $\Omega_{\rm c} = 2\pi\nu$. Then, with the additional assumption that there was no change in the actual moment of inertia of the star at the moment of a glitch – which might not be the case if, for example, a crustquake was involved in the process – we can obtain the limit

$$\frac{I_{\rm d}}{I_{\rm c}} \geq \frac{\dot{\nu}^{\rm post} - \dot{\nu}^{\rm pre}}{\dot{\nu}^{\rm post}}\,, \tag{21}$$

where $\dot{\nu}_{\rm c}^{\rm pre}$ is the spin-down rate immediately before the glitch and $\dot{\nu}_{\rm c}^{\rm post}$ is after the glitch and at time $t_{\rm post}$ so that $\nu^{\rm post} < \nu^{\rm pre}$. Whilst $I_{\rm c}$ is the moment of inertia that was strongly coupled before the glitch (comprising both the normal component and any superfluid regions that were spinning down with it), $I_{\rm d}$ represents the moment of inertia that decoupled due to the lag change at the glitch and has not yet had time to recouple by time $t_{\rm post}$. Typically, $\Delta\dot{\nu}/\dot{\nu} \approx I_{\rm d}/I_{\rm c} \sim 10^{-3}$; only a small fraction of the interior superfluid.

Another constraint arises when considering the changes $\Delta\nu$ in the spin frequency, which serves as a proxy of the angular momentum imparted from the glitch-driving superfluid component to the crust. Limits on the physical quantities involved are by necessity model-dependent; nonetheless useful information can be gained and demonstrate the

huge potential glitch studies have in progressing our understanding of dense matter physics.

The starting point is angular momentum conservation (for simplicity we will first assume distinct, rigidly rotating, components but the same arguments hold when this assumption is relaxed). If the glitch is driven by a superfluid region of moment of inertia $I_{\rm gl}$ which has spun down on average by $\delta\Omega^{\rm s}_{\rm gl}$ due to vortex unpinning, then angular momentum $I_{\rm gl}\delta\Omega^{\rm s}_{\rm gl}$ has been transferred to the crust and any other components tightly coupled to it. We denote the moment of inertia of the stellar component that accelerates at the glitch as I_α. The extent of this part of the star that initially spins up depends on the interplay of the various coupling times and the glitch rise timescale. A larger fraction of the star catches up with the spin-up if this happens gradually, and vice-versa. By angular momentum conservation, $I_\alpha\Delta\Omega_{\rm in} = I_{\rm gl}|\delta\Omega^{\rm s}_{\rm gl}|$, where $\Delta\Omega_{\rm in}$ is the initial spin-up magnitude. We stress again that observationally inferred 'instantaneous' $\Delta\nu$ at the glitch epoch can be an underestimate of the actual change of the crust's spin frequency in the presence of short-lived transients. In other words, when replacing $\Delta\Omega_{\rm in} = 2\pi\Delta\nu$ we take I_α to represent all regions that couple on timescales shorter than the observational resolution of each glitch. The observed spin-up size is then

$$\Delta\nu = \frac{I_{\rm gl}}{I_\alpha}\frac{|\delta\Omega^{\rm s}_{\rm gl}|}{2\pi}\ . \qquad (22)$$

The reduction in superfluid velocity $\delta\Omega^{\rm s}_{\rm gl}$ is unknown. Under the assumption though that the lag might vanish at a glitch but does not reverse in the region $I_{\rm gl}$, $\delta\Omega^{\rm s}_{\rm gl}$ is limited to a maximum by the critical lag $\omega_{\rm cr}$ for unpinning.

Several approaches have been followed to translate (22) to information about the stellar structure and the microphysical parameters that enter the glitch problem. One such method, which relies mostly on observational data rather than theoretical estimates of $\omega_{\rm cr}$, is using a parameter called the glitch activity [152]. The activity parameter $\mathcal{A}$ can be related to the rate of angular momentum transfer to the observed component due to glitches, $\dot{L}_{\rm gl}$, and is calculated by considering the cumulative spin-up of the crust over the years of observations $T_{\rm obs}$:

$$\mathcal{A} = \frac{1}{T_{\rm obs}}\sum_i \Delta\nu_i \qquad (23)$$

with the index i accounting for all glitches the pulsar displayed during $T_{\rm obs}$. For pulsars with spindown rate $10^{-14}\ {\rm Hz\,s^{-1}} < |\dot{\nu}| < 10^{-10.5}\ {\rm Hz\,s^{-1}}$, glitch activity increases with increasing $|\dot{\nu}|$, with an average ratio $\mathcal{A}/|\dot{\nu}| \sim 1.2\%$–$1.8\%$. This is the fraction of spindown that has been 'reversed' because of glitches in a given pulsar [153, 27, 157, 154]. There is an observational bias at play that could explain why pulsars with smaller spindown rates fail to follow this trend: the glitch rate falls quickly with decreasing $|\dot{\nu}|$ and so finite observational spans $T_{\rm obs}$ could underestimate the activity in this category. As evident from (23), the activity of a pulsar is dominated by its largest glitches: for the calculation of $\mathcal{A}$, a long total observing time span is more important than the sensitivity of the observations to small glitch detection. On the other extreme, the Crab pulsar and B0540-69 have among the highest $|\dot{\nu}|$ but display low activity compared to expectations of the above scaling.

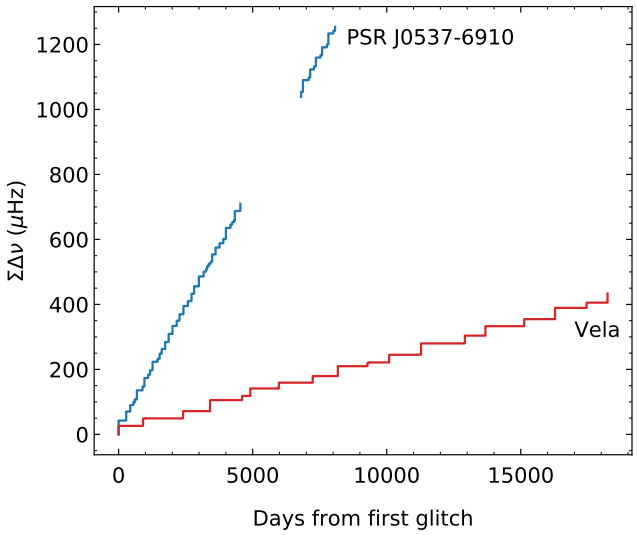

Figure 13. The accumulated glitch spin-up over 52 years of observations of the Vela pulsar (red); and over 21 years of observations of PSR J0537−6910 (blue). For the latter there is a period of 5.7 years in which no observations are available. The regularity of the glitches in these neutron stars allows an average glitch activity $\mathcal{A}$ to be approximated by the slope of the curves.

By definition, $\mathcal{A}$ is sensitive to the period over which there are observations geared to find glitches, and will not be constant if calculated over different periods of time for pulsars with irregular time intervals between glitches and diverse glitch sizes $\Delta\nu$. There are, however, some pulsars with regular glitch activity: their glitches have a characteristic size and interglitch waiting time (details on these distributions follow in Section 4.3). For the sources where these conditions are met, and provided there are many glitches detected, an average $\mathcal{A}$ can be well defined and calculated with reasonable accuracy; it is the slope of a linear fit in a graph as in figure 13. Owing to their special characteristics, the same few neutron stars can be used to relate the observed activity to a constraint for the internal parameters in equation (22), which in turn can be used to infer their bulk properties such as their mass given an equation of state as discussed below.

In the pulsars with regular activity, glitches not only have a characteristic (typically large) size and recurring time, they are also rather similar in their recovering behaviour. It is thus a reasonable assumption that both I_α and the product $I_{\rm gl}\delta\Omega^{\rm s}_{\rm gl}$ do not vary much from glitch to glitch. An average rate of angular momentum transfer to the component I_α will then be $\dot{L}_{\rm gl} = I_\alpha 2\pi\mathcal{A}$. Moreover, the fact that the waiting times are alike means that the lags building up during the interglitch intervals are also comparable, as the average spindown rate does not change much on such timescales. Assuming therefore that glitches are always triggered close to a specific value of the lag, e.g. when $\omega_{\rm cr}$ is reached, the reductions $\delta\Omega_{\rm gl}$ (and thus $I_{\rm gl}$) must also be varying little between glitches. The rate that the component $I_{\rm gl}$ stores angular momentum is limited by $I_{\rm gl}2\pi\dot{\nu}$ and so $\dot{L}_{\rm gl} \lesssim I_{\rm gl}2\pi\langle|\dot{\nu}|\rangle$, where $\langle|\dot{\nu}|\rangle$ is the average spindown rate in the course of observations $T_{\rm obs}$. It follows that

$$\frac{I_{\rm gl}}{I_\alpha} \gtrsim \frac{\mathcal{A}}{\langle|\dot{\nu}|\rangle} \ . \tag{24}$$

For the Vela pulsar $I_{\rm gl}/I_\alpha$ turns out to be close to 1.5% [168], and similar numbers close to 1% and up to 2% are obtained for other pulsars which present the regularity that validates this approach, and enough glitches to reliably estimate $\mathcal{A}$ [169]. Most of these pulsars have a characteristic age similar to Vela, around 10^4 yrs, except PSR J0537−6910 which is somewhat younger ($\sim 5\times10^3$ yrs) and has a relatively small inferred ratio $I_{\rm gl}/I_\alpha \sim 9\times10^{-3}$ [100].

The fact that the moments of inertia ratio $I_{\rm gl}/I_\alpha$ is about 1% lent support to the idea that glitches originate in the inner crust of the neutron star which, if I_α is close to the total moment of inertia $I_{\rm tot}$, indeed would amount to a similar percent. As reviewed in Section 2.3, vortices are expected to pin to the ion lattice thus the crustal superfluid can remain decoupled on long timescales. It is also possible for the core superfluid to be strongly coupled to the normal component (see Section 5.1 for the microphysical mechanism), validating the assumption $I_\alpha \simeq I_{\rm tot}$. This argument does not, however, account for the entrainment effect in the inner crust which can have an important impact on the dynamics. The unbound neutrons that form the inner crust superfluid might not experience a frictional force, but nonetheless they are not absolutely free to flow as they are Bragg reflected by the ions of the crustal lattice [170]. This non-dissipative mechanism essentially reduces the angular momentum reservoir available for glitches, as it couples the two components so that $\dot{\Omega}_{\rm s}$ is never zero, even in the idealised scenario of perfect vortex pinning. The angular momentum of the superfluid depends not only on $\Omega_{\rm s}$ but also on $\Omega_{\rm c}$ (see equations (7) and (10)). Using the band theory of solids, it has been shown that entrainment can be very strong in the inner crust, especially at baryon densities around $0.02-0.03$ fm^{-3} [171]. This will considerably limit the effectively decoupled moment of inertia of the crust that can contribute to $I_{\rm gl}$ [172]. The right-hand-side of equation (24) is increased by a factor $(1-\varepsilon_{\rm n})$.

The upper end of theoretical expectations for entrainment [171] implies a moments of inertia ratio as high as ∼10%, uncomfortably large to identify with the crustal moment of inertia for most equations of state and a canonical neutron star mass around 1.5 solar masses [172, 173]. Confining $I_{\rm gl}$ to the inner crust superfluid would imply a thick crust for the pulsars under consideration, which in general, for a given equation of state, translates to small stellar masses, sometimes unrealistically so, or to particularly stiff equations of state (resulting in larger radii). At the moment there are multiple unknown factors entering the problem, which leaves several conceivable solutions to this tension (see for example [174] for an exploration of the different effects using a Skyrme model). Each route can offer valuable information to be reaped when progress narrows down the possibilities. Given the uncertainties in the equation of state and the sensitivity of the crustal moment of inertia to the transition pressure from the core to the crust, strong entrainment could still be consistent with crust-driven glitches [175, 176]. Different calculations of the entrainment coefficient produce a more modest effect [177, 178], which is easier to accommodate.

A third possibility is that a fraction of the core is not tightly coupled to the crust, therefore I_α is smaller than what is typically assumed. The fast coupling mechanism between the core superfluid and the normal component is attributed to the electron scattering off magnetised vortex cores (details in Section 5.1) [179]. The nature of the protons (normal, type-I or type-II superconducting), of vorticity in the $^3P_2$ channel [180], the existence of more exotic condensates in the inner core of the star, or even a magnetically-decoupled region because of the magnetic field's topology [181] are all factors that affect the exact profile of coupling timescales in the core, and thus I_α. Information on this can also come from observations of the glitch recovery phase ([182], Section 5.4).

Last, part of the core might be so weakly coupled that it contributes to the glitch reservoir: $I_{\rm gl}$ can involve both the crust and core, or even be located in the core alone. This requires a mechanism for vortex pinning in the core, with the interaction of type-II proton fluxtubes and vortices being the natural candidate as discussed in Section 2.3. Strong pinning throughout the core raises the question why glitch $\Delta\nu$ is not, at least occasionally, much larger than what we

observe. A possible answer lies in the dependence of the pinning efficiency to the magnetic field topology; geometrical arguments indicate that the toroidal field region is the best candidate for strong pinning to occur [41]. This region has a spatially confined extent in the core that varies with the total (interior) magnetic field strength. Early results suggest that the mixed toroidal-poloidal field region might reside only in the crust for internal field intensities below $\sim 10^{13}$ G [183, 184]. Even if pinning to fluxtubes does not result in the core region to act as angular momentum reservoir, it can still raise the coupling timescales locally, reducing I_α and contributing to the observed recoveries [185]. In fact, analysis of a large Vela glitch which was observed in 2016, indicates that both last factors are at play (I_α being lower and $I_{\rm gl}$ having contributions from the core) [186]. We discuss this point further in Section 5.

Considering a scenario where glitches are driven by the singlet-state superfluid alone, a varying fraction of the outer core would supplement the crustal superfluid depending on the pairing gap (in some theoretical profiles, $^1S_0$ pairing occurs only at inner crust densities, in others it extends to the outer core). The exact superfluid moment of inertia in this state will also depend on the internal temperature of the star with respect to the critical temperature for superfluidity. A range of pairing profiles and equations of state, combined with glitch activity measurements of nine pulsars for which internal temperature estimates can be made (either from surface temperature measurements, when available, or simulations of minimal cooling in combination with information on the pulsar's age) was explored in [169], who found that, in the case of the pairing profile of [187], the singlet-state superfluid can provide the required moment of inertia $I_{\rm gl}$ in (24) to explain glitch observations given the entrainment parameters in [171]. On the basis of this idea, glitches can be a powerful test of nuclear physics predictions: by adopting an equation of state, the results can be translated into a mass estimate for these pulsars, information otherwise inaccessible for individual neutron stars [169].

Another way to calculate the extent of the angular momentum reservoir requires a density profile for the critical lag $\omega_{\rm cr}$ for unpinning, which limits $\delta\Omega^s_{\rm gl}$ in (22) and thus returns a maximum glitch size for a given moments of inertia ratio $I_{\rm gl}/I_\alpha$. A worked-out example of this has been presented in the context of the "snowplow" glitch model [188] using the mesoscopic pinning force calculations of [45] in order to obtain estimates of upper bounds for glitching pulsar masses [189, 190]. In this model, vortices pin only to the crust but extend through the core whose moment of inertia contributes to glitches – the hypothesis being that vortices can continuously thread the singlet and triplet superfluid state, and no layer of normal neutron matter exists beyond neutron drip densities. The maximum of the pinning force per vortex line is achieved in the inner crust, for vortices fully immersed in it. It is assumed that weakly-pinned vortices further in unpin, move out, and re-pin in the strong pinning region (hence the name, since vortices pile-up there as if pushed by a snowplow). Once the Magnus force exceeds even the maximum of the pinning force, all the accumulated vorticity is released, leading to a glitch [188]. An idealised scenario that maximises glitch impact is that all regions reach their critical lag and then the reservoir is fully depleted at the glitch, i.e. the average lag goes to zero. This allows one to obtain a strict upper limit on the amount of angular momentum that can be exchanged during a glitch which, once the pinning model and equation of state are fixed, depends only on the mass of the star. By comparing to data of the largest glitch measured in a pulsar, one can thus obtain a (model-dependent) upper limit on the mass of the star. Any larger glitches observed in the future by the same pulsar would lower this upper limit, as the angular momentum available in the reservoir is inversely proportional to the mass of the star. A note of caution is due, as currently published limits based on this method rest on the assumption that published glitch sizes $\Delta\nu$ always underestimate the true glitch size; this is however not universally reliable. Reported $\Delta\nu$ are calculated by extrapolating the post-glitch timing solution back to the (inferred) glitch epoch: depending on the model fitted to the recovery of each glitch in question and its accuracy, it is even possible to overestimate the true glitch size (leading, incorrectly, to a lower inferred maximum mass).

4.3. Glitch sizes and interglitch waiting times of individual pulsars

Besides population-wide studies, it is constructive to examine the glitch properties of each neutron star separately: as already apparent in figures 8 and 12, the distribution of glitch sizes and interglitch time intervals in individual pulsars varies. Analysis of these distributions for most glitching pulsars is challenging because of restricted sample size, nonetheless some first conclusions can be drawn for the sources with $\gtrsim 10$ recorded glitches.

First, the Vela pulsar (top panel in figure 14) and PSR J0537−6910 present a relatively narrow distribution in glitch amplitudes $\Delta\nu$, dominated by large glitches (peak around $15-20\mu$Hz), that can be modelled as a Gaussian [100, 191, 92]. Their waiting times distribution, that is, the time interval between successive glitches, also shows some normality. Other potential candidates for this category do exist, for example pulsars J1420-6048, J1801-2451, and J1803-

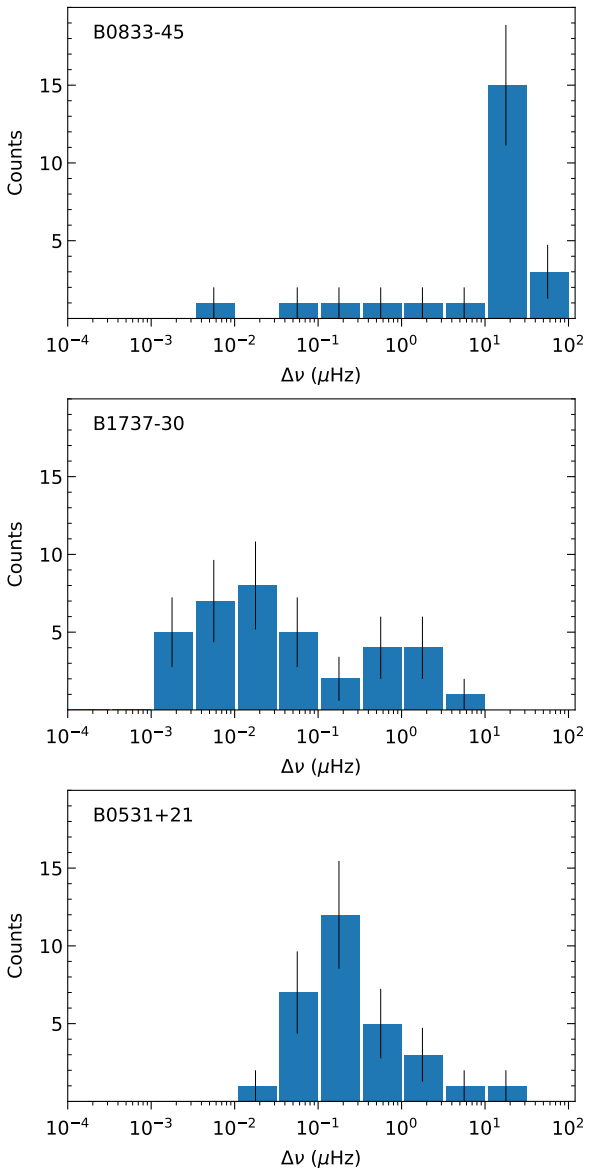

Figure 14. Histograms of glitch sizes $\Delta\nu$ presented in a logarithmic scale for the same three pulsars as in figure 8. Top panel: Glitch sizes of the Vela pulsar are best described by a Gaussian $\Delta\nu$ distribution which peaks at large, $\Delta\nu \sim 10$ μHz, amplitudes. Middle panel: PSR B1737−30 (J1740−3015) has a broad $\Delta\nu$ distribution, consistent with a power-law. Bottom panel: the Crab pulsar size distribution can be modelled as a Log-normal distribution or a power-law.

2137, but more glitches need to be observed from these pulsars before a statistically sound conclusion about their glitching nature can be reached.

On the other hand, the rest of pulsars with sufficiently "large" glitch sample have size distributions that are generally best-fitted with a log-normal distribution or a power-law. A prominent example is the Crab pulsar (bottom panel in figure 14), whose size distribution is best-fitted with a log-normal [63, 191], but can also be described with a bounded power-law, with index ~ 1.36 [63, 192]. The probability of glitches with sizes below $\sim 10^{-2}$ μHz seems to rapidly decrease, but not purely as a result of observational bias as the Crab pulsar is regularly monitored and glitch-like irregularities of smaller amplitudes can be, and have been, detected [63]. Other examples of pulsars with non-Gaussian size distributions are PSR J1341−6220 (possibly log-normal, although see below), J0631+1036 and B1737−30 (power-laws, middle panel in figure 14) [191, 68]. The waiting times distribution of pulsars in this category can be well represented by an exponential distribution [67, 68].

It is worth noting here that whilst the $\Delta\nu$ distributions must be complete above a certain detection threshold, and thus can be interpreted as long as this limitation on its lower end is taken into account, the waiting times distribution is much more sensitive to the completeness of the sample. Possibly undetected small glitches between larger, recorded, events can considerably change the inferred waiting times. This effect is, fortunately, unlikely to be greatly important for pulsars that are monitored sufficiently often.

Some pulsars with multiple glitches do not seem to fit well in either category. PSR J1341-6220 is potentially such a source, where although waiting times are broadly consistent with an exponential distribution, both sizes and waiting times distributions show a weak peak and there is a moderate excess of large glitches (compared to power-law expectations). In all such pulsars their glitching nature might become clear when more glitches are detected, but it cannot be yet excluded that a mixture model might be more appropriate. This could be perhaps the case if there are two mechanisms at play, giving rise to glitches of different properties: a population with a characteristic scale and a population with a broader range of (usually smaller) sizes. The scenario of two glitch populations has also been investigated for its possible effect on the inferred correlation between glitch parameters $(\Delta\nu, \Delta\dot{\nu})$ and forward or backwards waiting times [100, 191] which is considered below.

Both the size and waiting time distributions provide input and constraints on proposed glitch models. One aspect of the mechanism particularly illuminated by these analyses is the glitch trigger. The conjecture that glitches are driven by a build-up – deplete – repeat process in a fixed region that acts as the angular momentum reservoir until the same critical lag is always reached is not in accordance with the wide size and waiting times distributions observed.

In general, glitch sizes do not correlate with the time since the preceding event (with which stress will scale, e.g. the lag ω that is built between glitches) as could have been expected if glitches

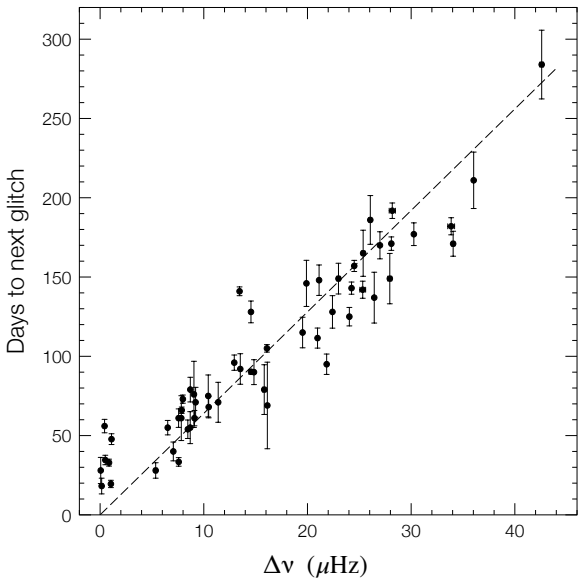

Figure 15. The time interval until the following glitch versus the size $\Delta\nu$ of the preceding glitch, for 56 glitches in PSR J0537−6910 as reported in [100, 101, 194]. The dashed line is a linear regression model forced to pass by the origin.

depleted the entire reservoir of angular momentum (and assuming the same moment of inertia of superfluid takes part in each event). By contrast, a strong positive correlation between glitch size $\Delta\nu$ and the time until the next glitch ΔT_f has been established for the pulsar J0537−6910 [193, 100], as manifested in figure 15. This correlation favours a threshold-regulated process as the glitch trigger – it takes longer to reach the threshold after a larger $\Delta\nu$ glitch – whilst glitches do not necessarily deplete the available reservoir. Non depletion means glitch sizes can be larger than expected from the lag built since the last event (as they can draw from residual lag), or smaller if they release only some of the lag accumulated. This is in accordance with observations of J0537−6910 ([100], interpreted by a simple superfluid model) and explains the lack of correlation between waiting times and size of following glitch, but leaves the question why the $\Delta\nu$–ΔT_f correlation is not seen in other pulsars. A similar, albeit weaker, trend is observed in few more pulsars [191] but not even the Vela, despite being the only other source with sizes and times distributions best fitted by Gaussians, presents a clear correlation.

The possibility that the lack of observed $\Delta\nu$–ΔT_f correlation is because of undetected glitches is disfavoured for the Crab pulsar, whose frequent monitoring allows discovery of very small events [63]. For even smaller glitches, the $\Delta\nu$–ΔT_f correlation would imply respectively small following time intervals. Missing those events will lead to only slight overestimation of the interglitch intervals, insufficient to completely obscure the inference of the underlying correlation (the same argument most likely holds for the Vela pulsar as well). Simulations of glitches with power-law distributed sizes $\Delta\nu$ show that the incompleteness of the glitch sample below a certain cut-off worsens the inferred correlation, but the effect is important only when the power-law is steep [191].

Another possibility which we mentioned earlier is that there are two populations of glitches: one consisting mostly of large events that are threshold-driven (and therefore will have correlated $\Delta\nu$–ΔT_f), and a set of smaller events occurring stochastically. This idea has been explored for the 8 pulsars with the largest glitch samples as of 2018 (including J0537−6910); results for most pulsars were mixed or inconclusive. For the Vela pulsar the correlation improves when glitches below $\Delta\nu \sim 2\mu$Hz are excluded [191], but remains weak compared to that of J0537−6910.

Another correlation has also been observed in J0537−6910, between the interglitch waiting time and the change $|\Delta\dot{\nu}|$ of the following glitch [193]. Later studies, using a nearly twofold sample of glitches, found less strong evidence for this correlation but confirmed the trend when the smallest glitches (5 out of 45) were excluded from the analysis [100]. Such a correlation could relate to the superfluid that is available to decouple (causing $\Delta\dot{\nu} < 0$) at each glitch. Longer waiting times will allow more of the decoupled superfluid to reach its steady-state before the next event is triggered, permitting for larger $|\Delta\dot{\nu}|$ at the next event. It might also hint at a different superfluid response depending on glitch size, something that can be explained by the non-linear dependence of the mutual friction on the lag ω [195]. This is covered along other data correlations and trends relating to the post-glitch recoveries in Section 5.

4.4. Glitch triggers and meta-models

The statistics presented above, while far from exhaustive given the still low number of glitches observed from many pulsars, can however be used to try to obtain constraints on proposed glitch models and in particular the nature of the trigger mechanism. The observation that, for several pulsars, the size and waiting time distributions lack a characteristic scale is suggestive of a scale-invariant process at work [67]. Crustquakes could be such a process if there is an analogy to several geological phenomena, but scale invariance can also arise from the other promising glitch trigger mechanism: vortex unpinning and knock-on effects which lead to vortex avalanches. In fact, simulations of quantum vortices in a spinning down container (obtained by solving the Gross-Pitaevskii equations) have shown that avalanches take place

and indeed result in powerlaw distributions for glitch sizes and exponential ones for waiting times [65, 66]. Nevertheless, scaling up these simulations, which typically model $\approx 10^3$ vortices, to a realistic setup for pulsars by including $\sim 10^{13}$ vortices is technically unfeasible. One has to rely on large scale models, that include some aspects of the microphysics, while allowing for local stresses to be released and/or accumulated in a stochastic manner [196, 197].

Testing each individual model against observations and understanding the impact of its underlying set of assumptions is prohibitive, however, so-called meta-models can be introduced. These are models that incorporate the features of a number of popular glitch models, yet remain agnostic about the details of the microphysics. Two such meta-models have been developed [198, 199]. One considers mechanisms that can trigger a glitch stochastically as a critical value is approached (such as in the vortex avalanche model, in which avalanches are more likely to occur if the system is close to $\omega_{\rm cr}$). The other meta-model considers mechanisms in which the glitch is triggered at a given threshold, e.g. when $\omega_{\rm cr}$ is reached (or the elastic limit of the crust exceeded, for crustquake models), but the system approaches this threshold in a 'noisy' manner. The meta-models allow to predict a number of quantities of physical interest, which can be compared to current and future observations. Of main interest is whether a model can explain the roughly two classes of probability density functions (PDFs) for the glitch size and waiting time distributions seen in individual pulsars: the first with exponentially distributed waiting times, and monotonically decreasing, scale-invariant (possibly power-law distributed in some cases) size distributions, and the second with unimodal distributions for both sizes and waiting times, i.e. Vela-like behaviour. Besides PDFs, the phenomenological meta-models can be used to predict waiting time and size cross-correlations, and also auto-correlations, making concrete and falsifiable predictions about long-term glitch statistics.

The first meta-model assumes a state-dependent Poisson process [198]. Stress accumulates in a deterministic fashion (driven by the external spin-down torque) in-between glitches, but glitches are triggered stochastically via a Poisson process whose rate increases with the stress. The instantaneous glitch rate $\lambda(t)$ is thus a function of the stress, quantified by a single variable $X(t)$, which in the case of superfluid-driven glitches is taken to be the spatially averaged lag between the superfluid component and rigid crust. A phenomenological prescription was adopted where

$$\lambda(X) = \lambda_0 (1 - X/X_{\rm cr})^{-1}\,, \tag{25}$$

with λ_0 a reference rate, and $X_{\rm cr}$ the maximum lag that the pinning force can support, but it was found that the specific functional form for the rate has little impact on the long-term statistics as long as it encodes two main physical expectations of glitch models: (1) $\lambda(t)$ increases monotonically with X and (2) $\lambda(t)$ diverges at $X = X_{\rm cr}$ (i.e. glitches become increasingly more likely as the critical lag at which all vortices unpin is approached) [198]. The lag evolves according to:

$$X(t) = X(0) + \frac{N_{\rm sd} t}{I_{\rm c}} - \sum_{i=1}^{N(t)} \Delta X^{(i)}\,, \tag{26}$$

where $X(0)$ is an arbitrary initial condition, $N_{\rm sd}$ is the spin-down torque, $I_{\rm c}$ the moment of inertia of the normal component, $N(t)$ the number of glitches that have occurred up to time t, and $\Delta X^{(i)}$ is the change in lag at the i-th glitch. The change $\Delta X^{(i)}$ is governed statistically by a conditional distribution $\eta(\Delta X|X_p)$, which depends on the lag immediately before the glitch X_p. The model remains agnostic of the physics governing η, and can describe a number of prescriptions. For example Gross-Pitaevskii simulations suggest a power-law for η if vortex avalanches are the dominant glitch trigger mechanism (but the picture is similar for crust quakes) [65, 66]. A number of alternative functional forms, in particular uni-modal distributions that predict typical scales for the glitch, as expected e.g. by models built around the snowplow paradigm [188], have also been explored [200]. Whilst η encodes the microphysical distribution of changes in the lag in the superfluid of moment of inertia $I_{\rm gl}$, which is not directly observable, it can be connected to the long term distribution of glitch sizes $\Delta\nu$ by demanding angular momentum conservation at each event

$$\Delta X^i = -\frac{2\pi(I_{\rm c} + I_{\rm gl})\Delta\nu^{(i)}}{I_{\rm gl}}\,, \tag{27}$$

and evolving X through a simple Monte Carlo automaton as follows [198]. A random waiting time Δt given the current lag is extracted from the distribution:

$$p[\Delta t|X(t)] = \lambda\left[X(t) + \frac{N_{\rm sd}\Delta t}{I_{\rm c}}\right] \times \exp\left\{-\int_t^{t+\Delta t} dt^{'}\lambda[X(t^{'})]\right\}\,. \tag{28}$$

The lag is updated to account for the spin-down during the time Δt and a random size ΔX is picked from η, and subtracted from the lag. The change in observable frequency is calculated, and the procedure repeated.

In general the model can be analysed in terms of two dimensionless variables that control the dynamics: $\tilde{X} = X/X_{\rm cr}$, the dimensionless stress, and $\alpha = (I_{\rm c}X_{\rm cr}\lambda_0)/N_{\rm sd}$ which encodes how rapidly the system is driven. For large values of α the system is slowly driven, and when powerlaw and unimodal distributions

are considered for η then the size distributions are approximately scale-invariant (although not exactly power-laws for unimodal η distributions) [198, 200]. Current observations do not yet allow us to distinguish between the different forms, but future data may make this possible, especially in combination with other constraints on cross and auto-correlations, as will be discussed in the following.

For small values of α, on the other hand, the system is rapidly driven and typically has time to climb close to the critical lag before a glitch is triggered. Small glitches will be followed by short ΔT_f before the next one is triggered, whilst larger events will statistically precede longer waiting times therefore, even though avalanches are uncorrelated, a strong correlation between size and waiting time until the next glitch will emerge. Note that whilst this matches observations of the rapidly spinning-down PSR J0537−6910, the other high $|\dot{\nu}|$ candidate for this behaviour is the Crab pulsar, which does not display a $\Delta\nu$-ΔT_f correlation. To overcome this challenge requires $X_{\rm cr}$ to vary between these pulsars. If η is unimodal, it also leads to unimodal distributions in the size and waiting time glitch distributions. Weak correlations are also predicted between size and waiting time from the previous glitch (ΔT_b), and auto-correlations in the size and waiting time distributions are also predicted but depend strongly on the physical input for η. As more data is collected, observations of such correlations and auto-correlations in the pulsar population may thus be used to constrain the physical input of the models [201]. The effect of a history-dependent rate η has also been studied [202], corresponding to the generally accepted picture of how pulsar glitches occur, i.e. vortices pin and unpin in an inhomogeneous pinning landscape, where the unpinning is governed by the evolution of the local thresholds of pinning sites as a function of time. Interestingly the model predicts aftershocks for large glitches, which are not observed (although given the scarcity of data, we may not have yet observed large, system resetting, glitches) and cross-correlations between sizes and waiting times, which are provisionally inconsistent with observations.

A different kind of studied meta-model is one where glitches are the result of a Brownian stress accumulation process [199]. In this case glitches are not triggered stochastically, but instead happen when a certain threshold is reached (as predicted e.g. for glitches triggered by fluid instabilities [69]). The evolution of the stress driving the system to that threshold is, though, stochastic following a Brownian process. These fluctuations in the stress might be connected to precursor events, such as the one possibly seen prior to the Vela 2016 glitch [110]. The stress evolves according to the Langevin equations:

$$\frac{dX(t)}{dt} = \xi + \sigma_{\rm d} B_{\rm wn}(t)\,, \tag{29}$$

where ξ is a drift coefficient (which models the external driving due to the spindown torque) and $\sigma_{\rm d}$ the diffusion coefficient, with $B_{\rm wn}(t)$ a white noise process with zero mean and unit variance [199]. A dimensional parameter can be defined as

$$\mu = \xi X_{\rm cr}/\sigma_{\rm d}^2\,, \tag{30}$$

with $X_{\rm cr}$ the critical lag at which the glitch is triggered, releasing an amount of stress ΔX which is once again drawn from a distribution η.

Contrary to the state-dependent Poisson process model, the Brownian model predicts no size and waiting time autocorrelations, and no cross correlation between size and ΔT_b. If any such quantity is determined to be non-zero for a pulsar, it would rule out this meta-model.

Thus, over time and with additional observations, we expect to be able to use sample properties like the distribution of sizes and waiting times, as well as their cross-correlations and auto-correlations, to falsify one or both meta-models, and to place constraints on the microphysical inputs, such as η. This would shed light on the physics acting in the star, and reveal whether the same mechanism is indeed acting in all glitching pulsars, or not [203].

Whilst this kind of meta-models are powerful tools to describe the long-term statistics of glitching pulsars, they do not fully encode the hydrodynamical coupling between the fluids in the star. For example glitches are treated as impulsive events with sizes determined from angular momentum conservation (giving essentially the asymptotic size of a glitch) and short term transients cannot be determined. A first attempt at addressing this is the development of a meta-model in which glitches are triggered stochastically but then the system is evolved following the multifluid equations [204]. Input avalanche sizes, as well as the strength of the mutual friction coupling by which they evolve, were drawn from underlying powerlaw distributions. The resulting glitch size distributions will differ significantly from power-laws, and present a lower size cutoff – such as the one observed in the Crab pulsar [63]. The cut-off appears because events in which only a small number of vortices unpin will appear in the data as gradual changes instead of displaying the standard glitch signature (see also [39] for an in-depth discussion on the effect of glitch size on the shape of the glitch recovery in this framework). Furthermore, when a realistic non-linear form is used for the mutual friction between the superfluid and the crust in the strong coupling regime (as will be described subsequently), it

can qualitatively change the response of the star to a glitch [195]. For small glitches, the transfer of angular momentum will be slow enough for the core to remain coupled. In large events, where vortex velocity will be high, the spin-up will be faster (see Section 5.1) and part of the core might decouple: I_α in equation (22) will be smaller, leading to an even larger observed $\Delta\nu$. This can create an excess of larger glitches and a bimodal size distribution, such as the one observed from the overall pulsar population.

5. The stages of a glitch

In this section we discuss the particulars of the individual stages of observed glitches and their theoretical implications. The first stage of a glitch is the spin-up phase. For the vast majority of glitches, this phase is not temporally resolved by observations and can be very fast: less than a minute for large Vela glitches. The next stage is the post-glitch response. Whilst some glitches appear only as steps in ν, without appreciable changes in $\dot{\nu}$ and $\ddot{\nu}$, most will be followed by a period of rotational evolution that differs from the pre-glitch one and can last for up to many years. This stage can present considerable complexity and remarkable variation from glitch to glitch, as can be seen in figure 6. Nonetheless, several phases can be recognised in the post-glitch response, the most common of which are illustrated in figure 18.

The observed timescales, as well as the functional form of the post-glitch response, have a high value as probes of the neutron star's internal state and dominant processes. On that account, we start by looking in some more detail at the underlying physics (Sections 5.1 and 5.2), before discussing insights gained by observations of the glitch rise (Section 5.3) and the rich phenomenology of post-glitch relaxation (Section 5.4).

5.1. Interactions between neutron stars' constituents

The great range of timescales involved in the glitch phenomenon can be understood in terms of the various microphysical mechanisms by which the constituents of neutron stars interact and couple to each other. In what follows we outline the most important ones together with order of magnitude estimates for the resulting timescales ‡.

The charged particles of the core (including the protons if they are not superconducting) respond to the spin-up of the crust either by (low frequency) hydromagnetic waves or the formation of an Ekman layer at the crust-core boundary. The timescales will be of the order of few to some tens of seconds for typical pulsar parameters ([205], see also [206] and [181] for a more detailed discussion and the effects of proton superconductivity and the magnetic field topology, as well as [179] for some considerations on the coupling of type-II superconducting protons). If neutrons are also normal, strong-interaction scattering between protons and neutrons brings them to rotational equilibrium almost instantaneously. If both protons and neutrons are superfluid, and without considering vortices, this coupling timescale becomes extremely large.

In general, the temperature T in the core of mature neutron stars will be well below the local critical temperature for superfluidity, so that not only neutrons are superfluid but also the contribution to coupling from scattering off neutron superfluid quasiparticles (excitations) becomes negligible. The dominant coupling mechanism is scattering involving the (normal) vortex cores.

The scattering of electrons from the thermal excitations in the normal vortex core (magnetic-dipole interaction) is relevant to both the inner crust and core of neutron stars. The relaxation timescale of this process strongly depends on the pairing gap, and thus can vary from very short to year-long with location within the star [207]. Whilst in the inner crust the superfluid will be in the singlet-state (neutron pairs have zero spin and orbital angular momentum), in the stellar core a $^3P_2$ phase appears. In this state, and in the case of non-superconducting protons, the vortex acquires a magnetic field of the order 10^{11} G as a result of the neutron magnetic moment. This provides additional coupling via the interaction of electrons with the vortex magnetic field; this coupling timescale between the vortex lattice and the electron 'fluid' scales as $\sim 1.3\times10^8 k_f \chi^{2/3}\Delta_{\rm n}^{-1}\nu^{-1}$ s where k_f the neutron Fermi wavenumber in fm^{-1}, the neutron pairing gap $\Delta_{\rm n}$ is in MeV, and χ is the electron fraction [208]. This timescale is relatively long ($\sim$ weeks) for typical pulsar parameters and becomes relevant only for very cold, roughly colder than 10^7 K, neutron star cores, for which (thermal) vortex core excitations will be negligible. Under the same assumption of normal protons and for realistic internal temperatures, the strong interaction between protons and vortex core excitations will provide the most efficient coupling, leading to very short timescales (typically under a pulsar period) [209].

In the case, however, of proton superconductivity in the core, the Andreev-Bashkin entrainment effect arises due to the long-range part of the strong interaction. It introduces a dissipationless coupling, in which neutrons "drag" along protons as they move (and vice-versa). This results in an induced current of protons around neutron vortices which sources a

‡ Exact formulae can be found in the referenced articles, see also [43] for a recent review.

much stronger (nearly 10^{15} G) magnetic field along the vortex axis – the exact magnetic flux carried by the neutron vortex will vary across the star as it depends on the entrainment parameters. The coupling timescale between vortices and electrons is then greatly reduced to the order of seconds for typical pulsar parameters [179], rendering the contribution of this mechanism to mutual friction the most important.

A magnetised neutron vortex will also interact with the proton vortices that carry the magnetic flux in a type-II superconducting region (fluxtubes). This interaction can potentially lead to pinning (Section 2.3) and decoupling of the core, at least partially, but once a large enough lag has developed for the Magnus force to exceed the pinning force, vortices will cut through fluxtubes – a highly dissipative process, leading to strong mutual friction [210]. In fact such form of mutual friction may also efficiently damp modes of oscillation of a neutron star, for example leading to an effective saturation amplitude for the r-modes [211] and explaining the lack of observed gravitational wave emission associated with such modes in the neutron stars in Low Mass X-ray Binaries [212, 213].

A transition to type I superconductivity is expected for densities exceeding 3×10^{14} g cm^{-3}, relevant for the inner core of neutron stars. The coexistence of a type-I superconductor with the neutron superfluid and vortices has not been studied yet in detail, mainly due to the unknown distribution, scale and shape of the non-superconducting regions in which the magnetic field will be confined. In the case in which the neutron vortices carry co-linear normal proton domains, however, the mutual friction coupling timescale can be calculated and is of the order of minutes [214].

In the inner crust vortices are immersed in the ion lattice. Dislocation of the crustal nuclei due to their interaction (pinning) with vortices can create a charge distribution around a vortex, allowing for Coulomb scattering of electrons, but early calculations indicate this mechanism will not dominate the coupling timescale [215] over the magnetic-dipole interaction. A moving vortex though will excite and couple to phonons in the lattice; the coupling will be relatively weak but is not highly sensitive to temperature, in contrast to processes relying on vortex core excitations, and can be in the correct range of tens of days observed in post-glitch recoveries [216]. The motion of a vortex through the pinning potential will also excite kelvons – quantised Kelvin modes on the vortex itself – which will interact with the lattice [217]. The resulting dissipation depends on the velocity of the vortex [217, 42]. In regions where vortices pin and a lag ω builds up, the vortex velocity can be high when it unpins (and is limited by the maximum $\omega_{\rm cr}$ for unpinning, see Section 2). Very strong coupling can then occur, with timescales of order of seconds, which can explain the fast rise of glitches as seen in the Vela pulsar [109] (see Section 5.3).

5.2. The superfluid response

Following the assumption of a linear drag force acting on vortices as in Section 2.2, the above microphysical estimates of the coupling timescales, for instance between the vortex lattice and the electron fluid or ion crystal, would enter the problem via the parameter $\mathcal{R}$ (equation (5)), or equivalently the dimensionless mutual friction parameter $\mathcal{B}$ (equation (9)).

In order to make connections to observations, these timescales should be converted to the dynamical timescales for the coupling between the macroscopic components of the star (which include the inertia of all constituents that relax to co-rotation). The short coupling timescales between the charged particles compared to the time resolution of observations and glitch recovery timescales, typically allows their treatment as a large-scale charge-neutral fluid. As discussed in Section 2, the most basic model of a superfluid neutron star is thus that of a system of two fluids: a neutron condensate with total moment of inertia $I_{\rm s}$, and the 'normal' fluid with moment of inertia $I_{\rm c}$, described by the equations (10). In these equations, the mutual friction coefficient $\mathcal{B}$ encodes the strength of the coupling and can be derived from the microphysical timescales (e.g. the electron-vortex interactions) $\tau_{\rm MF}$ as:

$$\mathcal{B} = \frac{1}{2\Omega_{\rm s}\tau_{\rm MF}} \,. \tag{31}$$

The dynamical timescale for the response of the superfluid is then given by equation (13), which we recall here:

$$\tau_c \approx \left(\frac{I_{\rm c}}{I_{\rm s}+I_{\rm c}}\right)\frac{1}{2\Omega_s\mathcal{B}} \,. \tag{32}$$

This expression neglects, for the sake of clarity, entrainment and any external torques, and $\Omega_{\rm s}$ is the initial angular velocity of the neutron fluid (see Section 2.2). We recall that the other assumptions going into equation (13) and the estimate of $\mathcal{B}$ were (i) that we can work in Newtonian gravity, (ii) that the flow is laminar, and (iii) that the response of the system is linear [218]. We now examine the consequences of these assumptions in a little more detail.

The first issue to consider is the effect of modelling the star in general relativity. The multifluid equations in Section 2 can be cast in a general relativistic form and have been studied by several authors (see [219] for a review). In particular it was shown early on that the vortex geometry can be affected significantly by the curvature of the spacetime [220], which will affect, among others, the glitch spin-up timescale.

Modelling the star as two rigid components with vortex pinning occurring throughout the star, the relativistic coupling timescale can differ from the Newtonian one significantly (by factors of up to 40%), depending on the equation of state of dense matter [221].

The analysis can be somewhat simplified if one considers pinning only in the crust of the neutron star. In this case many of the metric functions can be approximated in terms of their value in a thin outer shell. Following this procedure the difference between the relativistic coupling timescale, τ_{GR}, and the Newtonian estimate in (32), τ_c, can be expressed as a function of the stellar compactness only [222]. Explicitly, in units where $G = c = 1$:

$$\frac{\tau_{GR}}{\tau_c} = \left(1 - \frac{2M}{R_\star}\right)^{-1/2} \left(1 - \frac{2I}{R_\star^3}\right)^{-1} , \tag{33}$$

where M is the mass, $R_\star$ the radius and I the moment of inertia of the star. By using the universal relations of [223] to express I in terms of the relativistic compactness $C = M/R_\star$, one can obtain a relation which depends only on C and is, to a good approximation, independent of the equation of state. This can be seen in figure 16, from which it is clear that the relativistic correction can be sizable (even a factor 2 or more) for the most compact stars.

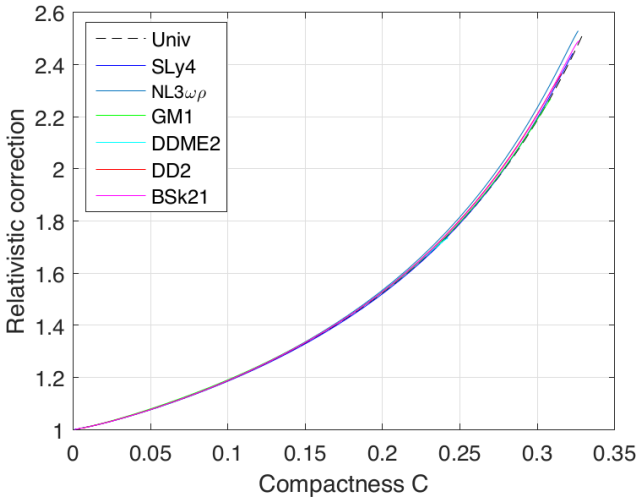

Figure 16. The relativistic factor τ_{GR}/τ_c as a function of compactness $C = M/R_\star$ of the star, plotted both for a number of realistic equations of state, as in [222], and using the universal relation obtained from (33). The relativistic estimate of the timescale can be significantly longer than the Newtonian one for the most compact stars.

Apart from the Newtonian approximation, the other main assumption that goes into (32) is that the equation for the response of the fluid is linear in the lag between the angular velocities of the components ω, i.e. that, assuming that all components are rotating around a common axis $\hat{z}$, the dissipative mutual friction force $F_{\rm MF}\hat{z}$ can be written in the form

$$F_{\rm MF} \propto \frac{\omega}{\tau_{\rm MF}} , \tag{34}$$

with $\tau_{\rm MF}$ the (constant) microphysical coupling timescale. This could be an acceptable approximation in some limiting cases, nevertheless it is not expected to hold in general (see e.g. [224, 217, 42, 195]). More broadly, one has that

$$F_{\rm MF} \propto f(\omega) , \tag{35}$$

where now f is a general function of ω.

The derivation of equation (13) in Section 2.2 assumes a constant number of free vortices (encoded in $\mathcal{B}$ via the free vortex fraction $\gamma_{\rm fv}$). Whilst for small deviations from the steady-state, $F_{\rm MF}$ could be expanded and approximated as linear in ω, this might not be the case for a pinned superfluid when large deviations from the steady-state can arise. Consider, for example, the steady-state case in which most regions of pinned vorticity are subcritical, i.e. are close to the critical threshold for the unpinning of vortices, so that any excess vorticity is expelled either by discreet avalanches [66] or vortex creep [224]. The effect of a glitch on these regions will be to move the system away from the threshold for unpinning, thus shutting down vortex motion. In the absence of free vortices the mutual friction vanishes, but as the normal fluid spins down, increasing the lag ω again and pushing it closer to the unpinning threshold, the fraction of free vortices $\gamma_{\rm fv}$ will start to increase.

In general $\gamma_{\rm fv}$ will depend non-linearly on the lag ω, and on the parameters of the system, such as density, pinning force and temperature. Simulations of vortex motion in pinning potentials show that $\gamma_{\rm fv}$ can be modelled as a sigmoid (an S-shaped curve) [50], and that for the particular case of a periodic pinning potential, can be approximated with an analytic formula:

$$\gamma_{\rm fv} \approx \left[(2+\sqrt{2})\left(\frac{\omega}{\omega_{\rm cr}} - \frac{1}{\sqrt{2}}\right)\right]^{1/2} , \tag{36}$$

where beyond $\omega_{\rm cr}$ (corresponding to the maximum of the pinning force) all vortices are free.

Similarly, in the thermal creep model, in which vortices unpin and 'creep' out due to thermal excitations, the fraction of free vortices sensitively depends on the lag ω as well as the temperature. This can be modelled in terms of a function

$$f(\omega) = \mathcal{C}\sinh(\omega/\tilde{\omega}_c) , \tag{37}$$

where $\mathcal{C}$ and $\tilde{\omega}_c$ are constants that depend on temperature, pinning energy, vortex density and the external spindown torque [224]. For $\omega/\tilde{\omega}_c \ll 1$ one has

$\sinh(\omega/\tilde{\omega}_c) \approx \omega/\tilde{\omega}_c$, and the above expression reduces to a linear model (although see [225] for a discussion of why the linear regime may not be realised in pulsars). If one is not in the linear limit, the solution to the full non-linear equations of motion (and observed response) is no longer an exponential, but rather takes the form of a Fermi function, so that the post-glitch spin-down rate of the crust evolves as:

$$\dot{\Omega}_{\rm c} = \frac{N_{\rm sd}}{I_{\rm c}} - \frac{I_{\rm s}}{I_{\rm c}}\frac{N_{\rm sd}}{I}\mathcal{F} \tag{38}$$

where $N_{\rm sd}$ is the external spin-down torque and $\mathcal{F}$ a function of time of the form:

$$\mathcal{F} = 1 - \frac{1}{1 + [\exp(t_o/\tau_{nl}) - 1]\exp(-t/\tau_{nl})} \tag{39}$$

with non-linear coupling timescales $t_o = I\Delta\Omega/|N_{\rm sd}|$ and $\tau_{nl} = (k_B T/E_p)(\omega_{cr} I/|N_{\rm sd}|)$, E_p the pinning energy, and $\Delta\Omega$ the perturbation (glitch) size.

More generally the mutual friction strength is expected to depend on vortex velocity, so that a realistic model which includes both phonon and kelvon modes for the response of the crust, will naturally be non linear [195] and not give rise to a clearly exponential relaxation. Furthermore, in the presence of strong pinning, vortex accumulation regions can form and the simple approximation for the vortex density $n_v \approx 2\Omega_{\rm s}$ will be no longer valid (Section 2). In this case one has $n_v = n_v(\Omega_{\rm s}, \partial\Omega_{\rm s}/\partial\varpi)$ with ϖ the cylindrical radius (assuming, for now, straight vortices and a laminar flow), and the equations of motion take the form of a Burgers equation [226] with propagating solutions, i.e. unpinned vortex 'fronts' which can propagate as avalanches, and also give rise to glitch precursors (see also Section 5.3).

The final assumption in equation (32) is that the flow is laminar, and the vortices straight. In a pulsar, however, both standard hydrodynamical and superfluid quantum turbulence may be excited [69, 71, 70]. The vortex array will be disrupted, leading to a polarized tangle and a different functional form of the mutual friction. In this sense, turbulence is a particular form of non-linear lag dependence of the response. In the case $\omega \ll \Omega_{\rm s}$, which is generally valid in neutron star interiors, the mutual friction may be decomposed, formally, as a sum of terms due to straight vortices, which carry the circulation, and a part due to vortex rings [73]:

$$F_{\rm MF} = L_s \langle f \rangle^s + L_r \langle f \rangle^r \ , \tag{40}$$

where L_s is the length of straight vortices in a fluid element, L_r the length of vortex rings, and the brackets indicate an average of the local force per unit length f taken over the straight vortices (superscript s) and rings (superscript r). Essentially this corresponds to adding an extra term to the mutual friction, which has the same form as the standard Gorter Mellink mutual friction considered in the study of isotropic turbulence in laboratory superfluids [227]. The dependence of the total mutual friction force on the lag ω and on the microphysical mutual friction timescale $\tau_{\rm MF}$ is thus:

$$f(\omega) \propto \alpha_T \frac{\omega}{\tau_{\rm MF}} + \beta \left(\frac{\omega}{\tau_{\rm MF}}\right)^3 , \tag{41}$$

where the coefficients α_T and β depend on the rotation rate of the superfluid $\Omega_{\rm s}$ [228, 73].

Having presented the theoretical assumptions underlying calculations of the coupling functional form and timescales in superfluid neutron star interiors, we now examine the different observed stages of the glitch spin-up and following recovery and the kind of constraints they can put on glitch models.

5.3. The glitch rise

At the typical observational resolution, glitches appear as discontinuous changes in ν and $\dot{\nu}$. By the time the first post-glitch observations are carried out, any superfluid component that reacts at a shorter timescale has caught up with the spin-up of the normal component. This means that in most cases it is not possible to use observations to probe the fast processes involved in the early glitch stages.

Fortunately, a few pulsars are being monitored frequently enough that their glitches have now been caught in the act; this has provided us with invaluable information both about the timescale of the spin-up (related to the glitch driving mechanism) and the very early relaxation, as angular momentum is being shared between stellar components (believed to be related to the faster coupling of all, or most, of the core superfluid). Additionally, the two pulsars for which constraining the rise timescale has been possible, the Vela pulsar and the Crab pulsar, present a very different picture.

Vela pulsar

The Vela pulsar (PSR J0835-4510) emits in the radio band and from mid-IR all the way up to > 50 GeV γ-rays. It is the brightest pulsar in the Southern radio sky and has been a target of intense monitoring since its discovery in 1968 [229]. Nearly continuous monitoring over decades has allowed the detection of many glitches including three large events that occurred during high time-resolution observations. The first such glitch had a magnitude of $\sim 3.5 \times 10^{-5}$ Hz and occurred in 2000. It was seen by the 14 m diameter antenna at the Mount Pleasant Radio Observatory, which had a parallel single-pulse observing system. Besides 2-min integrated pulse profiles, 10-sec folded data

were constructed using the single pulses. For the period around the glitch, the single-pulse data were also recorded and maintained in storage. These observations led to two very important results [109]. First, they placed a strong upper limit in the acceleration timescale of the crust: the glitch rise time $\tau_{\rm rise}$ was confined to under 40 seconds. Secondly, they revealed a fast decaying $\Delta\nu_{\rm d}$ component, associated with an exponential relaxation timescale of ≈ 70 s, on which we focus in the next section. Emission properties appeared unaltered around the glitch epoch: the profile shape, intensity or polarization in the radio showed no changes before or after the glitch [109], and a campaign to detect changes in the X-rays found no pulse profile or flux changes between observations taken 3 and 35 days post-glitch [230]. In 2004 another large ($\sim 2\times10^{-5}$ Hz) glitch was captured by the same telescope, for which a similar constraint for the spin-up timescale (< 30 s) was obtained [90]. A third glitch whose initial phase was observed by the Mount Pleasant Radio Observatory, this time with a 26 m telescope, took place in 2016. The frequency evolution around this glitch is displayed in figure 17. It is possible that the spin-up was preceded by a enhancement of the spin-down rate for some tens of seconds before the inferred glitch epoch [110]. The spin-up itself was again undetectable and a tighter limit of $\tau_{\rm rise} \leq 12.6$ s was put on the rise timescale [91, 110].

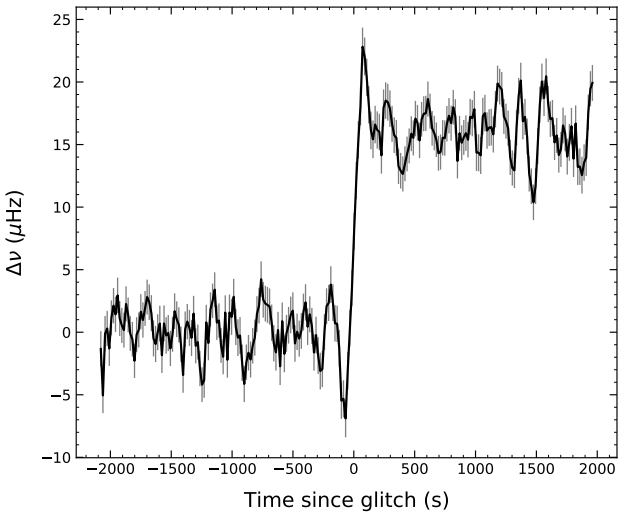

Figure 17. Spin frequency evolution of the Vela pulsar over 1.12 hours and across the large glitch of 2016. We plot $\Delta\nu = \nu - 11.186433306\,\text{Hz}$, using the Mount Pleasant ToAs made available by [91]. Each frequency data point is calculated from a timing fit to 1800 ToAs (equivalent to approximately 161 s) in which only ν was varied and $\dot{\nu} = 0$, using a moving window with a stride of 17.3 s.

Under the assumption that the angular momentum of the glitch-driving region is of the order of few percent of the total moment of inertia of the star, i.e. $I_{\rm gl} \approx 0.01(I_{\rm s} + I_{\rm c})$, then a simple analysis of (32) reveals that $\mathcal{B} \gtrsim 1/(2\Omega_{\rm s}\tau_{\rm rise})$ in the $I_{\rm gl}$ region. Using the constraint for $\tau_{\rm rise}$ of the Vela 2016 glitch, and taking $\Omega_{\rm s} \approx \Omega \approx 70$ rad/s, leads to $\mathcal{B} \gtrsim 10^{-4}$. This is consistent with strong coupling due to kelvon excitations in the crust, or vortex/flux tube cutting in the core of the neutron star, and still marginally consistent with electron scattering off vortex cores also in the core of the star. On such short timescales, however, the regions of the star that are coupled on a longer timescale (e.g. regions of the outer core, or in the crust where phonon excitations provide the main source of coupling [231]) can decouple from the spin-evolution, and a careful analysis of the density dependence of mutual friction and of the moments of inertia involved in the glitch is necessary to obtain quantitative constraints.

Crab pulsar

The Crab pulsar has been also extensively monitored from the radio band to high energies and timing of its pulses is regularly performed by a plethora of observatories and instruments. One of the most comprehensive sets of Crab timing data comes from the Jodrell Bank Observatory, at which daily observations of the Crab pulsar have been carried out over many years (since 1984) [80]. Recent twice-daily ToAs can reach a high time resolution; shorter timescales can also be probed by integrating around 30 mins of observations to obtain somewhat less precise, but denser ToAs [111].

Frequent observations made possible the temporal resolution of part of the spin-up for six of the Crab's largest glitches. These glitches began with an unresolved change in spin frequency, $\Delta\nu_{\rm inst}$, just like in the aforementioned Vela glitches (and as is the case for most observed glitches, although for most pulsars scarce ToA coverage must be factored in). Remarkably though, following this initial step, the pulsar kept spinning up for about one day, before the standard recovery process took over in the timing residuals.

The first such event was observed in 1989 [232]. The initial, unresolved, change in frequency was $\Delta\nu_{\rm inst} \simeq 1.85$ μHz and was completed in less than 2.5 hours. The spin frequency then gradually continued to rise by about 0.7 μHz for around 19 hours. The other partially-resolved glitches happened in 1996 (preceded by an enhanced slow down of the pulsar for around 20 days), 2004, 2011, 2017 and 2019 and had similar behaviour, with the rise times for the delayed spin-up spanning from about 12 hours to nearly 2 days, and the subsequent recoveries having characteristic timescales about 1 to 3 weeks long [112, 233, 192, 111]. All these events are amongst the largest Crab glitches in terms of $\Delta\nu$ amplitudes, the 2017 glitch being the greatest observed so far with $\Delta\nu \sim 14$ μHz.

Potentially of relevance due to a similar timescale

is the observation that several small glitches of the Crab pulsar (which lack the delayed rise component) present a deviation from the exponential-like decay that can be described as a gradual rise of the frequency over ~ 1 day. These "secondary" spin-ups take place 20 to 40 days after the typical glitch, which led to suggestions they are a form of aftershocks [112], although a timing noise origin cannot be excluded.

The difference between the rise in glitches of the Crab and Vela pulsars is quite striking, and it is natural to investigate whether the same process can cause these two kinds of behaviour, or whether their glitches have a different origin. Concerning the latter possibility, it has been proposed that crust quakes play a fundamental role in triggering glitches in the Crab pulsar [234]. The idea is that stresses built up in the crust, also due to pinning [235], and are released by a crust breaking event. One suggested model assumes that the crust will fracture into plates, which move closer to the rotational axis carrying with them a large number of pinned vortices [236]. Besides some initial unpinning responsible for the fast rise, the inward movement of vortices leads to an accelerated creep rate and an increased spin-down of the superfluid, corresponding (due to angular momentum conservation) to the observed extended spin-up of the normal component. Comparison of the model with observations shows that approximately 10^3 plates with typical length-scales $\approx$ 10 m, must move inwards, each carrying $\sim 10^{11}$ pinned vortices. It is not, however, addressed whether it is realistic for vortices to remain pinned to the moving plates; additionally, molecular dynamics simulations (although performed on a smaller scale) suggest that in the extreme conditions of pressure and gravity that characterize a neutron star, the crust will not fracture, but simply fail [237]. Future simulations of crustal pinning and stress release will be necessary to test whether such a scenario is possible.

The alternative has also been explored: that whilst the glitch mechanism is the same for both the Crab and Vela pulsar, internal conditions are different, for example due to the different age of the star, the strength of potential pinning in the core, or the fact that the glitch-driving superfluid $I_{\rm gl}$ is in a different stellar region. A more gradual rise can be naturally obtained if the dependence of the mutual friction on ω is non-linear [195]. The behaviour can differ depending on the fraction $\gamma_{\rm pin}$ of pinned vortices. Similarly, if the interaction of neutron vortices and fluxtubes leads to pinning involving a large number of fluxtubes, then a gradual rise can occur. If on the other hand vortices pin only to a small number of fluxtubes, the result can be a more rapid glitch rise [58]. On the premise that glitches are caused by sudden heating – which facilitates unpinning and leads to an increased creep rate – the internal temperature controls the observed signature [238]. For the same input energy the relative change in temperature will be smaller, and the resulting spin-up slower, for the young (and presumably hotter) Crab pulsar than in Vela. A two-phased rise is predicted, with the extended rise due to the diffusion of the heat pulse taking about one day for Crab-like parameters, as indeed observed [239].

In models which include vortex accumulation, e.g. in a sheet at the edge of the strong pinning crustal region [188], a dependence of the vortex areal density n_v on derivatives of the lag ω with respect to the cylindrical radius ϖ arises [231]:

$$n_v \propto \frac{\partial \omega}{\partial \varpi}\,, \tag{42}$$

which, as mentioned earlier, leads to an evolution equation for the lag of the form:

$$\frac{\partial \omega}{\partial t} \approx \varpi \mathcal{B} \omega \frac{\partial \omega}{\partial \varpi}\,. \tag{43}$$

The solutions of this equation are outward travelling waves, and the timescale it takes for the wave to cross the glitch region and re-establish equilibrium is the timescale on which the delayed rise occurs. Furthermore the additional non-linear terms can also give rise, in particular cases, to glitch precursors, such as the one observed in the Vela 2016 glitch [226].

The above model can be applied to both the Crab and Vela pulsar observations to understand their origin [240]. Assuming the unpinning region is in the inner crust, and solving for typical parameters of the Crab glitches with delayed rise, allows one to estimate a mutual friction parameter $\mathcal{B} \approx 10^{-5}$. This is in good agreement with estimates of mutual friction due to phonon coupling in the inner crust. A much lower value of $\mathcal{B} \approx 10^{-7}$ would be necessary to explain observations if the unpinning region is located in the core. Such weak mutual friction is not predicted for the core by most models [43] – it could, however, be consistent with a large number of pinned core vortices, as discussed previously [58]. Applying the same model to the Vela pulsar, the absence of a delayed rise implies $\mathcal{B} \gtrsim 0.1$ if the glitch is triggered in the crust, which would be an extreme value for kelvon mutual friction. If the glitch is triggered in the core however, one has $\mathcal{B} \approx 10^{-3}$, which is more consistent with expectations if mutual friction is due to electron scattering off vortex cores.

These models thus suggest a core origin for Vela glitches, and show that the same process can lead to the kind of delayed rise observed in the Crab pulsar if its glitches occur in the inner crust, or if there is strong vortex-fluxtube pinning in its core. Whilst more work and future observations are necessary in order

to distinguish between the possible scenarios for the Crab pulsar and pin down the location of its glitch driving region $I_{\rm gl}$, the notion of core-driven glitches for the Vela pulsar is further supported by the constraints on $I_{\rm gl}/I_\alpha$ presented in Section 4.2, which are obtained with a completely independent method.

5.4. Post-glitch response

The rotational evolution following glitches provides the strongest observed signature of superfluid dynamics and can be a distinguishing characteristic of glitches compared to other timing irregularities. At the same time, it presents an extremely diverse phenomenology, not just from pulsar to pulsar but also for glitches of the same neutron star (as shown in figure 6), pointing to the complexity of the underlying mechanism. A sketch of post-glitch recovery can be found in figure 18, where the main attributes are introduced. These features are not always present and can appear in different combinations. Furthermore, not all glitches are accompanied by measurable increases in $\dot{\nu}$ (although most large ones are).

The textbook post-glitch evolution is a partial recovery of the main parameters ν and $\dot{\nu}$ which can be modelled by one or more exponential terms, as in equation (15). Such exponential relaxation presents after both large and small glitches, with a vast range of characteristic timescales (from minutes to years) among neutron stars. An example of this behaviour can be seen in figure 5, in panel (d): the spin-down $\dot{\nu}$ decreases abruptly at the glitch epoch and then relaxes towards its pre-glitch value in a quasi-exponential way, leaving a persisting change $\Delta\dot{\nu}_{\rm p} < 0$. As a result, the spin frequency (panel (c)) decays quasi-exponentially after its initial abrupt change at the glitch, and eventually falls below the prediction of the pre-glitch solution.

In those pulsars where glitch activity appears rather regular, such as the Vela pulsar [29], glitches typically show a similar set of relaxation timescales (characteristic of each pulsar). Following the exponential relaxation, the evolution might continue with a quasi-linear increase in $\dot{\nu}$ (decreasing spindown rate $|\dot{\nu}|$, at an approximately constant $\ddot{\nu}$ rate); other times, this linear recovery is all that can be discerned observationally, most likely because the exponential phase cannot be resolved rather than due to a true absence of it [100]. It is worth noting that it is not always possible to differentiate in the data between an exponential decay with very long timescale ($\mathcal{O}(10^3)$ days) and a quasi-linear decay, and whenever the interglitch interval is long enough to allow deduction of $\ddot{\nu}$ at different post-glitch periods, there is some evolution – usually a decreasing $\ddot{\nu}$ – during the quasi-linear phase (e.g. [241] for the case of PSR J0537−6910).

It is often the case that glitch-induced shifts in the rotational parameters persist until the next event takes place, or until the end of the dataset. In some cases the change in $\dot{\nu}$ shows little to no relaxation (and so eventually the pulsar rotates slower than it would have if not for the glitch), in yet others $\Delta\dot{\nu} \approx 0$ and so the glitch spin-up does not recover at all. The lack of recovery is more common among older pulsars, which also tend to undergo smaller glitches than younger neutron stars [126, 27]. Much of the variety observed in post-glitch recoveries can be understood within the standard picture of glitches, as a result of different stellar regions involved and an interplay between their respective coupling timescales [39].

A parameter often quoted in the literature is the glitch recovery fraction (often called healing parameter) $Q \equiv \Delta\nu_{\rm d}/\Delta\nu$, that is, the fraction of the glitch amplitude that recovers exponentially. Whilst high values of Q are observed after both small and large glitches, low Q values – of the order of few percent – are typical of large (giant) glitches [28]. A collective distribution of measured Q values using data from the ATNF glitch catalogue § shows some bi-modality, with a lower broad peak around 10^{-2} and a high narrow peak of $Q \lesssim 1$; many of these high Q measurements concern glitches of the Crab pulsar [28, 85].

Atypical glitch recoveries are sometimes seen, mainly after small and intermediate size glitches as well as in highly-magnetised neutron stars and magnetars (Section 6.2). For example, the measured change in $\dot{\nu}$ at the glitch epoch is positive for nearly 15% of known glitches. Such $\Delta\dot{\nu} > 0$ measurements are not predicted by the simpler superfluid models. Usually however these values are small, sometimes consistent with being zero given their uncertainties, and may be reflecting timing noise instead of the response to the glitch (and as discussed earlier in Section 3, some of those recorded events might not be glitches at all). An example of such a "glitch", where an initial unresolved $\Delta\nu > 0$ was followed by a continuous increase in frequency relative to the pre-event solution ($\Delta\dot{\nu} > 0$), was observed in the relatively old ($\sim 10^7$ yr) PSR J0147+5922 [31]; the same study presents timing residuals for several pulsars, showing most of the kinds of post-glitch evolution discussed above.

In the Crab pulsar, the post-glitch evolution of $\dot{\nu}$ following its largest glitches is better described when an asymptotic exponential term is included (increasing $|\dot{\nu}|$), along the common decaying terms. Essentially, the post-glitch behaviour is as follows: the spin-down rate $|\dot{\nu}|$ increases at the glitch as it is common (after $\tau_{\rm rise}$, if resolved). Then there is an observed partial

§ `www.atnf.csiro.au/research/pulsar/psrcat/glitchTbl.html`

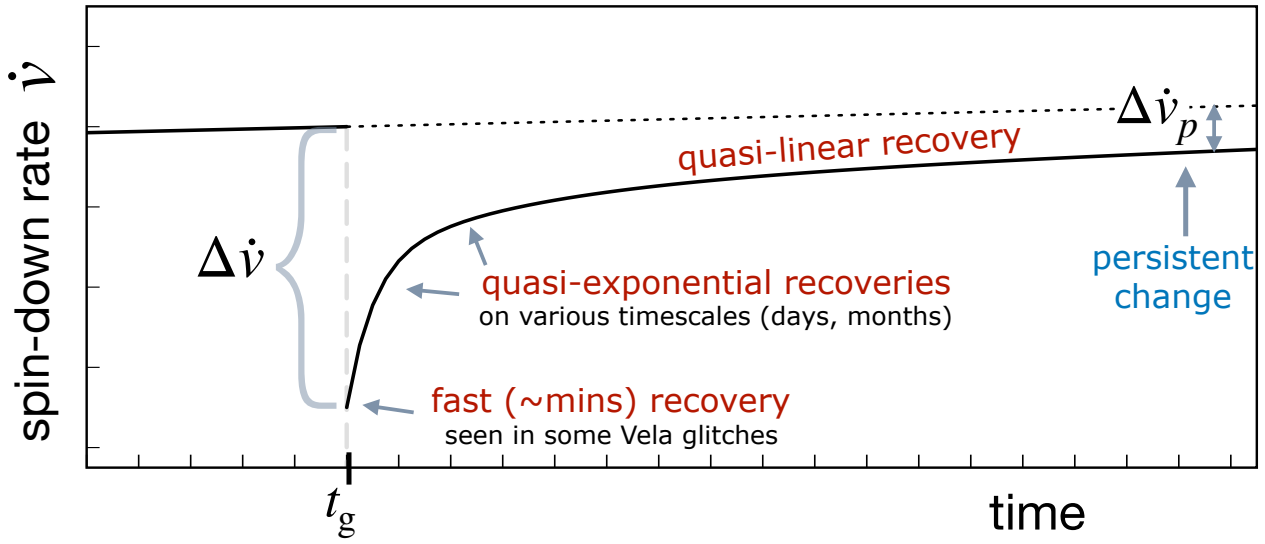

Figure 18. A sketch of the various components observed in pulsar glitch recoveries (time progresses to the right, and we recall that $\dot{\nu} < 0$ therefore lower values indicate faster spin-down). A very fast decay with characteristic timescale of ~ 1 min has been observed in three Vela pulsar glitches. In many pulsars post-glitch recovery can be characterised by a single or multiple exponentially decaying terms, on various timescales (usually $10 - 100$ days). In a subset of pulsars the evolution of $\dot{\nu}$ proceeds with a nearly constant $\ddot{\nu}$. Glitches often leave persisting offsets in the spin frequency and/or the spin-down rate.

recovery of this transient, i.e. $|\dot{\nu}|$ decreases, which dominates the residuals for the first ~ 100 days after the glitch. This recovery is not complete and leaves a negative $\dot{\nu}$ offset which grows quasi-exponentially over time (a $|\dot{\nu}|$ increasing term). This last term has a long characteristic timescale of ~ 320 d [80]. So far this feature appears unique to large glitches of the Crab pulsar; future observations will hopefully bring more information about its prevalence.

Other atypical examples include recovery that results in a net decrease in spindown rate ($\Delta\dot{\nu}_{\rm p} > 0$) as opposed to the usual persistent increase seen after most glitches. At least two glitches have shown this behaviour, one in the rotationally-powered pulsar J1119−6127 [108] and one in the Rotating Radio Transient (RRAT) J1819-1458 [159]. As both these neutron stars have high inferred magnetic fields and radio emission changes associated with those glitches, they will be discussed separately in Section 6.2.

After this brief overview of observed post-glitch recoveries, each attribute depicted in figure 18 is examined below.

Shortest-term transients

A rapidly decaying term has been discerned following the three Vela glitches for which high-resolution observations were available. It was first detected after the 2000 glitch and its associated timescale was estimated around 70 s [109]. The same feature was identified again in the 2004 giant glitch, when a big fraction of the initial spin-up decayed on a timescale of ~ 60 s [90]. Similarly, in the 2016, a large component $\Delta\nu_{\rm d} \sim 18$ μHz decayed on a ~ 50 s timescale [91, 110].

Both the upper limits on the glitch rise time and the rapid components of the post-glitch recovery suggest that part of the core does not remain coupled during the glitch. From the rise timescale $\tau_{\rm rise} \lesssim 12.6$ s we have already estimated $\mathcal{B}_{\rm gl} \gtrsim 10^{-4}$ (Section 5.3) in the glitch driving region $I_{\rm gl}$. Any superfluid region of the core with $\mathcal{B} \lesssim 10^{-4}$ will be coupled on a longer timescale, as $\tau \propto 1/\mathcal{B}$, therefore cannot follow the spin-up and its response can be associated with the short-term recovery observed [231]. As an example consider a region of the core, denoted by the index 'm', with an average $\mathcal{B}_{\rm m} \simeq 5 \times 10^{-5}$. The associated coupling timescale $\tau_{\rm m}$ is of the order of a minute. Before the glitch, the component $I_{\rm m}$ will have had ample time to couple to the spindown of the normal component so that $\dot{\Omega}_{\rm s,m} = \dot{\Omega}_{\rm c}$. The equilibrium lag between the two will be $\omega_{\rm eq,m} \simeq 10^{-8}$ rad/s, calculated from (3) and parameters appropriate for the Vela pulsar. During the glitch, however, this superfluid cannot keep up with the spin-up of the normal component, since $\tau_{\rm rise} < \tau_{\rm m}$. The glitch change, $\Delta\Omega_{\rm c} \approx 10^{-4}$ for large Vela glitches, is greater than $\omega_{\rm eq,m}$, so that immediately after the spin-up, the lag of this region is inverted. This is sometimes called an "overshoot" (as post-glitch $\Omega_{\rm c}$ "overshoots" the local $\Omega_{\rm s,m}$), and the subsequent recoupling of this region on $\tau_{\rm m}$ gives rise to the observed minute-timescale components of the recovery. Unfortunately,

the amplitude $\Delta\nu_{\rm d}^{\rm m}$ of this fast decaying component of the recovery is not well constrained from observations. If we assume that 2/3 of the initial $\Delta\nu$ decays on a $50-70$ s timescale, which is roughly consistent with the findings from the 2004 and 2016 glitches, then

$$\frac{I_{\rm m}}{I_{\rm c}} \simeq \frac{\Delta\nu_{\rm d}^{\rm m}}{\Delta\nu - \Delta\nu_{\rm d}^{\rm m}} \approx 2\,. \tag{44}$$

A full analysis of this scenario, allowing for differential rotation of the superfluid component, has been carried out with realistic equations of state and including strong crustal entrainment [182]. It has been shown that even in this case, where a notable part of the core decouples during the glitch (as suggested by (44), a crust-confined $I_{\rm gl}$ region cannot reproduce the glitch sizes seen in Vela if its mass is greater than $1M_\odot$.

This can also be seen with a simplified three component model, with three components $I_{\rm c}$, $I_{\rm gl}$ and $I_{\rm s}^{\rm r}$ as discussed in Section 2. Fitting such a model to the Vela 2016 glitch data allows one to constrain the average mutual friction parameter for the region of the core that decouples to indeed be in the range [186]:

$$\mathcal{B}_{\rm s}^{\rm r} \approx 2.8\times10^{5} - 6.3\times10^{5}\,. \tag{45}$$

The situation is less clear cut for the glitch driving region, as the inferred value of $\mathcal{B}$ depends on entrainment, and therefore on whether one assumes the region to be in the crust or core. Assuming for now that the glitch is triggered in the crust, one has:

$$\mathcal{B}_{\rm gl} \approx 2.4\times10^{-3} - 0.7\,, \tag{46}$$

using the entrainment values calculated by [171]. This estimate is in agreement with results obtained by contrasting the 2016 glitch observations to a model that uses a density profile $\mathcal{B}_{\rm crust}(r)$ for kelvon mutual friction in the crust based on realistic microphysical parameters (but without considering entrainment, and assuming a uniform $\mathcal{B}_{\rm core}$) [242]. However, despite the values of the mutual friction coefficients being consistent with the expectations for kelvon mutual friction in the crust, a crustal origin for the glitch is not fully compatible with the data. First of all, the above solution is not unique. If entrainment is weak (for example in the case of an amorphous crustal lattice [243]), or if the glitch driving region is assumed to extend into the core, then one has

$$\mathcal{B}_{\rm gl} \approx 5.6\times10^{-4} - 0.17\,, \tag{47}$$

which widens the possible range to that expected for electron scattering mutual friction in the core [179]. Secondly, and crucially, fitting for the relative moment of inertia $x_2 = I_{\rm gl}/I_{\rm tot}$ of the angular momentum reservoir (still with the same three component model [186] and the data of the 2016 Vela glitch), reveals that, even in the absence of entrainment, $I_{\rm gl}$ cannot be in the crust alone. This can be seen in figure 19, where the value of x_2 (the relative moment of inertia of the glitch triggering region) is compared to the moment of inertia available in the crust. It is clear that for all values of the stellar mass the crustal moment of inertia is not sufficient.

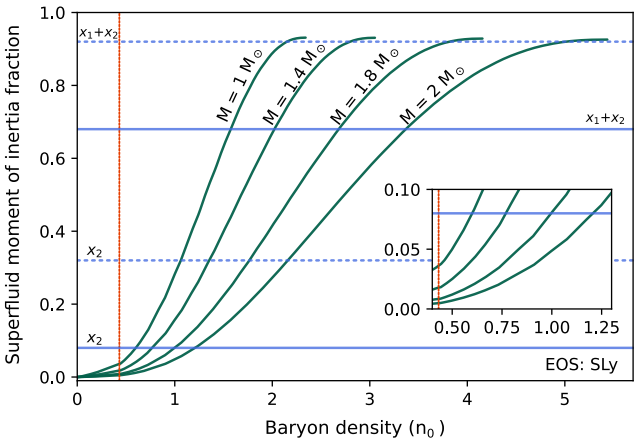

Figure 19. Relative moment of inertia of the superfluid triggering the glitch (x_2) and of the total superfluid including the fraction of inertia of the reacting superfluid (x_1), compared to the moment of inertia available in the star for several stellar masses calculated with the SLy4 equation of state [244]. The horizontal lines are the results of the fits for x_2 and (x_1+x_2) to the Vela 2016 glitch by [186], with full lines indicating the case without entrainment, the dotted lines including realistic entrainment in the crust. The red vertical line indicates the crust/core transition, with the inset zooming into this region for clarity.

The growing support to the idea that at least part of the angular momentum reservoir for the large glitches in Vela must extend beyond the inner crust, motivates further investigation to potential vortex pinning in the core. A three-component model including both a pinned and a free core superfluid has been developed [58], in which the mutual friction strength depends strongly on the number $N_{\rm p}$ of proton fluxtubes to which neutron vortices can pin [57, 59]. As outlined in Section 5.3, for low values of $N_{\rm p}$ the spin-up is rapid; and additionally, a fast decay is present. A lower limit can be set on the moment of inertia of the free superfluid $I_{\rm f}$ as:

$$\frac{I_{\rm f}}{I} \geq 1 - \frac{\Delta\nu_{\rm pt}}{\Delta\nu} \tag{48}$$

where $\Delta\nu_{\rm pt} = \Delta\nu - \Delta\nu_{\rm d}^{\rm m}$ is the change in ν after any initial rapid (on timescales from seconds to minutes) decay has finished and $\Delta\nu$ is the instantaneous glitch size. On the other hand, if $N_{\rm p}$ is larger than a critical threshold, the rise is much more gradual and the rapid decay absent, resembling the early glitch phase observed in Crab glitches.

A number of independent models, based on the multifluid formalism, thus point to the core of the neutron star playing a role in the large glitches of the Vela pulsar, and possibly other glitches. More refined models of core pinning, and vortex-fluxtube interactions are, however, necessary to investigate this issue fully. In particular, many models rely on toroidal components of the magnetic field in the interior to provide the pinning region and the angular momentum reservoir [245]. Realistic equilibrium models for magnetic fields in superconducting neutron stars, however, suggest that the toroidal field is entirely confined to the crust of the star for magnetic field intensities like the ones expected in regular pulsars [183, 184]. Additional theoretical work is required to understand if this scenario is viable.

Exponential-like relaxation on intermediate timescales

Whilst the very fast relaxation is hard to observe without dedicated timing campaigns for selected, bright, targets (like the nearly continuous, high resolution, Vela monitoring), quasi-exponential recoveries on various longer timescales are seen in many pulsars – mainly in the younger glitching sources.

Usually a single exponential term is used in (15), and fitted timescales $\tau_{\rm d}$ range from a few up to several hundred days. However it must be noted that there is observational bias (and degeneracy) in both the number of decaying terms that are fitted for and the resulting timescales. A low density of available ToAs disfavours the detection of relatively fast timescales; for typical ToA spacing of 2 to 4 weeks, recoveries with $\tau_{\rm d} \lesssim 30$ days can be missed. Another factor is the length of post-glitch interval examined, which can sometimes be too short to reveal any relaxation that occurs on timescales of months to years. Additionally, exponential timescales comparable to or longer than the time interval until the next glitch (or the time span examined, if shorter than the interglitch interval) can sometimes offer a similarly adequate description of the data as a $\Delta\ddot{\nu}$ term. Resolving this ambiguity is hard for most datasets. In glitches for which a second decaying term has been included in the timing model, often one component has timescale $\tau_{\rm d,1}$ between a few days up to about two weeks, whilst the second component has a longer timescale of the order $\tau_{\rm d,2} \sim 100$ d.

Relatively fast decays ($\tau_{\rm d} < 2$ weeks) are not commonly seen in older pulsars, although in many cases the observing cadence would have allowed their detection. There are hints for a positive correlation between $\tau_{\rm d}$ and the pulsar's characteristic age [31, 28]. This trend can be explained within the theoretical framework of Section 2. From equation (32), and noting that $\Omega_{\rm s}$ cannot be very different than $\Omega_{\rm c}$, an inverse scaling of $\tau_{\rm d}$ with the pulsar's frequency $\nu = \Omega_{\rm c}/2\pi$ is expected. At the same time, the characteristic age of glitching pulsars anti-correlates with ν, especially when high-magnetic field ($B \gtrsim 10^{13}$ G) neutron stars are excluded. Therefore the observed trend could simply reflect the fact that components will re-couple faster in a more rapidly rotating pulsar. Possibly it can also be amplified by the fact that younger pulsars have higher $|\dot{\nu}|$, therefore the rotational lag is restored faster after the glitch.

As already noted, the Crab pulsar displays a high recovery fraction $Q \lesssim 1$, with an almost complete decay of the initial $\Delta\nu$ within the first month after the glitch. Typical timescales, obtained from fitting (15) including two exponential terms, are $\tau_{\rm d,1}$ =2–4 d and $\tau_{\rm d,2}$ =10–20 d [112]. Contrary to ν, the spin-down rate does not show a complete recovery but instead there are persisting changes $\Delta\dot{\nu}_{\rm p}$ that appear to correlate with the glitch size (see Section 4.2); moreover on three occasions (all large glitches rather isolated in time from other events) the asymptotic decrease in $\dot{\nu}$ described earlier was observed [80].

In the Vela pulsar several recovery timescales are identified which are similar between its large glitches, as can be seen in figure 20. Besides the very fast, $\tau_{\rm d,1} \sim 1$ minute, decay, another rapid decay happens on $\tau_{\rm d,2} \sim 0.5$ d [109, 90]. For the 2000 and 2004 glitches additional components included a $\tau_{\rm d,3}$ ∼2–3 d and $\tau_{\rm d,4}$ ∼20–30 d. Following the 2016 glitch, timescales $\tau_{\rm d,2} \sim 1$ d [246], $\tau_{\rm d,3}$ ∼5–6 d [91, 246] and $\tau_{\rm d,4} \sim 32$ d [93] have been reported, consistent with the findings for previous glitches. Note that for this glitch no study has modelled all four relaxing components simultaneously as they use different observations (with different time resolution and ToA cadence) and varying lengths of the post-glitch time span were fitted for.

The multiple exponential components of the relaxation were interpreted early on within the vortex creep model as the signature of different regions in the linearly coupled regime responding to the glitch on their own intrinsic timescale [247, 166]. The need to fit multiple exponentials, with vastly different timescales, thus suggests that the spatial dependence of the mutual friction coefficients and pinning forces have a strong influence on the shape of the recovery, as shown also by simulations that allow for a radial dependence of these quantities [242] and differential rotation [231, 39]. Of course a one to one correspondence between observably inferred relaxing components and physically distinct stellar layers ignores both the physical complexity of the problem, as well as the uncertainties associated with the data fits (e.g. that two different descriptive timing models might fit observations equally well) . Nonetheless, the longer timescales of the order of 100 days or greater, typically seen in post-glitch recoveries,

are indicative of the response of superfluid regions of low effective $\mathcal{B} \sim 10^{-9}$, close to the lower estimates for coupling to phonon excitations [216, 44]. Very slow re-coupling timescales can characterise crustal pinning regions, for which the fraction of free vortices can be small when the local lag is far from its critical value for unpinning (recall that the effective $\mathcal{B}$ entering equation (32) connects to the microphysical $\mathcal{B}_0$ as $\mathcal{B} = \gamma_{\rm fv}\mathcal{B}_0$).

Once the exponential-like decaying components fade away, many pulsars still display persisting changes in their rotation (mainly a $\Delta\dot{\nu}_{\rm p} < 0$). In some sources, like the Vela pulsar and other neutron stars which show similar glitching behaviour, dominated by large ($\Delta\nu \gtrsim 10$ μHz) events, the longer term recovery consists of an approximately linear ($\ddot{\nu} \approx$const) evolution of $\dot{\nu}$ (preceded or not by distinguishable exponential components, see figures 18 and 20). Neither this quasi-linear evolution nor the persisting changes can be explained by the response of linearly coupled regions, as these would eventually lead to a full recovery of semi-exponential form. In the following we consider these two features in turn and discuss their theoretical implications. Before that, however, we caution again that there is often ambiguity in the post-glitch recovery results obtained from data fits. Different combinations of parameters in equation (15) often lead to equally good descriptions of the ToAs (some examples can be found in [126] and [108]). It is particularly difficult to exclude an exponential (or expansion in a basis of exponentials) with very long characteristic timescale (comparable or greater to the interglitch interval) based on the observational data; nonetheless such a functional form is in general not favoured over a simpler linear (in $\dot{\nu}$) evolution [100, 84]. The choice of functional parametrization of the recovery affects also the inferred persisting changes following a glitch. For example, a study of 21 years of rotation of the Vela pulsar which included 8 glitches, concluded that their recoveries can be described by two common exponential characteristic timescales of $\tau_{\rm d,1} \simeq 25$ and $\tau_{\rm d,2} \simeq 1300$ d (shorter timescales will be missed because of scarce, approximately monthly, ToAs used in that study) [107]. In this case, due to the presence of the $\tau_{\rm d,2} \simeq 1300$ d term, the extrapolated recovery fraction Q will be larger than previously thought (for models that do not include the long-timescale exponential but only shorter decaying terms and a $\ddot{\nu}$ term instead), leading to considerably smaller inferred persisting changes $\Delta\nu_{\rm p}$.

The quasi-linear $\dot{\nu}(t)$ recovery phase

The phase of quasi-linear $\dot{\nu}$ recovery is typically characterised by an interglitch $\ddot{\nu}_{\rm ig}$ which is larger than the prediction of equation (20) for dipole braking, and often has a similar, but not always the same, value when

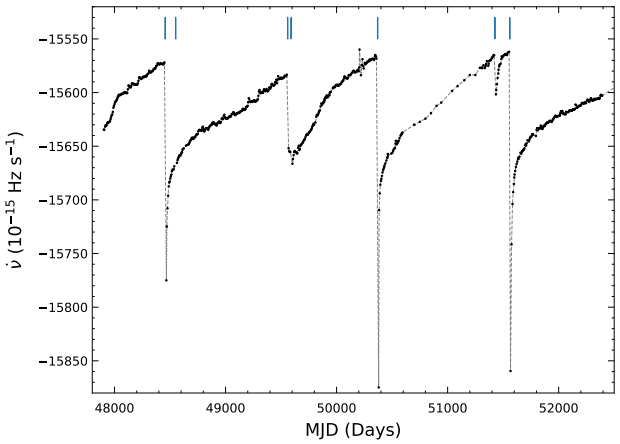

Figure 20. The $\dot{\nu}$ evolution of the Vela pulsar over nearly 13 years (data from [92] and [248]). The epochs of the seven glitches detected during this period are marked with blue vertical lines at the top of the plot. For most glitches the initial recovery is characterised by a few exponential-like decaying components, then a phase with approximately constant, high, $\ddot{\nu}_{\rm ig}$ follows until the next glitch.

measured in different interglitch intervals of the same pulsar. These attributes can be noticed in figure 20, which includes several Vela glitches. In general, measurements of $\ddot{\nu}$ are rather sensitive to timing noise and so fluctuations appear when $\ddot{\nu}_{\rm ig}$ is measured over different time spans of the interglitch interval, but overall a constant $\ddot{\nu}_{\rm ig}$ value provides a good description of the residuals after the exponential components have been subtracted and up to the epoch of the subsequent glitch (in pulsars with multiple glitches).

In PSR J0537−6910 there are hints for a monotonic decrease of $\ddot{\nu}_{\rm ig}$ as a function of post-glitch time [241], but the shortness of interglitch time intervals (~ 100 days on average) and density of ToAs impede a more refined analysis. Nonetheless, a linear $\dot{\nu}$ evolution provides an adequate fit [100], and even when including an exponential component following each glitch, a high $\ddot{\nu}_{\rm ig}$ is still required by the data [249].

Using a sample of 32 glitching pulsars for which the slope $\ddot{\nu}_{\rm ig}$ of the linear decay can be measured (of which some were found to be negative), a correlation is seen between $|\ddot{\nu}_{\rm ig}|$ and $|\dot{\nu}|$ [28]. For the same sample, it was also found that the quantity $(\ddot{\nu}_{\rm ig} - \ddot{\nu}_{n=3})$ is nearly proportional to $|\dot{\nu}|/\langle\Delta T\rangle$, where $\ddot{\nu}_{n=3}$ is the expected value according to equation (20), and $\langle\Delta T\rangle$ is the average interglitch interval. Another correlation was retrieved, this time for a sample of 16 pulsars, between $\ddot{\nu}_{\rm ig}$ and $|\Delta\dot{\nu}|/\Delta T_{\rm ig}$, where $\Delta\dot{\nu}$ the change at the previous glitch and $\Delta T_{\rm ig} \equiv \Delta T_f$ the time interval until the next event [85]. By definition $\ddot{\nu}_{\rm ig} = [\dot{\nu}(t_2) - \dot{\nu}(t_1)]/\Delta T_{\rm ig}$, where $\dot{\nu}(t_1)$ the spin-down rate after the preceding glitch (having subtracted any $\Delta\dot{\nu}_{\rm d}$ decaying terms, if present) and $\dot{\nu}(t_2)$ the spin-down

rate right before the subsequent glitch. Therefore the above correlation implies that the amount of $\Delta\dot{\nu}_{1,2} = \dot{\nu}(t_2) - \dot{\nu}(t_1)$ that recovers linearly until the next glitch correlates with its initial change, $|\Delta\dot{\nu}|$, and consequently, with the glitch size $\Delta\nu$ as we have seen that $\Delta\nu$ and $|\Delta\dot{\nu}|$ also appear to correlate (though this should be accessed on a pulsar-to-pulsar basis, see Section 4.2). It should also be noted that an anti-correlation between measured $\ddot{\nu}_{\rm ig}$ and $\Delta T_{\rm ig}$ could perhaps be expected (and indeed is weakly observed in the aforementioned sample of 16 pulsars [85]) if $\ddot{\nu}_{\rm ig}(t)$ decreases over the course of the interglitch interval, as has been suggested for PSR J0537−6910 [241, 101].

This phase of the recovery has been attributed to the response of non-linearly coupled regions. In the framework of vortex creep, following the model of [224], it has been shown that the response of the non-linear region will be the sum of Fermi functions (equation (39)) that will appear as a constant $\ddot{\nu}_{\rm ig}$ [250]. The existence of regions that respond in this step-like manner has found support in observations of Vela post-glitch recoveries [127]. Furthermore, the observed relationship between $\ddot{\nu}_{\rm ig}$ and $|\Delta\dot{\nu}|/\Delta T_{\rm ig}$ can naturally arise within this phenomenological model [251, 85]. Under the assumption that a glitch is triggered when the lag ω – reduced by $\Delta\omega$ at the glitch – has been restored in each and every region, which also means all regions have re-coupled and the recovery is complete and over at the epoch of the next glitch, the model can be used to estimate the epoch of a following glitch [250]. Comparison to observations of the Vela pulsar, although encouraging, unfortunately shows little predictive power for targeted observations. Discrepancies of observed versus predicted glitch epochs are on average over 4 months (and up to nearly 2 years) [250]. To lessen these discrepancies, large permanent shifts in $\dot{\nu}$, of the order $\Delta\dot{\nu}_{\rm p} \sim 2 \times 10^{-3}\dot{\nu}$, have to be assumed to concur with some of Vela's glitches [250].

Another situation where non-linear coupling might manifest is that of superfluid turbulence [40]. The development of turbulence in regions that respond passively to a glitch has been considered as another possible explanation for the quasi-linear recoveries of Vela-like glitches [73]. As vortices are pinned in these regions, vortex rings can be nucleated at pinning sites, adding vortex length to the fluid, and modifying the functional form of the coupling, as in equation (41). In this case one no longer has an exponential recovery, but rather a power-law form for the $\dot{\nu}(t)$ evolution which can be approximated as linear on the timescales of interest. The value of $\ddot{\nu}_{\rm ig}$ in this case depends on the initial value of the lag at the time of the glitch, which is unfortunately not probed by observations. Under the assumption, however, that a larger lag builds up the longer the waiting time to the glitch has been, the model suggests that a correlation of the form

$$\ddot{\nu}_{\rm ig} \propto \Delta T_{-}^{2}\,, \tag{49}$$

should be present between the measured $\ddot{\nu}_{\rm ig}$ for the linear post-glitch response and ΔT_{-}, the waiting time since the previous glitch. Data from PSR J0537−6910 and the Vela pulsar are consistent with such a correlation, and future observations will be able to confirm whether the relation holds [73]. Furthermore, the turbulent model can be used to fit the long term recoveries of Vela glitches, as seen in figure 21, where both the turbulent and laminar (linear) model are compared to the data. It is clear that the exponential solution to the laminar model is a poor fit, as weak mutual friction cannot reproduce the slope of $\dot{\nu}(t)$, while stronger values lead to a clearly exponential trend. The turbulent model, on the other hand, fits very well the quasi-linear part of the recoveries, with reasonable physical parameters for the crust.

Persistent post-glitch changes

Timing solutions of glitches according to (15) often include persistent changes in the timing parameters (ν, $\dot{\nu}$, and $\ddot{\nu}$). For example, the low recovery fractions Q generally observed indicate that the presence of $\Delta\nu_{\rm p}$ is not uncommon. It is often unclear if these are permanent changes or could be part of an ongoing recovery that is interrupted by another glitch or extends beyond the end of the data span. As the right choice of descriptive model is not always strongly constrained by observations, results for persistent changes can also be model-dependent – like in the earlier example for Vela's inferred $\Delta\nu_{\rm p}$, which are greater for a linear-$\dot{\nu}$ evolution compared to a long-timescale exponential [107]. Commonly reported are persisting changes in the spin-down rate, like the $\Delta\dot{\nu}_{\rm p}$ of the Crab pulsar which are generally between 10^{-14} – 10^{-12} s^{-2}. In pulsars where linear-$\dot{\nu}$ decay is seen, its slope $\ddot{\nu}_{\rm ig}$ can change between consecutive interglitch intervals, implying a $\Delta\ddot{\nu}_{\rm p}$; this can for example account for what is observed in PSR J0537−6910 [100]. Whilst $\Delta\ddot{\nu}_{\rm p}$ is generally understood to stem from the superfluid response to the glitch, the changes $\Delta\dot{\nu}_{\rm p}$ are often interpreted as permanent. Because of their magnitude, $\Delta\dot{\nu}_{\rm p}$ terms cannot be attributed to actual changes in the total stellar moment of inertia (e.g. a crustquake) as this would have left a large, permanent, imprint to the spin frequency as well, which is not observed. They can, however, result from an, either permanent or long-lived, change in the effective moment of inertia of the star – that is, the moment of inertia that is slowing down under the external torque – if there is progressive decoupling of the superfluid [100]. For example, this could happen in a sporadic manner due to formation of new strong

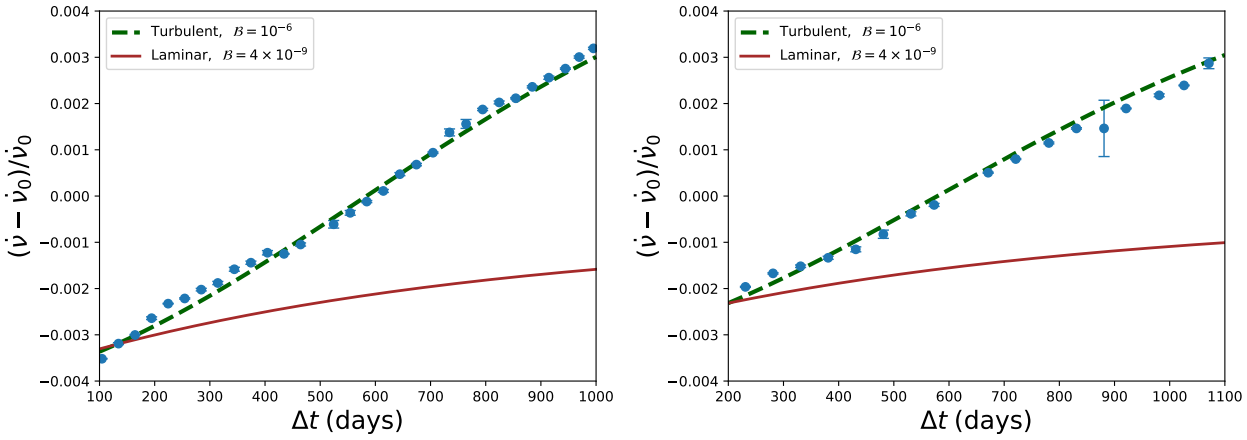

Figure 21. Recoveries of the Vela 1975 (left) and 1996 (right) glitches compared to the turbulent (green dashed curves) and laminar (red curves) models. Blue circles are the observed, normalised, difference in frequency $\dot{\nu}$ with respect to the longterm spindown $\dot{\nu}_0$ for long times Δt after the glitch (excluding the short term transients). The turbulent model provides a good fit for reasonable values of $\mathcal{B}$ in the crust. A more detailed discussion of the setup and the parameters can be found in [73].

pinning regions in the event of a crustquake that led to lattice deformation simultaneously with the glitch [252]. A varying external torque could also be the cause of observed $\Delta\dot{\nu}_{\rm p}$ changes; if (20) holds, a gain in magnetic dipole moment is required to explain $\Delta\dot{\nu}_{\rm p} < 0$, arising for example from a change in inclination angle α_B (see e.g. [253]). Such variations might have observable consequences to the magnetospheric emission (the behaviour of which around the time of glitches is discussed in the following) and could potentially be confirmed or ruled out with future observations.

The analyses presented in this section most often make the explicit or implicit assumption that, throughout the several glitch stages, the external torque acting on the star remains constant. Whilst this is a reasonable first approximation for the majority of neutron stars, based on the typically radiatively-quiet nature of glitches, it might not always be the case. The next section summarises some of the observational evidence, as well as hints, for magnetospheric activity around glitches, and very briefly touches upon their implications for glitch models.

6. Glitches and radiative changes

The general view that all glitches are radiatively quiet events, for which only the interior dynamics play a role whilst the magnetosphere is not altered, shifted when the first observations of magnetars came along. In these neutron stars, which present activity powered by magnetic rather than rotational energy, glitches often occur at the time of high-energy outbursts. Glitches in high-magnetic field pulsars, whether these appear predominantly as magnetars or as rotationally-powered pulsars, deserve a dedicated discussion (Section 6.2). There is, nonetheless, growing evidence that the whole system of neutron star-magnetosphere should be considered when trying to understand glitches, even for typical radio pulsars.

6.1. Glitches and emission changes in regular, rotationally-powered, pulsars

Although pulsars with modest inferred magnetic field intensities do not show dramatic emission changes associated with glitches, more subtle changes have been connected with their timing activity. For the majority of pulsars, observations surrounding the glitch might be several days before and after the glitch epoch; as a result, short-lived variations related to the glitch would be missed. Moreover, integrated pulse profiles will smear out any transient features at the single pulse level. Once again though, the nearby and well-studied Vela pulsar offers some insights.

Detailed studies of single pulses from the Vela pulsar are possible as it is nearby and radio-bright. The pulsed signal shows fluctuations in intensity and profile, in the form of microstructure, transient time-clustering of bright pulses and sporadic giant micro-pulses: a highly-polarised, very bright component leading the main pulse [254, 255, 256, 257, 258]. The X-ray lightcurve shape and radio ToAs appear correlated on a pulse-to-pulse analysis [259]. Brighter radio pulses tend to arrive earlier than the average pulse profile [260], and it was found that earlier radio ToAs also associate with significantly higher X-ray flux density in the main X-ray peak [259]. After a large glitch in 2010 [261], the emission rate of bright radio pulses (defined as 5 times brighter than the average pulse) was seen to increase [257]. An intense monitoring

campaign using the Mt Pleasant 26 m radio telescope (at 1376 MHz observing frequency) produced single-pulse and 10 s-folded data that span several months and coincided with two glitch-like irregularities, in 2014 and 2015‖. Both of these timing events appear associated with emission changes. At the epoch of the first, smallest, glitch-like event, the rate of bright pulses increases (after a rather quiet period) whilst no significant variation of the pulse width is observed. At the second, largest event, the width abruptly reduces whilst the rate of bright pulses does not change considerably [262]. Emission changes were also found around the 2016 giant glitch, timing observations of which have been already discussed in Section 5.3. Seconds before the inferred glitch epoch, an unusually broad pulse was detected, followed by a null (missing pulse) and then two pulses with below-average linear polarisation. Twenty pulses after the null, a series of pulses with increased mean and smaller variance of timing residuals started, which lasted for 29 rotations through the inferred epoch for the glitch. Surrounding the null, a 3σ decrease in pulse peak flux density was also observed, lasting about 2 minutes [91]. Nulls are not unusual in pulsars. By analysing other days from the same dataset, the same study found that approximately one every 77,700 pulses has a flux density below the detection threshold, and calculated a probability of 5×10^{-4} for a null to occur in the time window where the phenomena just described occurred.

There is less information about possible emission changes around the epochs of the Crab pulsar glitches in the radio band and other wavelengths. X-ray observations – typically with low time resolution – do not reveal any changes in integrated profile nor spectrum, and there are only very weak hints for a slight decrease in flux after large glitches, statistically consistent with zero variation [263, 264, 265]. Measurements of soft X-ray polarisation fraction showed a decrease following the 2019 Crab glitch compared to pre-glitch values, which is likely associated with the pulsar and not the nebula. The significance of the variation can be established only for periods after the first 30 days post-glitch (making it hard to confidently infer an association), and the polarisation fraction recovers on a timescale of about 100 days post-glitch [266, 267]. Indications for a changing magnetic field have also come from long-term studies of the radio pulse profile, which display an evolution as increasing rotational separation between pulse components (most prominent between the interpulse and main pulse) and decreasing flux density ratio between components and main pulse. However, this evolution appears secular, cumulative over about two decades, without signs of recovery, hence it might be unrelated to the glitch activity [268].

Connections between glitches and magnetospheric activity have also been drawn for pulsars with less stable pulse profile properties. One such category is state-switching pulsars, ones for which the pulse profile has two or more distinct states and alternates between them. The different states sometimes appear to correspond to different spin-down rates. For example, pulse width or intensity ratio between pulse components was found to correlate with timing variations, and in particular the inferred spin-down rate, for six state-switching pulsars, making for the first time a connection between timing noise and magnetospheric conditions [150]. One of these pulsars, J0742-2822, rapidly switches between two modes of emission (and also displays intermediate states in some profiles) that mostly differ in the emission at the central part of the profile. The rate of mode-switching has been varying over the years [150]. A shape parameter has been defined as the ratio between the intensity of the first and second pulse component at 1369 MHz, close to 1 for one mode whilst less than 1 for the mode with increased central part emission. The timing residuals, and in particular the spindown rate $\dot{\nu}$, show quasiperiodicity of about 100 days. Between 2007 and 2009, the fluctuations in shape parameter did not strongly correlate with $\dot{\nu}$. The situation changed when a glitch of moderate size took place in 2009: beyond the glitch epoch, and persisting for at least 1000 days, a much stronger correlation between the shape parameter and $\dot{\nu}$ was established [269].

A small glitch ($\Delta\nu \sim 12$ nHz) with a large change in spin-down rate, $\Delta\dot{\nu}/\dot{\nu} \sim 67 \times 10^{-3}$ and an unusual decrease in $|\ddot{\nu}|$ was reported for J2037+3621, a relatively old (~ 2 Myr) pulsar [270]. This timing event coincides with a mode change, where the newly emerged state shows a decrease in pulse width compared to the previous, stable, state [150]. The pulsar continued switching between a more narrow pulse profile and a wider one, both with decreased intensity ratio between leading and main components compared to the original state [270]. Another example of possibly glitch-associated emission switch is PSR J2021+4026, a radio-quiet $\gamma-$ray pulsar which displayed sudden radiative changes above 100 MeV, simultaneously with a timing irregularity that could be the result of a glitch. In 2011, around MJD 55850, a sudden increase in $|\dot{\nu}|$ occurred, of relative size $\sim 38 \times 10^{-3}$, on the upper end of observed glitch-related changes. A spin-up has not been detected but cannot be entirely ruled out. The rotational change coincided with a sudden decrease in flux by about

‖ These two events are much smaller than typical Vela glitches (4 nHz and ∼750 nHz respectively), with timing residuals consistent with that of another population of rotational features seen in this pulsar [92]. It is as yet unclear if they belong to the standard glitch population.

17% and significant changes in pulse profile [271]. The new state persisted for approximately 3 years and then both emission and spindown rate recovered over a few months to their previous state [272].

6.2. Glitches in high magnetic field pulsars and magnetars

The connection between glitches and emission changes becomes clearer for neutron stars with high inferred magnetic field intensities. The component of the surface magnetic field perpendicular to the pulsar's rotational axis can be calculated by assuming that the observed spin-down is entirely due to magnetic dipole radiation (equation (20)). Possibly better approximations are available, see for example [273], but pulsar inclination angles are largely unknown. Hence equation (20) is broadly used to get a rough estimate of the magnetic dipole field intensity B at the pulsar's surface, which is easily done by measuring ν and $\dot{\nu}$ via pulsar timing. MSPs are the pulsars with the weakest inferred B known, with values that accumulate just above 10^8 G but can reach up to $\sim 10^{10}$ G. For the common pulsar population, estimated B values range between 10^{10} and 10^{13} G, whilst magnetars exhibit the strongest fields known, with intensities over 10^{13}–10^{14} G in general.

Magnetars undergo outbursts in which their X-ray and soft γ-ray luminosity can exceed the available rotational power. Supported by the inferred high values for the surface dipole component of their magnetic field, the common stance is that the energy required for the observed activity comes from their magnetic energy reservoirs. Whilst the total energy release during short X-ray bursts and the most rare bright flares depends on certain assumptions, for example about how isotropic the emission is, typical estimates place the energy outputs at 10^{40}–10^{41} erg and up to 10^{45}–10^{46} erg respectively. On the other hand, the rotational energy reservoir will be

$$E_{\rm rot} = \frac{1}{2} I \Omega^2 \approx 2 \times 10^{44} \left(\frac{I}{10^{45}\,{\rm g\,cm^2}} \right) \left(\frac{\nu}{0.1\,{\rm Hz}} \right)^2 {\rm erg} \quad (50)$$

whilst the energy due to the magnetic field

$$E_{\rm mag} = \int \frac{B^2}{8\pi} {\rm d}V \approx 3 \times 10^{45} \left(\frac{\bar{B}}{10^{14}\,{\rm G}} \right)^2 \left(\frac{R_\star}{12\,{\rm km}} \right)^3 {\rm erg} \quad (51)$$

where $\bar{B}$ is the volume-averaged magnetic field. Comparing the two expressions above, we see that the main free energy reservoir for magnetar activity must be their magnetic field rather than their rotation.

Most magnetars were discovered through serendipitous detection of very bright¶, millisecond-scale X-ray flares, which are abundant during outburst episodes [274, 275]. In many cases the enhanced emission and flaring X-ray activity will decay slowly after a few weeks or months, resulting in luminosities that can be up to two orders of magnitude weaker after 1-2 yr. The quiescent X-ray luminosity of these magnetars is often too low and would require very long integration times for regular observations to be carried out. Only a handful of sources are bright enough during quiescence to be systematically monitored. In one of the most valuable studies to date concerning magnetar rotational evolution, 16 years of observations from five magnetars were investigated [98]. It was found that magnetars can be remarkably unstable rotators, and they can enter periods of enhanced levels of timing noise which are sometimes related to the X-ray outbursts. Furthermore, all five sources displayed glitches, some of which are in close temporal proximity to X-ray bursting activity.

The association of some glitches with X-ray flares is unequivocal (e.g. [98]). Nonetheless, while it appears that most X-ray emission changes – like pulse profile variations, flux enhancement, and bursts – can be associated with some sort of timing anomaly (either glitches or spin-down rate changes), not all glitches are accompanied by enhanced emission; and there are some magnetars for which no radiative changes have been observed in proximity to any of their glitches [276]. It is unclear what the causal relationship between the two phenomena is, but a possibility is that a crustal failure event can be the connection between the two. Besides spin-down driven crustquakes like those discussed for regular pulsars, a magnetar's crust is expected to fail under the substantial magnetic stresses that build up as a result of the evolution of their strong magnetic fields [277]. During such an event, vortex unpinning can occur either directly because of the crustal motion [278], or indirectly as the result of heating (thermally exciting vortices over their pinning barrier, see e.g. [239]). Thus crustal failure could both trigger a glitch and generate the magnetospheric conditions required to produce bursts (e.g. by twisting of magnetic field lines which are anchored to the crust [279]). The aforementioned study of five magnetars concluded that about 20-30% of glitches and timing anomalies can be associated with radiative changes ([98], also [280]), and that the size (or any other property) of the glitches is irrelevant with respect to the presence or lack of a radiative counterpart. However, later studies using an extended dataset, found a possible tendency for events related to X-ray outbursts to involve larger frequency changes [280]. No differences between the timing events that are associated with emission changes and those which are not has been found.

When the glitch activity of magnetars is compared to that of the pulsar population (as defined in equation

¶ 10^{34}-10^{35} erg s^{-1} in the 2-10 keV band

(23)), it falls precisely at the level expected by their spin-down rates [158, 157]. For normal pulsars, glitch activity does not seem to depend on the inferred surface magnetic field [157]. Glitch sizes $\Delta\nu$ of magnetars are not exceptional [280], but in line with expectations for their $\dot{\nu}$ values, and it is only when considering $\Delta\nu/\nu$ values that magnetar glitches appear particularly large because of their low spin frequencies. The modest $\Delta\nu$ of magnetars can also be explained by the fact that they are young neutron stars and the decay of their high magnetic fields provides a powerful heating source, potentially preventing large lags ω from building up. This is because, in hotter stars, the large thermal energy lowers the effective pinning barrier vortices must overcome in order to unpin. Therefore $\dot{\Omega}_{\rm s}$ can be greater than in a cooler object. The post-glitch recovery can involve large $\Delta\dot{\nu}$ changes which can be understood within the standard picture since re-coupling timescales (equation (32)) scale inversely with spin frequency and magnetars are slow rotators [39]. Nonetheless, at least for those glitches that coincide with outbursts, the approximation of a constant external torque is no longer valid and the contribution of a changing magnetosphere must be accounted for when examining the post-glitch evolution. The effect of the evolving external torque on the spin-down can be non-trivial and last for timescales comparable to those typical of post-glitch recoveries (e.g. [281]).

The rotation of magnetars has historically been characterised by long periods $P = 1/\nu$, in the range 2-10 s, and very high period derivatives $\dot{P} > 10^{-12}$, which are rather extreme in comparison to the rotational parameters of other pulsars and imply very high dipole magnetic field strengths as already mentioned [282]. The few magnetars that have smaller period derivatives, hence low inferred B intensity, in the range of standard pulsar values, have been explained in the magnetar model as powered by intense toroidal magnetic field components that would reside under the surface [283, 284]. Interestingly, the two rotation-powered pulsars with the highest period derivatives ($\dot{P} = 4 - 7 \times 10^{-12}$), and with some of the highest magnetic fields ($B = 4.1 - 4.9 \times 10^{13}$ Gauss), displayed large glitches accompanied by magnetar-like high energy outbursts after years of normal pulsar behaviour. Both are very young ($\tau_{\rm ch} < 2$ kyr), like some magnetars, but have periods ~ 0.3 s, which is shorter than the periods of all magnetars. The glitches in these high-B pulsars have some unique characteristics, as is described next.

In May 2006 the Rossi X-ray Timing Explorer (*RXTE*) detected four short bursts (< 0.1 s) coming from the direction of PSR J1846−0258 [285]. This radio-quiet, X-ray rotation-powered pulsar had been discovered in 2000 [286] and since then never exhibited such behaviour. Analysis of Chandra observations showed that the X-ray spectrum became softer and its (unabsorbed) flux increased by a factor of ~ 6 [287]. The pulsed X-ray flux, as recorded by *RXTE*, suddenly increased and slowly decayed to normal values in approximately two months [285, 287, 288]. The radiative outburst remarkably resembled magnetar behaviour. Moreover, together with the unusual radiative activity, there was a large glitch ($\Delta\nu = 12\,\mu$Hz [289]) and the rotation became less stable after the event [288, 289]. This was the first time that a rotation-powered pulsar showed magnetar-like activity, thereby connecting the magnetar phenomenology – including their glitches – with that of high magnetic field pulsars.

Whilst the 2006 PSR J1846−0258 glitch had a size common among Vela-like pulsars, it is unusual for young pulsars ($\tau_{\rm ch} \lesssim 5$ kyr) to exhibit such large glitches (e.g. [27]). However, the main peculiarity of that glitch was that it involved a very large recovery fraction $Q = 9$. The exponential recovery time-scale was found to be $\tau_{\rm d} = 127$ d, and a large persisting shift in spindown rate meant that the net effect of the glitch (after the transient recovery) was a very large reduction of ν ($\Delta\nu_{\rm p} = -95\,\mu$Hz) [289]. Only magnetars had shown such peculiar recoveries at the time. Additionally, once the rotation became more stable and the emission was back to normal, the braking index $n = \nu\ddot{\nu}/\dot{\nu}^2$ (see Section 7) which for this pulsar had been clearly established at $n = 2.65$, was measured to have a lower value $n = 2.19$ (a 17% reduction [290, 291]). Braking index changes are commonly observed when n is calculated using inter-glitch $\ddot{\nu}_{\rm ig}$ values in Vela-like pulsars or in PSR J0537−6910 [85, 108], but are not very common among very young pulsars such as the Crab or J0540−6919 [80, 99]. A change of n for PSR J1846−0258 was not observed after a previous, smaller, glitch [292]. More recently, after 14 years of calm following the first detected burst, the source was again seen to enter a period of activity in which the pulsed fraction increased, the spin-down rate was enhanced and a new glitch of similar size took place [293].

PSR J1119−6127 is a pulsar very similar to J1846−0258 in terms of rotational parameters (hence their inferred magnetic fields are very similar too), including a well determined braking index $n = 2.684$ [113]. However, this pulsar is not only detected in X-rays but also in the radio, and it has been monitored for more than two decades with the Parkes radio telescope. The radio pulse profile is normally stable and single peaked, but it became double peaked in one out of 332 analysed observations, spanning ~ 12 yr [113]. It was also shown that data from the next observing session (at the same observing frequency

of 1400 MHz) contained bright and short bursts of emission which were present only in some of the pulsar rotations (hence not necessarily visible in the average pulse profile). These observations happened to be the first ones available after a large ($\Delta\nu = 13\,\mu$Hz) glitch that occurred in 2007, and became the first evidence for radio radiative changes associated with a glitch in a rotation powered pulsar. No irregularities in the emission were detected in the available observations that followed a (smaller) glitch back in 2004.

Nonetheless, both the 2004 and 2007 glitches exhibited similar anomalous recoveries which can be seen in the last subplot of figure 6. After those glitches, the rotation converged (quasi-exponentially) to spin-down rates considerably smaller than the predictions of the pre-glitch models. In other words, these glitches had $\Delta\dot{\nu}_{\rm p} > 0$, which is contrary to what is observed in other pulsars (including PSR J1846−0258). The only other pulsar that has shown a glitch with $\Delta\dot{\nu}_{\rm p} > 0$ is the Rotating Radio Transient (RRAT, [294, 159]) PSR J1819−1458, which is also a high-B pulsar ($B = 5 \times 10^{13}$ G). RRATs are neutron stars characterised by only emitting sporadic, bright and short bursts of radio waves that track the rotation of the star; just like the transient emission observed in PSR J1119−6127 after its 2007 glitch. Despite the overshoot of $\dot{\nu}$, PSR J1119−6127 had also $\Delta\nu_{\rm p} < 0$, thus $Q \gtrsim 1$, similarly to the 2006 glitch in PSR J1846−0258. Moreover, the braking index after the 2007 glitch changed too. A 15% reduction was measured by [108] when using a model like equation (15) with one exponential term. But the authors showed that the data was equally well fitted by a model without a persistent change $\Delta\ddot{\nu}_{\rm p}$ and with two exponential terms. In this last fit the second exponential term had a negative $\Delta\nu_{\rm d_2} < 0$ and a rather long characteristic timescale $\tau_2 \simeq 2300$ d, which would imply eventually a return of the braking index to pre-glitch values.

While the anomalous glitch recoveries can be attributed to superfluid dynamics ([108, 295]), the temporal association of the 2007 glitch with radio emission changes, as well as the high inferred B value for this pulsar, led to the suggestion that magnetospheric effects could be contributing to the timing behaviour, analogous to what is seen in magnetars [108]. Unfortunately observations at high energies were not available around the 2007 glitch to confirm the existence of magnetar-like emission, but eventually an X-ray outburst was detected in 2016, classifying the pulsar as a transient magnetar. The short, bright, X-ray bursts were associated with a flux increase by a factor of nearly 200, spectral hardening, and a large glitch ($\Delta\nu = 14\,\mu$Hz) accompanied by a spin-down change comparable to the one seen in 2007 [163]. Interestingly, the radio emission flux decreased to undetectable levels at the time of the outburst [296]. By the time the radio emission was restored, about 10 days later [297], it exhibited a different, double peaked profile at 1400 MHz [298]. Changes were also reported at 2.3 and 8.3 GHz observing frequencies [299]). The glitch exhibited a peculiar recovery which did not resemble the ones observed after the 2004 and 2007 glitches. This time the spin-down rate took considerable time to increase (unlike in previous glitches, where $\Delta\dot{\nu}$ appeared instantaneous) and there was not an obvious overshoot in $\dot{\nu}$ (see figure in [300]). It was claimed that the recovery could not be fitted by an exponential function as in equation (15) [301] and the whole event was compared to the activity of the magnetar XTE J1810−197 [300].

The magnetospheric conditions and activity of magnetars can have a strong impact on the spin-down torque. For example, the crust is expected to shear under magnetic stresses [302, 303], twisting the magnetic field lines; this is found to result in an enhanced spindown [304] and can cause short-lived, large $\Delta\dot{\nu} < 0$ changes which can appear as anti-glitches [279], like those described in Section 3.6. Such effects must be accounted for together with the superfluid response when trying to explain the recoveries of glitches in magnetars and high magnetic field pulsars. Besides the mechanical motion of a failing crust, heating during magnetar activity could also lead to vortex unpinning, triggering a glitch [238] and altering the response of the superfluid. To fully understand magnetar glitches most likely will require advanced magneto-thermoplastic simulations (see e.g. [305] for a 1-D model), eventually including the neutron component and coupling to the rotational evolution equations. Crucial to this progress will be input from timing of magnetars, and especially of high-magnetic field radio pulsars whose rotation can be followed closely.

7. Impact of glitches on pulsar long-term rotational evolution

Glitches introduce lasting changes in the rotation of pulsars and their accumulated effect can affect the longer-term rotational evolution of neutron stars.

The persisting fractional changes $\Delta\nu_{\rm p}/\nu$ of the frequency are generally small, but as discussed in Section 4.2 a series of glitches over time $T_{\rm obs}$ can reverse $\sim 1\%$ of the secular spin-down in most pulsars, with the possible exception of some very young neutron stars with high spin-down rates, such as the Crab pulsar and PSR J0540−6919, for which this amounts to less than 0.01% [157].

Recovering the properties of the underlying spin-down mechanism is critical to understanding how

pulsars operate but can be difficult in the presence of glitches. The post-glitch relaxation(s) can dominate the spin evolution for periods comparable to or exceeding the duration of the observing span – which is often years or even decades, and it is unclear when the process is completed so that measured parameters (ν, $\dot{\nu}$, $\ddot{\nu}$) cease to contain glitch-induced contributions.

A parameter often used to characterise the rotational energy losses that cause the secular spin-down $\dot{\nu}$ and dictate the long-term evolution is the braking index n. This assumes a form $\dot{\Omega}_{\rm c} \propto -\Omega_{\rm c}^{n}$ and can therefore be directly calculated from observations as $n = \nu\ddot{\nu}/\dot{\nu}^2$ if the proportionality factor is to a good approximation constant over the time period of observations. Indicatively, expectations for the value of the braking index are $n = 3$ for magnetic-dipole braking (equation (20)), $n = 5$ for gravitational wave emission from the mass quadrupole mode, $n = 7$ for gravitational wave emission from an r-mode, and $n = 1$ if losses due to a pulsar wind dominate. A changing magnetic moment or/and a varying effective (coupled) moment of inertia will also affect the value of the braking index [306, 248].

Whenever $\ddot{\nu}$ has been quantified for large samples of pulsars, its absolute value is generally orders of magnitude larger than the expected $|\ddot{\nu}_{n=3}|$ for magnetic dipole braking (equation (20)) and approximately half of the measured values are negative [81]. Such values are not reflective of the underlying braking mechanism and are considered as artifacts of timing noise and potentially of glitching activity [251]. The scatter and magnitude of inferred braking indices greatly increases for pulsar characteristic ages $\tau_{\rm ch} \gtrsim 10^6$ yr [81] and a positive correlation between $|n|$ and $\tau_{\rm ch}$ has been reported to hold for pulsars with $\tau_{\rm ch} \gtrsim 10^5$ yr [307]. Whilst for these older objects similar numbers of $n > 0$ and $n < 0$ are observed, this is not true for younger ($\lesssim 10^5$ yr) neutron stars. For this group, the aforementioned correlation is not observed, values of n are mostly positive, and their modulus is smaller – though usually still $n > 3$ [307, 84]. Additionally, in very young pulsars ($\tau_{\rm ch} < 10^4\,$yr), a long-term $\ddot{\nu}$ is more likely to be steady and robust enough to be detected over the effects of timing noise, making such n measurements more reliable to represent the rotational long-term evolution. As these younger neutron stars are the same population for which glitches are most frequent (Section 4.1), it is important to understand the exact impact of glitches on inferred n values.

The post-glitch recovery is typically associated with large values of the braking index. During the initial post-glitch phase, when there is usually rapid, quasi-exponential, evolution of $\dot{\nu}$, it is not meaningful to think in terms of n. In pulsars which show the linear $\dot{\nu}$-recovery, $\dot{\nu}(t)$ has a steep positive slope which does not seem to flatten before the next glitch arrives (see Section 5.4 and the example of the Vela pulsar in figure 20). This phase is characterised by large, positive, inter-glitch braking indices $n_{\rm ig}$. For example $n_{\rm ig} \simeq 40$ for the Vela pulsar and $n_{\rm ig} \simeq 25$ for PSR J1803-2137 [248].

In known glitching pulsars like the above, $n_{\rm ig}$ cannot be confused with its long-term value. It has been speculated though that high ($n > 3$) measured braking indices of other pulsars can also be the result of a glitch, one that went undetected, for example if it occurred before the start of observations. In a sample of 19 pulsars where the favoured timing model (which incorporated a timing noise component) included a $\ddot{\nu}$ term, over half of the sources had inferred n values greater than 10 [84]. The possibility of a long-$\tau_{\rm d}$ (> 1000 days) exponential signature from an earlier glitch was explored, but found no support in the data. The effects of the quasi-linear recovery with high $\ddot{\nu}_{\rm ig}$ cannot, however, be distinguished from a high underlying, intrinsic, $\ddot{\nu}$. Hence the view that some of these $n > 10$ values are the result of unobserved glitches remains tenable. The correlation of glitch rate with spin-down rate (Section 4.1) can be used to inform on the likelihood of the presence of a glitch within a certain time interval before observations began.

Another impact of glitches comes from the cumulative effect of the persistent $\dot{\nu}$ changes, that can have fractional amplitudes $\Delta\dot{\nu}_{\rm p}/\dot{\nu}$ as high as $\sim 10^{-3}$. Overall they will decrease the inferred $\ddot{\nu}$ when this is obtained directly from timing measurements of long (glitch-containing) intervals, without the use of corrections related to glitch models and possible $\Delta\dot{\nu}_{\rm p}$ shifts. Indeed, it has been found that pulsars with regular large glitches (Vela-like) evolve with particularly low long-term braking indices $n_{\rm eff} \lesssim 2$ [248]. These are the effective braking indices for the specific pulsars: whilst they do not necessarily correspond to the value of n for the underlying braking mechanism, they nonetheless represent the rotational evolution of these pulsars during their recorded history (usually $\lesssim 30\,$yr) [248, 85]. The Vela pulsar and PSR J0537−6910 are clear examples of this behaviour, as their long-term trends have been studied over periods of times that involve more than 10 glitches. In the particular case of PSR J0537−6910, the inferred $n_{\rm eff}$ is not just under 3 but actually negative [193, 100, 101], which has been ascribed to its extraordinarily frequent glitches. Interestingly, other glitching pulsars which have glitch activities as expected (that is, as high as $\gtrsim 1\%$ of $\dot{\nu}$), despite exhibiting mostly small glitches, also appear to evolve with low long-term $\ddot{\nu}$ values that imply $n_{\rm eff} \lesssim 2$ [248, 191]. Even for the Crab pulsar, which does not follow the glitch activity trend, a period of increased glitch activity resulted in a decrease of

its inferred braking index, from 2.5 to $\sim$ 2.3 [80], thereby strengthening the argument for the rotational long-term effect of glitches (and in particular, $\Delta\dot{\nu}_{\rm p}$). For some neutron stars that are infrequently observed, undetected glitches that leave $\Delta\dot{\nu}_{\rm p} < 0$ changes can mimic an $n < 0$ trend, perhaps explaining some of the anomalous n measurements [251].

Glitches not only directly impact the rotational state of neutron stars, but can affect other aspects of their evolution such as the thermal one and that of the magnetic field. Dissipation during the glitch will result in heating, as will crust failure if that is what initiates some of the events, leading to a potentially different signature in the thermal spectrum of a star following a glitch, depending on the trigger [239]. Furthermore, heating due to mutual friction dissipation as vortices migrate outwards may explain why many millisecond radio pulsars are hotter than their age would predict [308, 309]. More speculatively, pinning of vortices to fluxtubes can tie the magnetic evolution to the rotational one [60, 61], leading to magnetic flux-expulsion and decay of the magnetic field as vortices move out.

8. Concluding remarks

Years before neutron stars were discovered, it was expected on theoretical grounds that these objects should host superfluid components. It must, however, have seemed unlikely that there would ever be a way to probe the detailed interior dynamics and structure of a kind of star which was minuscule, dark, and distant. Glitches provide us with this opportunity.

The last decade has seen many observational advances related to glitches, such as a great increase in the overall number of reported events (from searches both in archival and new data, figure 1) which led to a better understanding of their properties [157, 191, 68], detailed analyses of the rotational history of selected glitching pulsars [80, 100, 92], and the discovery of a possible emission disturbance associated with a Vela glitch [91], to name a few. The development of new algorithms which aim to identify and characterise glitches (Section 3.4), and the close monitoring of hundreds of sources by current and upcoming facilities, is expected to generate larger, consistently measured glitch samples (whose completeness can been assessed) for even more pulsars. The possible association of emission variability with glitches (or glitch precursors) in regular pulsars [91], as well as its link to timing noise [150], offers another promising area of future research.

The considerable progress and research output based on the nearly continuous (and of high temporal resolution) timing of the Crab and Vela pulsars underlines the importance of dedicated timing campaigns, which are designed to probe glitch aspects that are not accessible via regular timing programs. The latter are essential to uncover the prevalence of glitches in the pulsar population, the dependence of glitch rate and activity on pulsar parameters, and to increase the glitch sample for statistical studies. But only very close monitoring can provide information on a range of valuable features, such as the small glitches' activity, subtle short-lived magnetospheric changes related to glitches, and a rigorous tracking of glitch rises and recoveries on all relevant timescales.

Further gains may also come from gravitational wave observations, as glitch induced asymmetries can couple to the gravitational field and lead to a possibly observable signal [310, 311, 312, 313, 314, 315, 316, 317]. Searches have been carried out in LIGO and Virgo data following glitches in the Vela and Crab [318, 319], and in other pulsars that glitched during the third observational run of the detectors [320]. No signals have been detected to date, but future searches will continue to target glitching pulsars and, as the sensitivity of the instruments improves, have the potential to provide new input on the glitch mechanism.

At the same time, on the theoretical side, hydrodynamical models based on a multifluid formalism [219, 321] are becoming more established. Aided by calculations and simulations that aim to better understand pinning and unpinning of vortices [322, 50, 57], such models can provide a solid framework to interpret phenomenology and connect it naturally to the physical parameters governing the interior dynamics. Generating the appropriate microphysical inputs is a challenging task: while progress has been made on understanding the dynamics of individual vortices in a pinning potential, accurate modelling requires 3D simulations of a large numbers of vortices in the pinning landscape of the inner crust, and possibly in the presence of magnetic fluxtubes in the core. A way forward is likely to come from the development of sub-grid prescriptions to build a bridge between microscopic vortex dynamics simulations and large scale hydrodynamics. This will allow for the inclusion of fluid effects in metamodels, and help constrain the glitch trigger mechanism from observations. Comparison with terrestrial superfluid experiments always remains instructive, and can guide progress in less explored, but likely important, directions, like the role of superfluid turbulence. A further step for most models discussed here will be relaxing the assumption of rigid body rotation of the normal component, as Ekman pumping (e.g. at the crust-core boundary) could play a role in post-glitch recoveries [323].

The observational and theoretical developments described above have already improved our under-

standing of the glitch process in recent years. The principal postulation of an internal origin for glitches holds up against new observations and more advanced simulations and microphysical input, as does the notion that the response of the neutron superfluid interior is responsible for the post-glitch recoveries on long timescales.

In order to explain the inferred spin-down rate changes and relaxation on timescales of several days or longer, a small fraction of the superfluid (order 0.1% of the total moment of inertia) must decouple from the observed crust at a glitch. A much greater fraction might decouple initially, including part of the superfluid core, as evidenced by high time resolution observations of fast and large spin-ups in the Vela pulsar followed by a rapid decay in the first minutes post-glitch.

A superfluid angular momentum reservoir that amounts to $\sim$ 1–2% of the moment of inertia of the 'normal' component that has spun-up is needed to drive the glitch itself. Calculations of the maximum pinning force between vortices and the ion lattice suggest that it is strong enough to sustain the required velocity differences between components. However, the entrainment effect in the crust might actually couple the superfluid neutrons to the normal component strongly enough that a large lag cannot develop between glitches. Furthermore the latest data challenge the hypothesis that glitches are a purely crustal phenomenon. Observations of individual glitches in the Vela and Crab pulsar strongly suggest that at least part of the core superfluid must be involved, possibly indicating that vortices are pinned to fluxtubes in this region.

Implicated in glitches is the collective unpinning of a very large number of vortices, the exact cause of which remains unknown. Nonetheless, statistical glitch studies over the whole pulsar population can help discern its nature. Whilst some pulsars exhibit regularity in their glitch behaviour (having a characteristic glitch size and similar interglitch waiting times), and at least in one instance a strong correlation between glitch size and time until the next event is seen - implying some threshold-regulated process, several other pulsars show properties of a scale invariant process behind the glitch trigger. The latter suggests that vortex-vortex interactions can lead to vortex avalanches: the rapid release of pinned vortices which move together and transport angular momentum to the crust.

There have been proof of principle studies that show how glitch observations can be used to constrain the mass and equation of state of neutron stars, and determine whether quantum turbulence is being excited in the pinned superfluid. Nevertheless, much of the detailed glitch physics remains work in progress, and more research is necessary to make fully quantitative predictions.

With the wealth of already available data and the prospect of more from existing, as well as future, observing programs, it is essential for theoretical and observational efforts to move forward in tandem; each side is able to provide insights that inform the development of the other and will allow us to build a clearer picture of the physics of the neutron stars' extreme conditions.

Acknowledgments

We would like to thank Simon Johnston, Richard Dodson, and the members of the pulsar group at Jodrell Bank Centre for Astrophysics, for permission to use pulsar times of arrivals to produce several figures included in this work, as well as Patrick Weltevrede and Sam Lander for providing feedback on an earlier version of this manuscript. D.A. acknowledges support from an EPSRC/STFC fellowship (EP/T017325/1). B.H. acknowledges support from NCN Poland via SONATA BIS grant 2015/18/E/ST9/00577. C.M.E. acknowledges support from ANID FONDECYT 1211964.

[1] Landau L D 1932 *Phys. Zs. Sowjet* **1** 285
[2] Baade W and Zwicky F 1934 *Physical Review* **46** 76–77
[3] Baade W and Zwicky F 1934 *Proceedings of the National Academy of Science* **20** 254–259
[4] Hewish A, Bell S J, Pilkington J D H, Scott P F and Collins R A 1968 Nature **217** 709–713
[5] VERITAS Collaboration, Aliu E, Arlen T, Aune T, Beilicke M, Benbow W, Bouvier A, Bradbury S M, Buckley J H, Bugaev V, Byrum K, Cannon A, Cesarini A, Christiansen J L, Ciupik L, Collins-Hughes E, Connolly M P, Cui W, Dickherber R, Duke C, Errando M, Falcone A, Finley J P, Finnegan G, Fortson L, Furniss A, Galante N, Gall D, Gibbs K, Gillanders G H, Godambe S, Griffin S, Grube J, Guenette R, Gyuk G, Hanna D, Holder J, Huan H, Hughes G, Hui C M, Humensky T B, Imran A, Kaaret P, Karlsson N, Kertzman M, Kieda D, Krawczynski H, Krennrich F, Lang M J, Lyutikov M, Madhavan A S, Maier G, Majumdar P, McArthur S, McCann A, McCutcheon M, Moriarty P, Mukherjee R, Nuez P, Ong R A, Orr M, Otte A N, Park N, Perkins J S, Pizlo F, Pohl M, Prokoph H, Quinn J, Ragan K, Reyes L C, Reynolds P T, Roache E, Rose J, Ruppel J, Saxon D B, Schroedter M, Sembroski G H, entrk G D, Smith A W, Staszak D, Tei G, Theiling M, Thibadeau S, Tsurusaki K, Tyler J, Varlotta A, Vassiliev V V, Vincent S, Vivier M, Wakely S P, Ward J E, Weekes T C, Weinstein A, Weisgarber T, Williams D A and Zitzer B 2011 *Science* **334** 69–72 (*Preprint* `1108.3797`)
[6] Kaspi V M, Taylor J H and Ryba M F 1994 ApJ **428** 713
[7] Gold T 1968 Nature **218** 731–732
[8] Cole T W 1969 Nature **221** 29–31
[9] Davies J G, Hunt G C and Smith F G 1969 Nature **221** 27–29
[10] Reichley P E and Downs G S 1969 Nature **222** 229–230

[11] Radhakrishnan V and Manchester R N 1969 Nature **222** 228–229
[12] Boynton P E, Groth E J I, Partridge R B and Wilkinson D T 1969 IAU Circ. **2179** 1
[13] Durney B 1969 Nature **222** 1260–1261
[14] Börner G and Cohen J M 1971 *Nature Physical Science* **231** 146–147
[15] Scargle J D and Pacini F 1971 *Nature Physical Science* **232** 144–149
[16] Greenstein G S and Cameron A G W 1969 *nat* **222** 862
[17] Ruderman M A 1968 Nature **218** 1128–1129
[18] Ruderman M 1969 Nature **223** 597–598
[19] Migdal A B 1959 *Nuclear Physics* **13** 655–674
[20] Ginzburg V L and Kirzhnits D A 1965 *Journal of Experimental and Theoretical Physics* **20** 1346
[21] Baym G, Pethick C, Pines D and Ruderman M 1969 Nature **224** 872–874
[22] Greenstein G 1970 Nature **227** 791–794
[23] Kim Y B, Hempstead C F and Strnad A R 1963 *Phys. Rev.* **131**(6) 2486–2495 URL https://link.aps.org/doi/10.1103/PhysRev.131.2486
[24] Beasley M R, Labusch R and Webb W W 1969 *Phys. Rev.* **181**(2) 682–700 URL https://link.aps.org/doi/10.1103/PhysRev.181.682
[25] Anderson P W and Itoh N 1975 Nature **256** 25–27
[26] Haskell B and Melatos A 2015 *International Journal of Modern Physics D* **24** 1530008
[27] Espinoza C M, Lyne A G, Stappers B W and Kramer M 2011 *MNRAS* **414** 1679–1704 (*Preprint* 1102.1743)
[28] Yu M, Manchester R N, Hobbs G, Johnston S, Kaspi V M, Keith M, Lyne A G, Qiao G J, Ravi V, Sarkissian J M, Shannon R and Xu R X 2013 *MNRAS* **429** 688–724
[29] Cordes J M, Downs G S and Krause-Polstorff J 1988 ApJ **330** 847
[30] Wang N, Manchester R N, Pace R T, Bailes M, Kaspi V M, Stappers B W and Lyne A G 2000 MNRAS **317** 843–860
[31] Yuan J P, Wang N, Manchester R N and Liu Z Y 2010 MNRAS **404** 289–304
[32] Chamel N and Haensel P 2008 *Living Reviews in Relativity* **11** 10 (*Preprint* 0812.3955)
[33] Anderson P W 1962 *Phys. Rev. Lett.* **9**(7) 309–311 URL https://link.aps.org/doi/10.1103/PhysRevLett.9.309
[34] Anderson P W and Kim Y B 1964 *Reviews of Modern Physics* **36** 39–42
[35] Hall H E and Vinen W F 1956 *Proceedings of the Royal Society of London Series A* **238** 215–234
[36] Andersson N, Sidery T and Comer G L 2006 MNRAS **368** 162–170 (*Preprint* astro-ph/0510057)
[37] Campbell L J 1979 Phys. Rev. Lett. **43** 1336–1339
[38] Jahan-Miri M 2006 ApJ **650** 326–331
[39] Haskell B and Antonopoulou D 2014 MNRAS **438** L16–L20 (*Preprint* 1306.5214)
[40] Andersson N, Sidery T and Comer G L 2007 MNRAS **381** 747–756 (*Preprint* astro-ph/0703257)
[41] Sidery T, Passamonti A and Andersson N 2010 MNRAS **405** 1061–1074 (*Preprint* 0910.3918)
[42] Jones P B 1992 MNRAS **257** 501–506
[43] Haskell B and Sedrakian A 2018 *Superfluidity and Superconductivity in Neutron Stars* vol 457 p 401
[44] Jones P B 1991 ApJ **373** 208
[45] Seveso S, Pizzochero P M, Grill F and Haskell B 2016 MNRAS **455** 3952–3967
[46] Wlazłowski G, Sekizawa K, Magierski P, Bulgac A and Forbes M M 2016 Phys. Rev. Lett. **117** 232701 (*Preprint* 1606.04847)
[47] Avogadro P, Barranco F, Broglia R A and Vigezzi E 2007 Phys. Rev. C **75** 012805 (*Preprint* nucl-th/0602028)
[48] Avogadro P, Barranco F, Broglia R A and Vigezzi E 2008 Nucl. Phys. A **811** 378–412 (*Preprint* 0804.1765)
[49] Howitt G, Melatos A and Haskell B 2020 MNRAS **498** 320–331 (*Preprint* 2008.00365)
[50] Antonelli M and Haskell B 2020 MNRAS **499** 3690–3705 (*Preprint* 2007.11720)
[51] Drummond L V and Melatos A 2017 MNRAS **472** 4851–4869 (*Preprint* 1709.02254)
[52] Drummond L V and Melatos A 2018 MNRAS **475** 910–920 (*Preprint* 1712.02938)
[53] Alpar M A 2017 *Journal of Astrophysics and Astronomy* **38** 44 (*Preprint* 1709.07912)
[54] Gügercinoğlu E and Alpar M A 2014 ApJ **788** L11
[55] Gügercinoğlu E and Alpar M A 2016 MNRAS **462** 1453–1460 (*Preprint* 1607.05092)
[56] Gusakov M E 2019 MNRAS **485** 4936–4950 (*Preprint* 1904.01363)
[57] Sourie A and Chamel N 2020 MNRAS **493** 382–389 (*Preprint* 2001.08964)
[58] Sourie A and Chamel N 2020 MNRAS **493** L98–L102 (*Preprint* 2001.09668)
[59] Sourie A and Chamel N 2021 MNRAS **503** 1407–1417 (*Preprint* 2102.11054)
[60] Ruderman M, Zhu T and Chen K 1998 ApJ **492** 267–280 (*Preprint* astro-ph/9709008)
[61] Glampedakis K, Andersson N and Samuelsson L 2011 MNRAS **410** 805–829 (*Preprint* 1001.4046)
[62] Sedrakian A D and Sedrakian D M 1995 ApJ **447** 305
[63] Espinoza C M, Antonopoulou D, Stappers B W, Watts A and Lyne A G 2014 MNRAS **440** 2755–2762 (*Preprint* 1402.7219)
[64] Rencoret J A, Aguilera-Gómez C and Reisenegger A 2021 A&A **654** A47 (*Preprint* 2106.12604)
[65] Warszawski L and Melatos A 2011 MNRAS **415** 1611–1630 (*Preprint* 1103.6090)
[66] Warszawski L and Melatos A 2013 MNRAS **428** 1911–1926 (*Preprint* 1210.2203)
[67] Melatos A, Peralta C and Wyithe J S B 2008 ApJ **672** 1103–1118 (*Preprint* 0710.1021)
[68] Howitt G, Melatos A and Delaigle A 2018 ApJ **867** 60 (*Preprint* 1809.06038)
[69] Glampedakis K and Andersson N 2009 Phys. Rev. Lett. **102** 141101
[70] Khomenko V, Antonelli M and Haskell B 2019 Phys. Rev. D **100** 123002 (*Preprint* 1906.05177)
[71] Peralta C, Melatos A, Giacobello M and Ooi A 2006 ApJ **651** 1079–1091 (*Preprint* astro-ph/0607161)
[72] Melatos A and Peralta C 2007 ApJ **662** L99–L102 (*Preprint* 0710.2455)
[73] Haskell B, Antonopoulou D and Barenghi C 2020 MNRAS **499** 161–170 (*Preprint* 2007.02748)
[74] Manchester R N and Taylor J H 1977 *Pulsars*
[75] Lorimer D R and Kramer M 2004 *Handbook of Pulsar Astronomy* vol 4
[76] Manchester R N, Lyne A G, Camilo F, Bell J F, Kaspi V M, D'Amico N, McKay N P F, Crawford F, Stairs I H, Possenti A, Kramer M and Sheppard D C 2001 MNRAS **328** 17–35 (*Preprint* astro-ph/0106522)
[77] Keith M J, Jameson A, van Straten W, Bailes M, Johnston S, Kramer M, Possenti A, Bates S D, Bhat N D R, Burgay M, Burke-Spolaor S, D'Amico N, Levin L, McMahon P L, Milia S and Stappers B W 2010 MNRAS **409** 619–627 (*Preprint* 1006.5744)
[78] Ray P S, Kerr M, Parent D, Abdo A A, Guillemot L, Ransom S M, Rea N, Wolff M T, Makeev A, Roberts M S E, Camilo F, Dormody M, Freire P C C, Grove J E, Gwon C, Harding A K, Johnston S, Keith M, Kramer M, Michelson P F, Romani R W, Saz Parkinson P M, Thompson D J, Weltevrede P, Wood K S and Ziegler M 2011 ApJS **194** 17 (*Preprint* 1011.2468)
[79] Deneva J S, Ray P S, Lommen A, Ransom S M, Bogdanov

S, Kerr M, Wood K S, Arzoumanian Z, Black K, Doty J, Gendreau K C, Guillot S, Harding A, Lewandowska N, Malacaria C, Markwardt C B, Price S, Winternitz L, Wolff M T, Guillemot L, Cognard I, Baker P T, Blumer H, Brook P R, Cromartie H T, Demorest P B, DeCesar M E, Dolch T, Ellis J A, Ferdman R D, Ferrara E C, Fonseca E, Garver-Daniels N, Gentile P A, Jones M L, Lam M T, Lorimer D R, Lynch R S, McLaughlin M A, Ng C, Nice D J, Pennucci T T, Spiewak R, Stairs I H, Stovall K, Swiggum J K, Vigeland S J and Zhu W W 2019 ApJ **874** 160 (*Preprint* 1902.07130)

[80] Lyne A G, Jordan C A, Graham-Smith F, Espinoza C M, Stappers B W and Weltevrede P 2015 MNRAS **446** 857–864 (*Preprint* 1410.0886)

[81] Hobbs G, Lyne A G, Kramer M, Martin C E and Jordan C 2004 MNRAS **353** 1311–1344

[82] Hobbs G, Lyne A G and Kramer M 2010 MNRAS **402** 1027–1048 (*Preprint* 0912.4537)

[83] Parthasarathy A, Shannon R M, Johnston S, Lentati L, Bailes M, Dai S, Kerr M, Manchester R N, Osłowski S, Sobey C, van Straten W and Weltevrede P 2019 MNRAS **489** 3810–3826 (*Preprint* 1908.11709)

[84] Parthasarathy A, Johnston S, Shannon R M, Lentati L, Bailes M, Dai S, Kerr M, Manchester R N, Osłowski S, Sobey C, van Straten W and Weltevrede P 2020 MNRAS **494** 2012–2026 (*Preprint* 2003.13303)

[85] Lower M E, Johnston S, Dunn L, Shannon R M, Bailes M, Dai S, Kerr M, Manchester R N, Melatos A, Oswald L S, Parthasarathy A, Sobey C and Weltevrede P 2021 MNRAS (*Preprint* 2109.07612)

[86] Wang N, Manchester R N, Zhang J, Wu X J, Yusup A, Lyne A G, Cheng K S and Chen M Z 2001 MNRAS **328** 855–866 (*Preprint* astro-ph/0111006)

[87] Hu Y, Li L, Yuan J P, Dang S J, Wang S Q, Wang Z J and Yuen R 2020 Ap&SS **365** 143 (*Preprint* 2011.02641)

[88] Shabanova T V, Pugachev V D and Lapaev K A 2013 ApJ **775** 2 (*Preprint* 1307.0297)

[89] Starovoit E D and Suleymanova S A 2019 *Astronomy Reports* **63** 310–315

[90] Dodson R, Lewis D and McCulloch P 2007 Ap&SS **308** 585–589 (*Preprint* astro-ph/0612371)

[91] Palfreyman J, Dickey J M, Hotan A, Ellingsen S and van Straten W 2018 Nature **556** 219–222

[92] Espinoza C M, Antonopoulou D, Dodson R, Stepanova M and Scherer A 2021 A&A **647** A25 (*Preprint* 2007.02921)

[93] Basu A, Joshi B C, Krishnakumar M A, Bhattacharya D, Nandi R, Bandhopadhay D, Char P and Manoharan P K 2020 MNRAS **491** 3182–3191 (*Preprint* 1911.04934)

[94] Jankowski F, Bailes M, van Straten W, Keane E F, Flynn C, Barr E D, Bateman T, Bhandari S, Caleb M, Campbell-Wilson D, Farah W, Green A J, Hunstead R W, Jameson A, Osłowski S, Parthasarathy A, Rosado P A and Venkatraman Krishnan V 2019 MNRAS **484** 3691–3712 (*Preprint* 1812.04038)

[95] Johnston S, Karastergiou A, Keith M J, Song X, Weltevrede P, Abbate F, Bailes M, Buchner S, Camilo F, Geyer M, Hugo B, Jameson A, Kramer M, Parthasarathy A, Reardon D J, Ridolfi A, Serylak M, Shannon R M, Spiewak R, van Straten W, Venkatraman Krishnan V, Jankowski F, Meyers B W, Oswald L, Posselt B, Sobey C, Szary A and van Leeuwen J 2020 MNRAS **493** 3608–3615 (*Preprint* 2002.10250)

[96] Ng C 2018 Pulsar science with the CHIME telescope *Pulsar Astrophysics the Next Fifty Years* vol 337 ed Weltevrede P, Perera B B P, Preston L L and Sanidas S pp 179–182 (*Preprint* 1711.02104)

[97] CHIME/Pulsar Collaboration, Amiri M, Bandura K M, Boyle P J, Brar C, Cliche J F, Crowter K, Cubranic D, Demorest P B, Denman N T, Dobbs M, Dong F Q, Fandino M, Fonseca E, Good D C, Halpern M, Hill A S, Höfer C, Kaspi V M, Landecker T L, Leung C, Lin H H, Luo J, Masui K W, McKee J W, Mena-Parra J, Meyers B W, Michilli D, Naidu A, Newburgh L, Ng C, Patel C, Pinsonneault-Marotte T, Ransom S M, Renard A, Scholz P, Shaw J R, Sikora A E, Stairs I H, Tan C M, Tendulkar S P, Tretyakov I, Vanderlinde K, Wang H and Wang X 2021 ApJS **255** 5 (*Preprint* 2008.05681)

[98] Dib R and Kaspi V M 2014 ApJ **784** 37 (*Preprint* 1401.3085)

[99] Ferdman R D, Archibald R F and Kaspi V M 2015 ApJ **812** 95 (*Preprint* 1506.00182)

[100] Antonopoulou D, Espinoza C M, Kuiper L and Andersson N 2018 MNRAS **473** 1644–1655 (*Preprint* 1708.09459)

[101] Ho W C G, Espinoza C M, Arzoumanian Z, Enoto T, Tamba T, Antonopoulou D, Bejger M, Guillot S, Haskell B and Ray P S 2020 MNRAS **498** 4605–4614 (*Preprint* 2009.00030)

[102] Gügercinoğlu E, Ge M Y, Yuan J P and Zhou S Q 2020 *arXiv e-prints* arXiv:2011.14788 (*Preprint* 2011.14788)

[103] Namkham N, Jaroenjittichai P and Johnston S 2019 MNRAS **487** 5854–5861 (*Preprint* 1906.06992)

[104] Lower M E, Bailes M, Shannon R M, Johnston S, Flynn C, Osłowski S, Gupta V, Farah W, Bateman T, Green A J, Hunstead R, Jameson A, Jankowski F, Parthasarathy A, Price D C, Sutherland A, Temby D and Venkatraman Krishnan V 2020 MNRAS **494** 228–245 (*Preprint* 2002.12481)

[105] Shannon R M and Cordes J M 2010 ApJ **725** 1607–1619 (*Preprint* 1010.4794)

[106] Arzoumanian Z, Brazier A, Burke-Spolaor S, Chamberlin S, Chatterjee S, Christy B, Cordes J M, Cornish N J, Crawford F, Thankful Cromartie H, Crowter K, DeCesar M E, Demorest P B, Dolch T, Ellis J A, Ferdman R D, Ferrara E C, Fonseca E, Garver-Daniels N, Gentile P A, Halmrast D, Huerta E A, Jenet F A, Jessup C, Jones G, Jones M L, Kaplan D L, Lam M T, Lazio T J W, Levin L, Lommen A, Lorimer D R, Luo J, Lynch R S, Madison D, Matthews A M, McLaughlin M A, McWilliams S T, Mingarelli C, Ng C, Nice D J, Pennucci T T, Ransom S M, Ray P S, Siemens X, Simon J, Spiewak R, Stairs I H, Stinebring D R, Stovall K, Swiggum J K, Taylor S R, Vallisneri M, van Haasteren R, Vigeland S J, Zhu W and NANOGrav Collaboration 2018 ApJS **235** 37 (*Preprint* 1801.01837)

[107] Shannon R M, Lentati L T, Kerr M, Johnston S, Hobbs G and Manchester R N 2016 MNRAS **459** 3104–3111 (*Preprint* 1604.02661)

[108] Antonopoulou D, Weltevrede P, Espinoza C M, Watts A L, Johnston S, Shannon R M and Kerr M 2015 MNRAS **447** 3924–3935 (*Preprint* 1412.5853)

[109] Dodson R G, McCulloch P M and Lewis D R 2002 ApJ **564** L85–L88 (*Preprint* astro-ph/0201005)

[110] Ashton G, Lasky P D, Graber V and Palfreyman J 2019 *Nature Astronomy* **3** 1143–1148

[111] Shaw B, Keith M J, Lyne A G, Mickaliger M B, Stappers B W, Turner J D and Weltevrede P 2021 MNRAS **505** L6–L10 (*Preprint* 2103.13180)

[112] Wong T, Backer D C and Lyne A G 2001 ApJ **548** 447–459 (*Preprint* astro-ph/0010010)

[113] Weltevrede P, Johnston S and Espinoza C M 2011 MNRAS **411** 1917–1934 (*Preprint* 1010.0857)

[114] Janssen G H and Stappers B W 2006 A&A **457** 611–618 (*Preprint* astro-ph/0607260)

[115] Melatos A, Dunn L M, Suvorova S, Moran W and Evans R J 2020 *apj* **896** 78

[116] Dunn L, Lower M E and Melatos A 2021 MNRAS **504** 3399–3411 (*Preprint* 2104.07363)

[117] Cognard I and Backer D C 2004 ApJ **612** L125–L127

(*Preprint* astro-ph/0407546)
[118] Singha J, Basu A, Krishnakumar M A, Joshi B C and Arumugam P 2021 *arXiv e-prints* arXiv:2106.01942 (*Preprint* 2106.01942)
[119] Sokolova E V and Panin A G 2020 *arXiv e-prints* arXiv:2006.15956 (*Preprint* 2006.15956)
[120] Atwood W B, Abdo A A, Ackermann M, Althouse W, Anderson B, Axelsson M, Baldini L, Ballet J, Band D L, Barbiellini G, Bartelt J, Bastieri D, Baughman B M, Bechtol K, Bédérède D, Bellardi F, Bellazzini R, Berenji B, Bignami G F, Bisello D, Bissaldi E, Blandford R D, Bloom E D, Bogart J R, Bonamente E, Bonnell J, Borgland A W, Bouvier A, Bregeon J, Brez A, Brigida M, Bruel P, Burnett T H, Busetto G, Caliandro G A, Cameron R A, Caraveo P A, Carius S, Carlson P, Casandjian J M, Cavazzuti E, Ceccanti M, Cecchi C, Charles E, Chekhtman A, Cheung C C, Chiang J, Chipaux R, Cillis A N, Ciprini S, Claus R, Cohen-Tanugi J, Condamoor S, Conrad J, Corbet R, Corucci L, Costamante L, Cutini S, Davis D S, Decotigny D, DeKlotz M, Dermer C D, de Angelis A, Digel S W, do Couto e Silva E, Drell P S, Dubois R, Dumora D, Edmonds Y, Fabiani D, Farnier C, Favuzzi C, Flath D L, Fleury P, Focke W B, Funk S, Fusco P, Gargano F, Gasparrini D, Gehrels N, Gentit F X, Germani S, Giebels B, Giglietto N, Giommi P, Giordano F, Glanzman T, Godfrey G, Grenier I A, Grondin M H, Grove J E, Guillemot L, Guiriec S, Haller G, Harding A K, Hart P A, Hays E, Healey S E, Hirayama M, Hjalmarsdotter L, Horn R, Hughes R E, Jóhannesson G, Johansson G, Johnson A S, Johnson R P, Johnson T J, Johnson W N, Kamae T, Katagiri H, Kataoka J, Kavelaars A, Kawai N, Kelly H, Kerr M, Klamra W, Knödlseder J, Kocian M L, Komin N, Kuehn F, Kuss M, Landriu D, Latronico L, Lee B, Lee S H, Lemoine-Goumard M, Lionetto A M, Longo F, Loparco F, Lott B, Lovellette M N, Lubrano P, Madejski G M, Makeev A, Marangelli B, Massai M M, Mazziotta M N, McEnery J E, Menon N, Meurer C, Michelson P F, Minuti M, Mirizzi N, Mitthumsiri W, Mizuno T, Moiseev A A, Monte C, Monzani M E, Moretti E, Morselli A, Moskalenko I V, Murgia S, Nakamori T, Nishino S, Nolan P L, Norris J P, Nuss E, Ohno M, Ohsugi T, Omodei N, Orlando E, Ormes J F, Paccagnella A, Paneque D, Panetta J H, Parent D, Pearce M, Pepe M, Perazzo A, Pesce-Rollins M, Picozza P, Pieri L, Pinchera M, Piron F, Porter T A, Poupard L, Rainò S, Rando R, Rapposelli E, Razzano M, Reimer A, Reimer O, Reposeur T, Reyes L C, Ritz S, Rochester L S, Rodriguez A Y, Romani R W, Roth M, Russell J J, Ryde F, Sabatini S, Sadrozinski H F W, Sanchez D, Sander A, Sapozhnikov L, Parkinson P M S, Scargle J D, Schalk T L, Scolieri G, Sgrò C, Share G H, Shaw M, Shimokawabe T, Shrader C, Sierpowska-Bartosik A, Siskind E J, Smith D A, Smith P D, Spandre G, Spinelli P, Starck J L, Stephens T E, Strickman M S, Strong A W, Suson D J, Tajima H, Takahashi H, Takahashi T, Tanaka T, Tenze A, Tether S, Thayer J B, Thayer J G, Thompson D J, Tibaldo L, Tibolla O, Torres D F, Tosti G, Tramacere A, Turri M, Usher T L, Vilchez N, Vitale V, Wang P, Watters K, Winer B L, Wood K S, Ylinen T and Ziegler M 2009 ApJ **697** 1071–1102 (*Preprint* 0902.1089)
[121] Zou W Z, Wang N, Manchester R N, Urama J O, Hobbs G, Liu Z Y and Yuan J P 2008 MNRAS **384** 1063–1068 (*Preprint* 0711.1196)
[122] Hobbs G B, Edwards R T and Manchester R N 2006 MNRAS **369** 655–672 (*Preprint* astro-ph/0603381)
[123] Ellis J A, Vallisneri M, Taylor S R and Baker P T 2020 ENTERPRISE: Enhanced Numerical Toolbox Enabling a Robust PulsaR Inference SuitE URL https://zenodo.org/record/4059815
[124] Lentati L, Alexander P, Hobson M P, Feroz F, van Haasteren R, Lee K J and Shannon R M 2014 MNRAS **437** 3004–3023 (*Preprint* 1310.2120)
[125] Lentati L, Alexander P and Hobson M P 2015 MNRAS **447** 2159–2168 (*Preprint* 1412.1427)
[126] Shemar S L and Lyne A G 1996 MNRAS **282** 677–690
[127] Buchner S and Flanagan C 2008 Glitches in the Vela Pulsar *40 Years of Pulsars: Millisecond Pulsars, Magnetars and More* (*American Institute of Physics Conference Series* vol 983) ed Bassa C, Wang Z, Cumming A and Kaspi V M pp 145–147
[128] Xie Y and Zhang S N 2013 ApJ **778** 31 (*Preprint* 1309.2537)
[129] Içdem B, Baykal A and Inam S Ç 2012 MNRAS **419** 3109–3114
[130] Archibald R F, Kaspi V M, Ng C Y, Gourgouliatos K N, Tsang D, Scholz P, Beardmore A P, Gehrels N and Kennea J A 2013 Nature **497** 591–593 (*Preprint* 1305.6894)
[131] Şaşmaz Muş S, Aydın B and Göğüş E 2014 MNRAS **440** 2916–2921 (*Preprint* 1411.5274)
[132] Hu Y M, Pitkin M, Heng I S and Hendry M A 2014 ApJ **784** L41 (*Preprint* 1311.2955)
[133] Younes G, Ray P S, Baring M G, Kouveliotou C, Fletcher C, Wadiasingh Z, Harding A K and Goldstein A 2020 ApJ **896** L42 (*Preprint* 2006.04854)
[134] Klochkov D, Staubert R, Postnov K, Shakura N and Santangelo A 2009 A&A **506** 1261–1267 (*Preprint* 0908.0053)
[135] Ray P S, Guillot S, Ho W C G, Kerr M, Enoto T, Gendreau K C, Arzoumanian Z, Altamirano D, Bogdanov S, Campion R, Chakrabarty D, Deneva J S, Jaisawal G K, Kozon R, Malacaria C, Strohmayer T E and Wolff M T 2019 ApJ **879** 130 (*Preprint* 1811.09218)
[136] Woods P M, Kouveliotou C, van Paradijs J, Finger M H, Thompson C, Duncan R C, Hurley K, Strohmayer T, Swank J and Murakami T 1999 ApJ **524** L55–L58 (*Preprint* astro-ph/9907173)
[137] Gavriil F P, Dib R and Kaspi V M 2011 ApJ **736** 138 (*Preprint* 0905.1256)
[138] Garcia F and Ranea-Sandoval I F 2015 MNRAS **449** L73–L76 (*Preprint* 1402.0848)
[139] Mastrano A, Suvorov A G and Melatos A 2015 MNRAS **453** 522–530 (*Preprint* 1507.07308)
[140] Kantor E M and Gusakov M E 2014 ApJ **797** L4 (*Preprint* 1411.2777)
[141] Ducci L, Pizzochero P M, Doroshenko V, Santangelo A, Mereghetti S and Ferrigno C 2015 A&A **578** A52 (*Preprint* 1505.00548)
[142] Stark M J, Heffner C M, Baykal A and Swank J H 2004 Pulsar frequency glitches in the Accreting pulsar SAX J2103.5+4545 *American Astronomical Society Meeting Abstracts* (*American Astronomical Society Meeting Abstracts* vol 205) p 83.03
[143] Galloway D K, Morgan E H and Levine A M 2004 ApJ **613** 1164–1172 (*Preprint* astro-ph/0401476)
[144] Serim M M, Şahiner Ş, ćerri-Serim D, Inam S ć and Baykal A 2017 MNRAS **471** 4982–4989 (*Preprint* 1706.00727)
[145] Howitt G and Melatos A 2022 MNRAS **514** 863–874 (*Preprint* 2205.05896)
[146] Shabanova T V 2010 ApJ **721** 251–258 (*Preprint* 1007.0125)
[147] Zhou S Q, Zhou A A, Zhang J, Liu M Q, Liu H Y, Zhang L, Feng Z W, Zhu X D and Wu D 2019 Ap&SS **364** 173
[148] Shabanova T V 2007 Ap&SS **308** 591–593 (*Preprint* astro-ph/0607248)
[149] Shabanova T V 2009 ApJ **700** 1009–1016 (*Preprint* 0906.3641)

[150] Lyne A, Hobbs G, Kramer M, Stairs I and Stappers B 2010 *Science* **329** 408 (*Preprint* 1006.5184)
[151] Shabanova T V 2009 *Astronomy Reports* **53** 465–471
[152] McKenna J and Lyne A G 1990 Nature **343** 349–350
[153] Lyne A G, Shemar S L and Smith F G 2000 *MNRAS* **315** 534–542
[154] Basu A, Shaw B, Antonopoulou D, Keith M J, Lyne A G, Mickaliger M B, Stappers B W, Weltevrede P and Jordan C A 2022 MNRAS **510** 4049–4062 (*Preprint* 2111.06835)
[155] McKee J W, Janssen G H, Stappers B W, Lyne A G, Caballero R N, Lentati L, Desvignes G, Jessner A, Jordan C A, Karuppusamy R, Kramer M, Cognard I, Champion D J, Graikou E, Lazarus P, Osłowski S, Perrodin D, Shaifullah G, Tiburzi C and Verbiest J P W 2016 MNRAS **461** 2809–2817
[156] Perera B B P, DeCesar M E, Demorest P B, Kerr M, Lentati L, Nice D J, Osłowski S, Ransom S M, Keith M J, Arzoumanian Z, Bailes M, Baker P T, Bassa C G, Bhat N D R, Brazier A, Burgay M, Burke-Spolaor S, Caballero R N, Champion D J, Chatterjee S, Chen S, Cognard I, Cordes J M, Crowter K, Dai S, Desvignes G, Dolch T, Ferdman R D, Ferrara E C, Fonseca E, Goldstein J M, Graikou E, Guillemot L, Hazboun J S, Hobbs G, Hu H, Islo K, Janssen G H, Karuppusamy R, Kramer M, Lam M T, Lee K J, Liu K, Luo J, Lyne A G, Manchester R N, McKee J W, McLaughlin M A, Mingarelli C M F, Parthasarathy A P, Pennucci T T, Perrodin D, Possenti A, Reardon D J, Russell C J, Sanidas S A, Sesana A, Shaifullah G, Shannon R M, Siemens X, Simon J, Spiewak R, Stairs I H, Stappers B W, Swiggum J K, Taylor S R, Theureau G, Tiburzi C, Vallisneri M, Vecchio A, Wang J B, Zhang S B, Zhang L, Zhu W W and Zhu X J 2019 MNRAS **490** 4666–4687 (*Preprint* 1909.04534)
[157] Fuentes J R, Espinoza C M, Reisenegger A, Shaw B, Stappers B W and Lyne A G 2017 *A&A* **608** A131
[158] Dib R, Kaspi V M and Gavriil F P 2008 ApJ **673** 1044–1061
[159] Lyne A G, McLaughlin M A, Keane E F, Kramer M, Espinoza C M, Stappers B W, Palliyaguru N T and Miller J 2009 MNRAS **400** 1439–1444 (*Preprint* 0909.1165)
[160] Gotthelf E V and Halpern J P 2020 ApJ **900** 159 (*Preprint* 2007.10431theSAO/NASAAstrophysicsDataSystem)
[161] Viganò D and Pons J A 2012 MNRAS **425** 2487–2492 (*Preprint* 1206.2014)
[162] Gotthelf E V, Halpern J P and Alford J 2013 ApJ **765** 58 (*Preprint* 1301.2717)
[163] Archibald R F, Kaspi V M, Tendulkar S P and Scholz P 2016 ApJ **829** L21 (*Preprint* 1608.01007)
[164] Carlin J B, Melatos A and Vukcevic D 2019 MNRAS **482** 3736–3743 (*Preprint* 1810.12458)
[165] Ashton G, Prix R and Jones D I 2017 Phys. Rev. D **96** 063004 (*Preprint* 1704.00742)
[166] Alpar M A, Chau H F, Cheng K S and Pines D 1993 ApJ **409** 345
[167] Datta B and Alpar M A 1993 A&A **275** 210–212
[168] Link B, Epstein R I and Lattimer J M 1999 Phys. Rev. Lett. **83** 3362–3365
[169] Ho W C G, Espinoza C M, Antonopoulou D and Andersson N 2015 *Science Advances* **1** e1500578–e1500578 (*Preprint* 1510.00395)
[170] Carter B, Chamel N and Haensel P 2005 Nucl. Phys. A **748** 675–697 (*Preprint* nucl-th/0402057)
[171] Chamel N 2012 Phys. Rev. C **85** 035801
[172] Chamel N 2013 Phys. Rev. Lett. **110** 011101
[173] Andersson N, Glampedakis K, Ho W C G and Espinoza C M 2012 Phys. Rev. Lett. **109** 241103 (*Preprint* 1207.0633)
[174] Hooker J, Newton W G and Li B A 2015 MNRAS **449** 3559–3567 (*Preprint* 1308.0031)
[175] Piekarewicz J, Fattoyev F J and Horowitz C J 2014 Phys. Rev. C **90** 015803 (*Preprint* 1404.2660)
[176] Steiner A W, Gandolfi S, Fattoyev F J and Newton W G 2015 Phys. Rev. C **91** 015804 (*Preprint* 1403.7546)
[177] Watanabe G and Pethick C J 2017 Phys. Rev. Lett. **119** 062701
[178] Martin N and Urban M 2016 Phys. Rev. C **94** 065801
[179] Alpar M A, Langer S A and Sauls J A 1984 ApJ **282** 533–541
[180] Leinson L B 2020 MNRAS **498** 304–309 (*Preprint* 2008.06328)
[181] Glampedakis K and Lasky P D 2015 MNRAS **450** 1638–1650 (*Preprint* 1501.05473)
[182] Newton W G, Berger S and Haskell B 2015 MNRAS **454** 4400–4410 (*Preprint* 1506.01445)
[183] Lander S K 2014 MNRAS **437** 424–436 (*Preprint* 1307.7020)
[184] Sur A and Haskell B 2021 PASA **38** e043
[185] Gügercinoğlu E and Alpar M A 2014 ApJ **788** L11
[186] Montoli A, Antonelli M, Magistrelli F and Pizzochero P M 2020 A&A **642** A223
[187] Schwenk A, Friman B and Brown G E 2003 Nucl. Phys. A **713** 191–216 (*Preprint* nucl-th/0207004)
[188] Pizzochero P M 2011 ApJ **743** L20
[189] Pizzochero P M, Antonelli M, Haskell B and Seveso S 2017 *Nature Astronomy* **1** 0134
[190] Antonelli M, Montoli A and Pizzochero P M 2018 MNRAS **475** 5403–5416
[191] Fuentes J R, Espinoza C M and Reisenegger A 2019 A&A **630** A115 (*Preprint* 1907.09887)
[192] Shaw B, Lyne A G, Stappers B W, Weltevrede P, Bassa C G, Lien A Y, Mickaliger M B, Breton R P, Jordan C A, Keith M J and Krimm H A 2018 MNRAS **478** 3832–3840 (*Preprint* 1805.05110)
[193] Middleditch J, Marshall F E, Gotthelf E V, Wang Q D and Zhang W 2006 Predicting The Glitches In PSR J0537-6910 *American Astronomical Society Meeting Abstracts #208* (*American Astronomical Society Meeting Abstracts* vol 208) p 4.01
[194] Abbott R *et al* (LIGO Scientific Collaboration, Virgo Collaboration, KAGRA Collaboration), Antonopoulou D, Arzoumanian Z, Enoto T, Espinoza C M and Guillot S 2021 ApJ **913** L27 (*Preprint* 2012.12926)
[195] Celora T, Khomenko V, Antonelli M and Haskell B 2020 MNRAS **496** 5564–5574 (*Preprint* 2002.04310)
[196] Warszawski L and Melatos A 2008 MNRAS **390** 175–191 (*Preprint* 0806.4738)
[197] Melatos A and Warszawski L 2009 ApJ **700** 1524–1540 (*Preprint* 0904.2998)
[198] Fulgenzi W, Melatos A and Hughes B D 2017 MNRAS **470** 4307–4329
[199] Carlin J B and Melatos A 2020 MNRAS **494** 3383–3391 (*Preprint* 2004.00168)
[200] Carlin J B and Melatos A 2019 MNRAS **483** 4742–4750 (*Preprint* 1901.01633)
[201] Carlin J B and Melatos A 2019 MNRAS **488** 4890–4896
[202] Carlin J B and Melatos A 2021 *arXiv e-prints* arXiv:2105.13588 (*Preprint* 2105.13588)
[203] Melatos A and Drummond L V 2019 ApJ **885** 37 (*Preprint* 1910.05503)
[204] Haskell B 2016 MNRAS **461** L77–L81
[205] Easson I 1979 ApJ **228** 257–267
[206] Easson I 1979 ApJ **233** 711–716
[207] Feibelman P J 1971 Phys. Rev. D **4** 1589–1597
[208] Sauls J A, Stein D L and Serene J W 1982 Phys. Rev. D **25** 967–975
[209] Sedrakian A 1998 Phys. Rev. D **58** 021301 (*Preprint* astro-ph/9806156)

[210] Link B 2003 Phys. Rev. Lett. **91** 101101 (*Preprint* astro-ph/0302441)
[211] Haskell B, Glampedakis K and Andersson N 2014 MNRAS **441** 1662–1668 (*Preprint* 1307.0985)
[212] Ho W C G, Andersson N and Haskell B 2011 Phys. Rev. Lett. **107** 101101 (*Preprint* 1107.5064)
[213] Haskell B, Degenaar N and Ho W C G 2012 MNRAS **424** 93–103 (*Preprint* 1201.2101)
[214] Sedrakian A 2005 Phys. Rev. D **71** 083003 (*Preprint* astro-ph/0408467)
[215] Bildsten L and Epstein R I 1989 ApJ **342** 951
[216] Jones P B 1990 MNRAS **243** 257–262
[217] Epstein R I and Baym G 1992 ApJ **387** 276
[218] Andersson N, Sidery T and Comer G L 2006 MNRAS **368** 162–170 (*Preprint* astro-ph/0510057)
[219] Andersson N and Comer G L 2021 *Living Reviews in Relativity* **24** 3 (*Preprint* 2008.12069)
[220] Rothen F 1981 A&A **98** 36–38
[221] Sourie A, Chamel N, Novak J and Oertel M 2017 MNRAS **464** 4641–4657 (*Preprint* 1607.08213)
[222] Gavassino L, Antonelli M, Pizzochero P M and Haskell B 2020 MNRAS **494** 3562–3580 (*Preprint* 2001.08951)
[223] Breu C and Rezzolla L 2016 MNRAS **459** 646–656 (*Preprint* 1601.06083)
[224] Alpar M A, Pines D, Anderson P W and Shaham J 1984 ApJ **276** 325–334
[225] Link B 2014 ApJ **789** 141 (*Preprint* 1311.2499)
[226] Khomenko V and Haskell B 2018 PASA **35** e020 (*Preprint* 1801.01413)
[227] Gorter C J and Mellink J H 1949 *Physica* **15** 285–304
[228] Andersson N, Sidery T and Comer G L 2007 MNRAS **381** 747–756 (*Preprint* astro-ph/0703257)
[229] Large M I, Vaughan A E and Mills B Y 1968 Nature **220** 340–341
[230] Helfand D J, Gotthelf E V and Halpern J P 2001 ApJ **556** 380–391 (*Preprint* astro-ph/0007310)
[231] Haskell B, Pizzochero P M and Sidery T 2012 MNRAS **420** 658–671 (*Preprint* 1107.5295)
[232] Lyne A G, Smith F G and Pritchard R S 1992 Nature **359** 706–707
[233] Ge M Y, Zhang S N, Lu F J, Li T P, Yuan J P, Zheng X P, Huang Y, Zheng S J, Chen Y P, Chang Z, Tuo Y L, Cheng Q, Güngör C, Song L M, Xu Y P, Cao X L, Chen Y, Liu C Z, Zhang S, Qu J L, Bu Q C, Cai C, Chen G, Chen L, Chen M Z, Chen T X, Chen Y B, Cui W, Cui W W, Deng J K, Dong Y W, Du Y Y, Fu M X, Gao G H, Gao H, Gao M, Gu Y D, Guan J, Guo C C, Han D W, Hao L F, Huo J, Jia S M, Jiang L H, Jiang W C, Jin C J, Jin J, Jin Y J, Kong L D, Li B, Li D, Li C K, Li G, Li M S, Li W, Li X, Li X B, Li X F, Li Y G, Li Z W, Li Z X, Liu Z Y, Liang X H, Liao J Y, Liu G Q, Liu H W, Liu X J, Liu Y N, Lu B, Lu X F, Luo Q, Luo T, Ma X, Meng B, Nang Y, Nie J Y, Ou G, Sai N, Shang R C, Song X Y, Sun L, Tan Y, Tao L, Wang C, Wang G F, Wang J, Wang J B, Wang M, Wang N, Wang W S, Wang Y D, Wang Y S, Wen X Y, Wen Z G, Wu B B, Wu B Y, Wu M, Xiao G C, Xiao S, Xiong S L, Xu Y H, Yan W M, Yang J W, Yang S, Yang Y J, Yang Y J, Yi Q B, Yin Q Q, You Y, Yue Y L, Zhang A M, Zhang C M, Zhang D P, Zhang F, Zhang H M, Zhang J, Zhang T, Zhang W C, Zhang W, Zhang W Z, Zhang Y, Zhang Y F, Zhang Y J, Zhang Y, Zhang Z, Zhang Z, Zhang Z L, Zhao H S, Zhao X F, Zheng W, Zhou D K, Zhou J F, Zhou X, Zhuang R L, Zhu Y X and Zhu Y 2020 ApJ **896** 55 (*Preprint* 2004.00791)
[234] Akbal O and Alpar M A 2018 MNRAS **473** 621–624 (*Preprint* 1707.01160)
[235] Chamel N and Carter B 2006 MNRAS **368** 796–808 (*Preprint* astro-ph/0503044)
[236] Gügercinoğlu E and Alpar M A 2019 MNRAS **488** 2275–2282 (*Preprint* 1905.08180)
[237] Horowitz C J and Kadau K 2009 Phys. Rev. Lett. **102** 191102 (*Preprint* 0904.1986)
[238] Link B and Epstein R I 1996 ApJ **457** 844 (*Preprint* astro-ph/9508021)
[239] Larson M B and Link B 2002 MNRAS **333** 613–622 (*Preprint* astro-ph/0105461)
[240] Haskell B, Khomenko V, Antonelli M and Antonopoulou D 2018 MNRAS **481** L146–L150 (*Preprint* 1806.10168)
[241] Andersson N, Antonopoulou D, Espinoza C M, Haskell B and Ho W C G 2018 ApJ **864** 137 (*Preprint* 1711.05550)
[242] Graber V, Cumming A and Andersson N 2018 ApJ **865** 23 (*Preprint* 1804.02706)
[243] Sauls J A, Chamel N and Alpar M A 2020 *arXiv e-prints* arXiv:2001.09959 (*Preprint* 2001.09959)
[244] Douchin F and Haensel P 2001 A&A **380** 151–167 (*Preprint* astro-ph/0111092)
[245] Gügercinoğlu E and Alpar M A 2014 ApJ **788** L11
[246] Xu Y H, Yuan J P, Lee K J, Hao L F, Wang N, Wang M, Yu M, Li Z X, Yue Y L, Liu Z Y, Yuen R, Wen Z G, Dang S J, Bai J M, Chen W, Huang Y X and Yan W M 2019 Ap&SS **364** 11
[247] Alpar M A, Anderson P W, Pines D and Shaham J 1984 ApJ **278** 791–805
[248] Espinoza C M, Lyne A G and Stappers B W 2017 MNRAS **466** 147–162 (*Preprint* 1611.08314)
[249] Ferdman R D, Archibald R F, Gourgouliatos K N and Kaspi V M 2018 ApJ **852** 123 (*Preprint* 1708.08876)
[250] Akbal O, Alpar M A, Buchner S and Pines D 2017 MNRAS **469** 4183–4192 (*Preprint* 1612.03805)
[251] Alpar M A and Baykal A 2006 MNRAS **372** 489–496 (*Preprint* astro-ph/0608055)
[252] Alpar M A, Chau H F, Cheng K S and Pines D 1996 ApJ **459** 706
[253] Link B and Epstein R I 1997 ApJ **478** L91–L94 (*Preprint* astro-ph/9701096)
[254] Krishnamohan S and Downs G S 1983 *Astrophysical Journal* **265** 372–388
[255] Johnston S, van Straten W, Kramer M and Bailes M 2001 ApJ **549** L101–L104
[256] Kramer M, Johnston S and van Straten W 2002 *MNRAS* **334** 523–532
[257] Palfreyman J L, Hotan A W, Dickey J M, Young T G and Hotan C E 2011 ApJ **735** L17
[258] Chen J L, Wen Z G, Hao L F, Yuan J P, Li J, Wang H G, Yan W M, Lee K J, Wang N, Xu Y H, Li Z X, Huang Y X, Yuen R and Mijit M 2020 ApJ **899** 118
[259] Lommen A, Donovan J, Gwinn C, Arzoumanian Z, Harding A, Strickman M, Dodson R, McCulloch P and Moffett D 2007 ApJ **657** 436–440
[260] Krishnamohan S and Downs G S 1983 ApJ **265** 372–388
[261] Buchner S J 2010 *The Astronomer's Telegram* **2768** 1
[262] Palfreyman J L, Dickey J M, Ellingsen S P, Jones I R and Hotan A W 2016 ApJ **820** 64
[263] Vivekanand M 2015 ApJ **806** 190
[264] Zhang X, Shuai P, Huang L, Chen S and Du Y 2018 ApJ **866** 82
[265] Vivekanand M 2020 A&A **633** A57
[266] Feng H, Li H, Long X, Bellazzini R, Costa E, Wu Q, Huang J, Jiang W, Minuti M, Wang W, Xu R, Yang D, Baldini L, Citraro S, Nasimi H, Soffitta P, Muleri F, Jung A, Yu J, Jin G, Zeng M, An P, Brez A, Latronico L, Sgro C, Spandre G and Pinchera M 2020 *Nature Astronomy* **4** 511–516
[267] Long X, Feng H, Li H, Zhu J, Wu Q, Huang J, Minuti M, Jiang W, Wang W, Xu R, Costa E, Yang D, Citraro S, Nasimi H, Yu J, Jin G, Zeng M, An P, Baldini L, Bellazzini R, Brez A, Latronico L, Sgrò C, Spandre G, Pinchera M, Muleri F and Soffitta P 2021 ApJ **912** L28

(*Preprint* 2104.11391)
[268] Lyne A, Graham-Smith F, Weltevrede P, Jordan C, Stappers B, Bassa C and Kramer M 2013 *Science* **342** 598–601 (*Preprint* 1311.0408)
[269] Keith M J, Shannon R M and Johnston S 2013 MNRAS **432** 3080–3084 (*Preprint* 1304.4644)
[270] Kou F F, Yuan J P, Wang N, Yan W M and Dang S J 2018 MNRAS **478** L24–L28 (*Preprint* 1801.01248)
[271] Allafort A, Baldini L, Ballet J, Barbiellini G, Baring M G, Bastieri D, Bellazzini R, Bonamente E, Bottacini E, Brandt T J, Bregeon J, Bruel P, Buehler R, Buson S, Caliandro G A, Cameron R A, Caraveo P A, Cecchi C, Chaves R C G, Chekhtman A, Chiang J, Chiaro G, Ciprini S, Claus R, D'Ammando F, de Palma F, Digel S W, Di Venere L, Drell P S, Favuzzi C, Ferrara E C, Franckowiak A, Fusco P, Gargano F, Gasparrini D, Giglietto N, Giroletti M, Glanzman T, Godfrey G, Grenier I A, Guiriec S, Hadasch D, Harding A K, Hayashida M, Hayashi K, Hays E, Hewitt J, Hill A B, Horan D, Hou X, Jogler T, Johnson A S, Johnson T J, Kerr M, Knödlseder J, Kuss M, Lande J, Larsson S, Latronico L, Lemoine-Goumard M, Longo F, Loparco F, Lubrano P, Malyshev D, Marelli M, Mayer M, Mazziotta M N, Mehault J, Mizuno T, Monzani M E, Morselli A, Murgia S, Nemmen R, Nuss E, Ohsugi T, Omodei N, Orienti M, Orlando E, Paneque D, Pesce-Rollins M, Pierbattista M, Piron F, Pivato G, Porter T A, Rainò S, Rando R, Ray P S, Razzano M, Reimer O, Reposeur T, Romani R W, Sartori A, Saz Parkinson P M, Sgrò C, Siskind E J, Smith D A, Spinelli P, Strong A W, Takahashi H, Thayer J B, Thompson D J, Tibaldo L, Tinivella M, Torres D F, Tosti G, Uchiyama Y, Usher T L, Vandenbroucke J, Vasileiou V, Venter C, Vianello G, Vitale V, Winer B L and Wood K S 2013 ApJ **777** L2 (*Preprint* 1308.0358)
[272] Zhao J, Ng C W, Lin L C C, Takata J, Cai Y, Hu C P, Yen D C C, Tam P H T, Hui C Y, Kong A K H and Cheng K S 2017 ApJ **842** 53 (*Preprint* 1706.00236)
[273] Spitkovsky A 2006 ApJ **648** L51–L54 (*Preprint* astro-ph/0603147)
[274] Mereghetti S, Pons J A and Melatos A 2015 Space Sci. Rev. **191** 315–338 (*Preprint* 1503.06313)
[275] Kaspi V M and Beloborodov A M 2017 ARA&A **55** 261–301 (*Preprint* 1703.00068)
[276] Scholz P, Archibald R F, Kaspi V M, Ng C Y, Beardmore A P, Gehrels N and Kennea J A 2014 ApJ **783** 99 (*Preprint* 1401.5000)
[277] Lander S K, Andersson N, Antonopoulou D and Watts A L 2015 MNRAS **449** 2047–2058 (*Preprint* 1412.5852)
[278] Ruderman M 1991 ApJ **382** 587
[279] Parfrey K, Beloborodov A M and Hui L 2013 ApJ **774** 92 (*Preprint* 1306.4335)
[280] Hu C P and Ng C Y 2019 *Astronomische Nachrichten* **340** 340–345 (*Preprint* 1903.09736)
[281] Beloborodov A M 2009 ApJ **703** 1044–1060 (*Preprint* 0812.4873)
[282] Olausen S A and Kaspi V M 2014 ApJS **212** 6 (*Preprint* 1309.4167)
[283] Rea N, Esposito P, Turolla R, Israel G L, Zane S, Stella L, Mereghetti S, Tiengo A, Götz D, Göğüş E and Kouveliotou C 2010 *Science* **330** 944 (*Preprint* 1010.2781)
[284] Perna R and Pons J A 2011 ApJ **727** L51 (*Preprint* 1101.1098)
[285] Gavriil F P, Gonzalez M E, Gotthelf E V, Kaspi V M, Livingstone M A and Woods P M 2008 *Science* **319** 1802 (*Preprint* 0802.1704)
[286] Gotthelf E V, Vasisht G, Boylan-Kolchin M and Torii K 2000 ApJ **542** L37–L40 (*Preprint* astro-ph/0008097)
[287] Kumar H S and Safi-Harb S 2008 ApJ **678** L43–L46
[288] Kuiper L and Hermsen W 2009 A&A **501** 1031–1046 (*Preprint* 0904.4376)
[289] Livingstone M A, Kaspi V M and Gavriil F P 2010 ApJ **710** 1710–1717 (*Preprint* 0905.3567)
[290] Livingstone M A, Ng C Y, Kaspi V M, Gavriil F P and Gotthelf E V 2011 ApJ **730** 66 (*Preprint* 1007.2829)
[291] Archibald R F, Kaspi V M, Beardmore A P, Gehrels N and Kennea J A 2015 ApJ **810** 67 (*Preprint* 1506.06104)
[292] Livingstone M A, Kaspi V M, Gavriil F P, Manchester R N, Gotthelf E V G and Kuiper L 2007 Ap&SS **308** 317–323 (*Preprint* astro-ph/0702196)
[293] Kuiper L, Harding A K, Enoto T, Ho W C G, Gendreau K and Arzoumanian Z 2020 *The Astronomer's Telegram* **13985** 1
[294] McLaughlin M A, Lyne A G, Lorimer D R, Kramer M, Faulkner A J, Manchester R N, Cordes J M, Camilo F, Possenti A, Stairs I H, Hobbs G, D'Amico N, Burgay M and O'Brien J T 2006 Nature **439** 817–820 (*Preprint* astro-ph/0511587)
[295] Akbal O, Gügercinoğlu E, Şaşmaz Muş S and Alpar M A 2015 MNRAS **449** 933–941 (*Preprint* 1502.03786)
[296] Burgay M, Possenti A, Kerr M, Esposito P, Rea N, Coti Zelati F, Israel G L and Johnston S 2016 *The Astronomer's Telegram* **9286** 1
[297] Burgay M, Possenti A, Kerr M, Esposito P, Rea N, Coti Zelati F, Israel G L and Johnston S 2016 *The Astronomer's Telegram* **9366** 1
[298] Archibald R F, Burgay M, Lyutikov M, Kaspi V M, Esposito P, Israel G, Kerr M, Possenti A, Rea N, Sarkissian J, Scholz P and Tendulkar S P 2017 ApJ **849** L20 (*Preprint* 1710.03718)
[299] Majid W A, Pearlman A B, Dobreva T, Horiuchi S, Kocz J, Lippuner J and Prince T A 2017 ApJ **834** L2 (*Preprint* 1612.02868)
[300] Wang H H, Lin L C C, Dai S, Takata J, Li K L, Hu C P and Hou X 2020 ApJ **902** 96 (*Preprint* 2008.12585)
[301] Archibald R F, Kaspi V M, Tendulkar S P and Scholz P 2018 ApJ **869** 180 (*Preprint* 1806.01414)
[302] Thompson C and Duncan R C 1995 MNRAS **275** 255–300
[303] Gourgouliatos K N and Lander S K 2021 MNRAS **506** 3578–3587 (*Preprint* 2106.03869)
[304] Thompson C, Lyutikov M and Kulkarni S R 2002 ApJ **574** 332–355 (*Preprint* astro-ph/0110677)
[305] Li X, Levin Y and Beloborodov A M 2016 ApJ **833** 189 (*Preprint* 1606.04895)
[306] Blandford R D and Romani R W 1988 MNRAS **234** 57P–60
[307] Dang S J, Yuan J P, Manchester R N, Li L, Wang N, Wang J B, Hobbs G, Liu Z Y and Kou F F 2020 ApJ **896** 140 (*Preprint* 2005.02200)
[308] Larson M B and Link B 1999 ApJ **521** 271–280 (*Preprint* astro-ph/9810441)
[309] Gonzalez D and Reisenegger A 2010 A&A **522** A16 (*Preprint* 1005.5699)
[310] Melatos A, Douglass J A and Simula T P 2015 ApJ **807** 132
[311] Warszawski L and Melatos A 2012 MNRAS **423** 2058–2074 (*Preprint* 1203.4466)
[312] Bennett M F, van Eysden C A and Melatos A 2010 MNRAS **409** 1705–1718 (*Preprint* 1008.0236)
[313] van Eysden C A and Melatos A 2008 *Classical and Quantum Gravity* **25** 225020 (*Preprint* 0809.4352)
[314] Yim G and Jones D I 2020 MNRAS **498** 3138–3152 (*Preprint* 2007.05893)
[315] Ho W C G, Jones D I, Andersson N and Espinoza C M 2020 Phys. Rev. D **101** 103009 (*Preprint* 2003.12082)
[316] Singh A 2017 Phys. Rev. D **95** 024022 (*Preprint* 1605.08420)
[317] Stopnitzky E and Profumo S 2014 ApJ **787** 114 (*Preprint* 1305.2466)

[318] Abadie J *et al.* (LIGO Scientific Collaboration) 2011 Phys. Rev. D **83** 042001 (*Preprint* 1011.1357)
[319] Keitel D, Woan G, Pitkin M, Schumacher C, Pearlstone B, Riles K, Lyne A G, Palfreyman J, Stappers B and Weltevrede P 2019 Phys. Rev. D **100** 064058 (*Preprint* 1907.04717)
[320] Abbott R *et al* (LIGO Scientific Collaboration, Virgo Collaboration, KAGRA Collaboration), Antonopoulou D, Arzoumanian Z, Basu A, Bogdanov S, Cognard I, Crowter K, Enoto T, Espinoza C M, Flynn C M L, Fonseca E, Good D C, Guillemot L, Guillot S, Harding A K, Keith M J, Kuiper L, Lower M E, Lyne A G, McKee J W, Meyers B W, Ng C, Palfreyman J L, Shannon R M, Shaw B, Stairs I H, Stappers B W, Tan C M, Theureau G and Weltevrede P 2022 ApJ **932** 133
[321] Andersson N and Comer G L 2006 *Classical and Quantum Gravity* **23** 5505–5529 (*Preprint* physics/0509241)
[322] Warszawski L, Melatos A and Berloff N G 2012 Phys. Rev. B **85** 104503 (*Preprint* 1203.5133)
[323] van Eysden C A and Melatos A 2010 MNRAS **409** 1253–1268 (*Preprint* 1007.4360)