IMPROVING PULSAR TIMING PRECISION

A dissertation submitted to the University of Manchester
for the degree of Master of Physics
in the Department of Physics and Astronomy

2025

Jade Salisbury
(10867885)
Performed under the supervision of
Patrick Weltevrede

Department of Physics and Astronomy

Contents

Abstract

There are currently hundreds of known pulsars being tracked and regularly observed. Often, pulsars with more interesting activity such as the Crab or Vela pulsar are observed nearly daily, while others as infrequently as months. We used TEMPO3 (unreleased, Patrick Weltevrede) to generate times of arrival (TOAs) for fake pulsars at near periodic intervals twice per day. To investigate four different observation strategies (periodic, logarithmic, geometric, arithmetic) we separated specific TOAs from the datasets using four governing equations. These TOA sets were then used to explore the abilities of TEMPO2 to retrieve glitch parameters for glitches of various complexities. We found that at frequent observation cadence, there is little effect from strategy on fitted glitch parameters. However, as observations become less frequent, there is evidence to suggest that periodic observations retrieve glitch parameters most effectively. We discuss the possible implications of our findings for use in measuring real pulsar glitches.

Acknowledgments

Importantly I would like to first acknowledge the usage of GitHub's Copilot for AI assistance in code suggestion and line completion during some ($\simeq 50\%$) of the Python programming for this project. It needs to be stressed that outside of this functionality, no other AI tools, generative or otherwise, were intentionally used in this project. Our full code is open source and can be found on our GitHub repository: `https://github.com/Nyft-i/pulsar_timing_precision`.

I would also like to thank my project partner, Lauren Maclean. We had met only extremely briefly before pairing for this project, so it feels quite lucky that we work so well together and share a sense of humour.

It has been a great pleasure to have the opportunity to explore an area which I have become profoundly interested in very quickly. Before undertaking this project I was not convinced that I wanted to pursue a PhD in physics, but now the opposite is true. It has been beyond great to be surrounded by so many knowledgeable people in the field, to the point where a good portion of the names cited in the bibliography are people I have had the pleasure of meeting or interacting with in some capacity. The PET group at Manchester have been beyond friendly, encouraging and funny; it has been incredibly easy to submerge myself in this project, often to the detriment of my fourth year modules.

Finally I think it is important to mention by name my other supervisors on this project, as they do not all appear on the front cover and have each played a large part in helping us to work on this project: Danai Antonopoulou, Avishek Basu, Benjamin Shaw and Patrick Weltevrede.

1 Introduction

Neutron stars are compact objects formed from massive stars in core collapse supernovae. Their masses are then just low enough to avoid collapsing further into black holes. Due to conservation of angular momentum, neutron stars often spin notoriously quickly. They are highly magnetised and emit beams of radiation which follow the rotation of the star. Finding the rotation frequency, ν, and the time evolution of this number, is often a more difficult task than one might think. With some work, the rotational properties of pulsars far from earth can be precisely known, making them useful tools in physics, especially in testing Einstein's general relativity.

Due to their extreme properties (strong magnetic field, high density, etc.) the behaviour of pulsars' rotational evolution with time is of particular interest for us to better understand physics at these extremes. Pulsars are expected to slow with time due to conservation of energy as they emit radiation. Consequently, various phenomena which cause spurious or initially unexpected results in the observable properties of a slowing pulsar can allow us to learn the physical structure in the interior of these stars. One such phenomenon is known as a 'glitch': a sudden spin-up event causing an increase in the rotational frequency, ν, near instantaneously. With a glitch there is often also an associated change in frequency derivative, $\dot{\nu}$. Sometimes, there are one or several exponential recovery components before $\dot{\nu}$ arrives at a semi-permanent value. Understanding the origin and underlying physics causing each of the properties of a glitch can allow us to probe the deep interior of a pulsar, furthering theories which cannot yet model such conditions [1].

In the pursuit of new science through pulsar astronomy, its limitations must be well understood. Many pulsars emit in the radio frequency (sometimes in other parts of the electromagnetic spectrum also) and so are observed from earth with radio telescopes. The amount by which we can effectively observe pulsars is limited by several factors: there are more pulsars than radio telescopes and telescope time must be shared fairly between radio astronomy projects. This poses a problem of deciding how and when to observe pulsars in order to retrieve the most accurate estimations of their parameters. In this project we specifically explore how observation strategies impact the retrieval of glitch parameters in radio pulsars.

For a long time it has been accepted that observing a pulsar more often improves the ability to retrieve parameters via algorithmic fitting. Pulsars deemed interesting, like the Crab pulsar* or others which glitch regularly, are observed sometimes with near daily cadence (where cadence refers to the frequency or rate of observations) [2]. Basu et al. (2019) [3] discuss at length the difficulty in differentiation of pulsar glitches from the similar phenomenon, timing noise, at a multitude of different observation cadences. In several cases they conclude that a higher rate of observations is required to identify smaller glitches. Additionally, it was shown by Dunn et al. (2021) [4] that rigidly periodic observations can cause ambiguity in retrieved parameters. Marshall et al. (2004) [5] observed a glitching pulsar with a logarithmic cadence in order to keep phase uncertainties low. Therefore, it has been shown that other other observation strategies are worth investigating in a quantitatively comparative way: the inspiration for this project.

To make objective statements about particular observation strategies, one must observe the same pulsar and attempt to parametrise the same glitch. For each strategy to be fairly tested, they must be working with data over the same amount of time, and with roughly the same number of total observations. A strategy observing at a much higher overall cadence is bound to perform better. To ensure each strategy is measuring the same glitch, higher cadence data can be sampled with given strategies; observations not in accordance to that strategy should be removed from all parameter retrieval processes. Two possible methods emerge: the sampling of high density observations from frequently observed pulsars (such as the Crab) or the sampling of simulated pulsars. The latter option is chosen as we are able to control the parameters of a glitch and know them perfectly.

We present four observation methods to compare for this project: arithmetic, geometric, logarithmic and periodic. Arithmetic observations increase the gap between observations by a constant amount of time after each observation. Logarithmic is similar, but observation time increases in logarithmic space. A geometric strategy is achieved by multiplying the observation gap by a constant (> 1) after each observation. Periodic observations wait a constant amount of time between observations.

The goal of this project is to come to quantitative conclusions on the effects of observation strategy on the retrieved and fitted values of pulsar glitch parameters. We aim to investigate each of the strategies above at several different overall cadences on several glitches of various complexity. This way, in addition to comparing strategies to one another, they can be compared to themselves in specific scenarios, reducing the possibility of making a conclusion that may only hold sometimes. Pulsars such as the Crab and Vela are observed so often that to employ alternative strategies such as the ones described in this project, they must be studied under separate assumptions (see Section 3.3). Instead, we look to improve the quality of glitch pulsar astronomy for the hundreds of pulsars observed on timescales of days to months, where average time between observations

*A famous pulsar in the Crab Nebula: a supernova remnant first spotted in 1054AD exceeds a few days. We intend to better understand if the structure of observations can be rearranged to retrieve more accurate parameters and eliminate ambiguities without changing the overall number of observations.

2 Theory

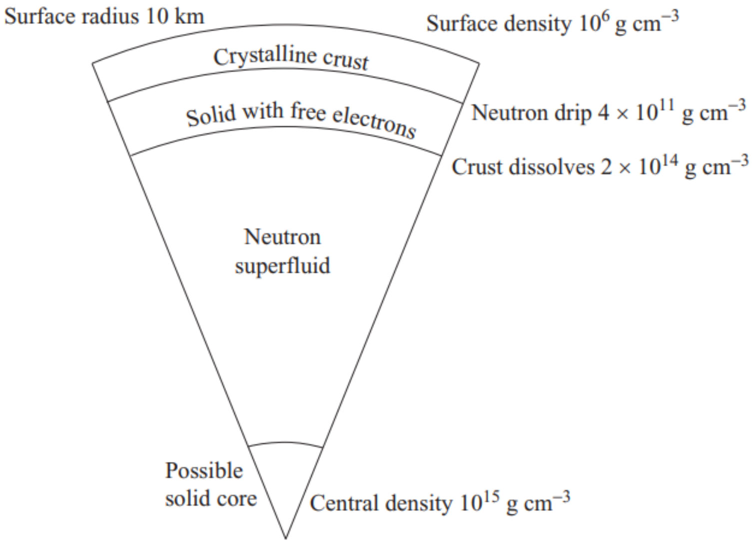

Figure 2.1: Assumed cross-section of a pulsar. Figure from [6].

2.1 Pulsars and Radio Astronomy

A pulsar is a type of rapidly spinning, highly magnetised neutron star. They are characterised by sweeping beams of radiation visible for brief moments during their rotation and are often compared to lighthouses as a result. Often, the radiation from the beams (or pulses, hence pulsars) lies in the radio spectrum, and therefore can be observed from Earth using radio telescopes. The formation of neutron stars follows from the end of massive progenitor stars with core masses above $1.4\mathrm{M}_{\odot}$, which collapse in type II supernovae, leaving behind comparatively tiny compact objects*. In cases where the neutron star still exceeds the TOV mass limit† it will collapse further into a black hole [6].

Following a core collapse, as neutron stars have such a small radius compared to their progenitor stars, by conservation of angular momentum they must increase the speed at which they rotate by a significant margin. Frequently pulsars rotate with periods less than a second. Those with periods shorter than order 1ms are referred to as millisecond pulsars (MSPs). Coupled tightly to the fast rotating star exists its magnetic field, extending beyond the radii of the star and dominating all nearby processes due to its strength. The dipole of the magnetosphere is misaligned with the neutron star's rotational axis, causing all associated emission processes to be directional and co-rotating with the star, providing the observed pulses.

Nearly all pulsars are always observed to be spinning down [7]. That is to say, the rate of their rotations are reducing with time. Due to energy conservation, a moving magnetic field will induce a counter-effect, causing the pulsar to lose energy with time. This loss of angular momentum is characterised by the breaking index, n, defined by $\dot{\nu} = -k\nu^n$, where ν and $\dot{\nu}$ are the frequency and spin-down rate respectively, and k is some arbitrary constant. The value of n is expected to be exactly 3 for only magnetic dipole breaking, but most pulsars have a breaking index lower than this, implying other methods of energy loss such as particle outflow [6].

While the full internal structure of a pulsar is not known, a simple model contains two main components: a rigid crystalline exterior and a superfluid interior, as shown in Figure 2.1. Starting from the outermost point, the star's crust is initially made from heavy nuclei like iron. Deeper into the crust, it becomes energetically favourable for electrons and protons to form neutrons, resulting in an abundance of nuclei with an abnormally high number of neutrons. Further still into the star, these heavy nuclei become unstable and the majority of matter is found as a neutron fluid, tightly coupled to the core [6].

At higher densities closer to the centre of the pulsar, our models of physics begin to break down and so theory cannot fully predict what can be expected. It has been suggested that a fluid made entirely of quarks [1],

*Around 10km in radius [6].

†The TOV limit is around $2\mathrm{M}_{\odot}$: See https://en.wikipedia.org/wiki/Tolman-Oppenheimer-Volkoff_limit

known as quark matter, may be present at the high densities in centre of the star, though no solid conclusions have been made.

2.2 Mechanisms for Pulsar Irregularity

Pulsars are remarkably stable. Their rotational periods can be found extremely precisely and therefore large arrays of tracked pulsars are often used as galactic arrays of clocks ‡, the measuring of which can be used to detect gravitational radiation [6, 8]. However, there exist two mechanisms by which the rotational rate of a pulsar will deviate from a simple slow with constant $\dot{\nu}$. This section will discuss both phenomena: glitches and timing noise.

Rarely, and somewhat unpredictably, a pulsar can rapidly experience a spin-up event known as a glitch. The exact mechanisms by which a glitch is triggered are unknown but it is assumed to be due to one of two causes: a sudden failure of the solid crust or a large-scale unpinning of superfluid vortices, causing a transferral of angular momentum from the core to the crust. In the former case, it is assumed that as the neutron slows, it at some point becomes energetically favourable for it to reduce its oblateness, causing crust-quakes [9]. Pulsars such as the Crab appear to be well described by this method; however, many other pulsars, the Vela included, are not [10].

The latter explanation is much more involved both mathematically and conceptually. Inner segments of a neutron star are known to be superfluid; rather than traditional bulk motion one would expect in a rotating fluid, angular momentum is carried in quanta by superfluid vortices, meaning that the total angular momentum of the fluid is proportional to the density of these vortices. As the pulsar slows, the only way by which the superfluid core can remain coupled to the crystal crust is through the migration of these vortices outwards. However, upon reaching the nuclei at the crust, a vortex may become pinned to them, fixing their contribution to the angular momentum and therefore decoupling the fluid from the crust. As the crust continues to slow, which is an observable property due to the coupling of the magnetic field to it, the difference between the spin rate of the crust and core increases, until eventually there is, for reasons not fully understood, a catastrophic failure and rapid transferral of angular momentum to the core via the unpinning of the vortices [6]. Some models propose this mass unpinning is triggered by the aforementioned crust-quakes [11].

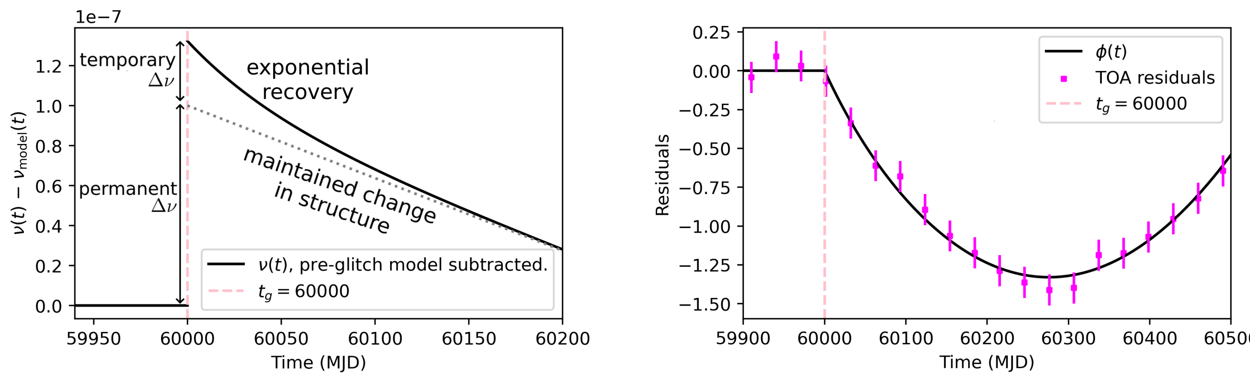

(a) Frequency deviations/residuals of a standard glitch with the pre-glitch model subtracted. Magnitudes of $\Delta\nu$ (the gap), $\Delta\dot{\nu}$ (the slope change) and τ_d (timescale of recovery) all differ between pulsars and individual glitches.

(b) A plot of simulated residuals after a glitch occurs at $t_g = 60000$. Data points show a potential set of TOAs which could be described by this glitch.

Figure 2.2: Representations of glitches in pulsar timing. 2.2a shows common glitch structure in terms of ν and 2.2b shows a potential residuals plot. Note the differing scales along the x-axis. Data is from simulated pulsars, and described by glitch B in Table 3.2.

In addition to instantaneous changes in ν and $\dot{\nu}$, some glitches exhibit a form of recovery behaviour (known as the post-glitch recovery or post-glitch response), where some fraction of the changes made to $\dot{\nu}$ are not permanent, and will tend towards (but not fully back to) their pre-glitch values on timescales of seconds to years [12]. The full structure of a pulsar glitch containing a single recovery component is shown in Figure 2.2a. This project focuses on glitches with only one or two exponential response terms.

The second mechanism by which pulsars may exhibit irregularities in their rotation is timing noise: a continuous wandering deviation in a pulsar's expected spin properties. The origins of timing noise are not as well understood§ as pulsar glitches, though it does appear to be a quasi-periodic switching between two modes

‡See pulsar timing arrays, or PTAs.

§Assuming that we believe current glitch models are at least partially correct.

of radio emission [6]. Sometimes, as a result of infrequent observation cadences, discrete jumps in ν and $\dot{\nu}$, like those belonging to pulsar glitches, can appear the same as the smooth transitions belonging to timing noise between two observations. Therefore there are often [3] cases where glitches cannot be distinguished from timing noise: especially when the glitches are small or timing noise is strong. We will discuss in Section 3.3 the reasoning behind not including timing noise in our pulsar models.

2.3 Timing a Pulsar

Simply observing a pulse and noting down its arrival time cannot be done. Individual pulses themselves are too weak to be individually identified in noise (i.e they have a low signal to noise, S/N, ratio). For a pulsar to be found at all, a fast Fourier transform (FFT) must be performed on some portion of the sky to find a periodic signal in an observation and form a guess at its period before many pulses over a much longer period of time (minutes to hours) can be 'folded' onto one another, greatly improving the S/N ratio of the pulse.

The folded data reveals a pulse profile: the shape traced by the intensity of the radiation beam while it is visible. The pulse profile can then be used to more accurately return arrival times of pulses in later observations [13]. Recently, Wang et al. (2022) [14] quantitatively compared several methods used over the last few decades by which pulsar arrival times can be calculated, the most common being the Phase Gradient Scheme (PGS) introduced by Taylor (1992) [15]. This project does not delve into collection of arrival times, but it is worth noting that methods akin to the PGS often collate uncertainties.

There is a significant dispersion effect on the light from pulsars due to frequency dependent scattering effects caused by electrons in the line of sight of observation, an effect magnified by the large distances between a pulsar and observer. Pulses need to be corrected (dedispersed) so that the they do not smear. Additionally, several positional effects cause pulses to arrive at a telescope differently on varying timescales as a result of rotation both in a day and around the sun, so the times of arrival (TOAs) are always adjusted so that they appear to be measured from the solar system's barycentre. Similarly, corrections are made to the TOAs to counteract relativistic effects¶ caused by massive bodies near the line of observation [13].

Given some set of measured TOAs, a pulsar can have its properties estimated through fitting with timing software packages such as TEMPO2 [16] or PINT‖. Methods such as these find a timing model which describe the pulsar at some specific epoch, containing information on values of ν, $\dot{\nu}$, its position on the sky, any proper motion components it may have, a measure of dispersion effects and more. These models get progressively more complicated as more information is known about a pulsar, such as information about specific glitches, timing noise, and whether the pulsar belongs to a binary.

When an observation is made of a pulsar, one arrival time (TOA) of a single pulse is retrieved, and has an associated error formulated using some scheme (such as PGS) from sources such as the telescope uncertainty and uncertainties in the known values required for each adjustment made to the arrival times as described above. Knowing many TOAs can build up a picture of how a pulsar rotates. For instance, a pulsar modelled to have only $\nu \neq 0$ would have each pulse arrive at some integer number of rotations after the previous. If acting predictably enough, every pulse, even those not observed, can be numbered exactly. Considering that on average pulsar spin periods are of order < 1s, precision is important and to solve for further complexity much data is required.

With a number of pulse TOAs, the spin evolution can be fitted to an equation of form

$$\nu(t) = \nu_0 + \dot{\nu}_0(t - t_0) + \frac{1}{2}\ddot{\nu}_0(t - t_0)^2 + \dots, \tag{2.1}$$

from a Taylor expansion of the pulsar's spin behaviour, frequently ignoring terms of order t^3 or higher [6, 13]. ν_0 is the value of ν at some period epoch $t = t_0$. $\dot{\nu}_0$ and $\ddot{\nu}_0$ describe both the spin-down and rate of change of spin-down. The information in Equation 2.1 is often represented in terms of phase, ϕ, instead:

$$\phi_{\text{model}}(t) = \phi_0 + \nu_0(t - t_0) + \frac{1}{2}\dot{\nu}_0(t - t_0)^2 + \dots, \tag{2.2}$$

with ϕ_0 representing a phase reference point at $t = t_0$. A pulsar which follows its timing model perfectly will have its TOAs occur at exactly integer numbers of ϕ. Pulses are therefore represented by residuals: their deviation from expected arrival times (given by integer values of phase in a model), with respect to time.

Should a glitch occur, there exists a discontinuity** in ϕ and a post-glitch behaviour described by

$$\phi_g(t \geq t_g) = \Delta\phi_g + \Delta\nu_g(t - t_g) + \frac{1}{2}\Delta\dot{\nu}_g(t - t_g)^2 + \sum_i \Delta\nu_i\tau_i(1 - \exp(-(t - t_g)/\tau_i), \tag{2.3}$$

¶E.g. Shapiro delay and other curved space effects. See `https://en.wikipedia.org/wiki/Shapiro_time_delay`

‖`https://github.com/nanograv/PINT`

**Pulsars are rarely observed mid-glitch. In such scenarios they contain apparent rising term(s) [12].

added to the model for $t > t_g$. Similarly to Equation 2.2, a glitch occurring at $t = t_g$ is a Taylor expansion of its effects on the spin, ν. $\Delta\phi_g$ is a measure of the amount of pulsar turns lost during a glitch. $\Delta\nu_g$ and $\Delta\dot{\nu}_g$ refer to the the instantaneous and permanent changes in parameters ν and $\dot{\nu}$. The final term describes the exponential recovery components in the glitch of which there may be one ($i = 1$) or multiple ($i > 1$). $\Delta\nu_i$ is the instantaneous change(s) in ν and τ_i is the decay time over which $\Delta\nu_i$ decays [16, 17]. The structure of a glitch in ν can be seen in Figure 2.2a. Typically, the pre-glitch model, ν_{model}, is subtracted, manifesting as the pre-glitch structure being flat and the post-glitch structure being fully representative of the glitch; glitch parameters can often be numerically small compared to ν and $\dot{\nu}$, therefore they can be better investigated this way. Similarly, a residuals plot as described above can be seen in Figure 2.2b.

In the work of Zhou et al. (2022) [11] it is noted that ~6% of pulsars have been seen to glitch, manifesting in the pulsars residuals decreasing over time with expected pulses. Eventually, a pulsar model will 'lose coherence' with data measurements, meaning that the model is more than a full rotation out of phase with the real pulsar behaviour. Sufficiently large glitches can cause immediate loss of coherence, and are particularly difficult to retrieve parameters for. The prevalence of timing noise in a pulsar increases the separating small glitch events in the data, and increase the likelihood of losing coherence in larger glitches.

3 Method

This project details 3 main steps which we have used to come to objective comparisons of different observation strategies: generated TOAs of pulsars described by parameters we know perfectly; code to sample the TOAs at intervals described by the strategies, and TEMPO2 to time and retrieve parameters for the pulsar and its glitches via algorithmic fitting.

3.1 Observation Strategies

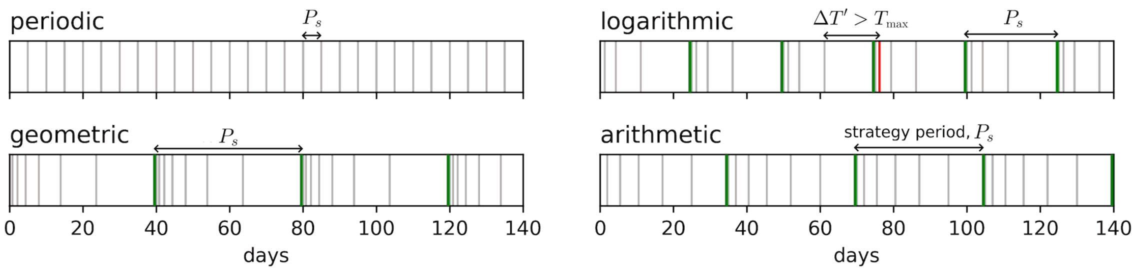

Figure 3.1: A comparative event plot corresponding to three examples of observation strategies with constant and equal average cadences of 5 ($k_p = 5$, $k_l = 25.7196$, $k_g = 1.6394$). An observation on any plot is represented by a gray vertical line. Periodic observation strategy has equally distributed observation times. Both logarithmic and geometric observation strategies exhibit periodic structure, represented by the strategy period, P_s. The strategies pictured restart (shown by the green/thicker line) at a cadence of 0.5d when $\Delta T'$ exceeds $T_{\max}$. Shown once by the red line in the logarithmic strategy is where the next observation would have occurred had the strategy not restarted (for instance if $T_{\max}$ was higher).

Periodic observations are done by the majority of pulsar experiments to date. To the best of our knowledge, only two other groups have considered doing otherwise for reasons relating to improving precision. For instance, the MeerKAT survey observes newly discovered pulsars with a pseudo-logarithmic cadence to establish a good model quickly before reverting to a periodic strategy [18]; the goal of the non-periodic observations was to reduce the cadence over time rather than improve precision or eliminate degeneracy. A similar but differently motivated example is given by Marshall et al. (2004) [5], where they instead chose a logarithmic* strategy in order to keep phase uncertainties below 0.1 cycles. Dunn et al. (2021) [4] talks in great detail about how exceptionally periodic observations can cause degeneracy in the retrieved results and mention that software like TEMPO2 can sometimes underestimate the sizes of glitches.

We propose three alternative observing strategies for this project: arithmetic, geometric and logarithmic. We additionally perform all of our experiment on fake, periodically observed pulsars in order to create fair comparison and a direct analogue to the methods by which real pulsars are observed in research. Periodic observations can be described by

$$\Delta T' = \Delta T = k_p, \tag{3.1}$$

*It is unclear if their definition of logarithmic observations is similar to ours.

where, assuming we have just performed an observation, ΔT is the time since the previous observation, and $\Delta T'$ is the time until the next observation. Parameter k_p is an arbitrary constant equal to the number of days between observations. An equation similar to 3.1 can describe all strategies, where some operation acts on the time between previous observations to find the time to wait until the next observation, and each strategy can be parametrised by an arbitrary constant k, with a subscript denoting the strategy. The need for other constants will emerge later.

Similarly, geometric observation cadence can be described by

$$\Delta T' = k_g \Delta T. \tag{3.2}$$

Geometrically observing multiplies the TOA gap by a constant, k_g, after each TOA is taken. Geometric cadence at $k_g = 2$ would cause the time between observations to double after each.

Continuing, an arithmetic strategy takes the form of

$$\Delta T' = \Delta T + k_a, \tag{3.3}$$

as one might expect. Parameter k_a describes the amount of time the observation gap increases by with each successive measurement.

Finally, the logarithmic strategy takes the form of

$$\Delta T' = k_l \ln\left(\frac{\Delta T}{10} + 1\right), \tag{3.4}$$

where the division and addition by constants 10 and 1 respectively are to reduce runaway and convergence effects found in the equation $\Delta T' = k_l \ln(\Delta T)$. Constants k_p, k_g, k_a and k_l are somewhat arbitrary and govern how a strategy changes its observation gaps. It may be clear that the above equations can only work if there is an initial value assigned to ΔT, given as ΔT_{start}. In this project our lowest start cadence of $\Delta T_{\text{start}} = 0.5$d was motivated by the observation cadence of the most frequently observed pulsars [2]. In lower cadence scenarios, other values of ΔT_{start} were chosen and are detailed in Figure 3.1.

In addition to the values of k and ΔT_{start}, there exists one other parameter important to all strategies except periodic: the "maximum gap", $T_{\max}$. The need to introduce this constant is best demonstrated through an example: geometric sampling at $k_g = 2$ would quickly cause observations to be spaced at intervals of extremely high powers of 2 (i.e. 1 day, 2d, ...128d, 256d, etc.). Despite the earlier high sampling of every few days, this strategy becomes near useless at late times. To combat this runaway effect, we implement a maximum time between observations, $T_{\max}$: if the next observation would occur at a time beyond $T_{\max}$, it instead happens at $\Delta T_{\text{start}}$for that strategy as described in Table 3.1. The effect of this can be seen in Figure 3.1. The total time between the first observation in a strategy and the last before $T_{\max}$ is exceeded is referred to as the strategy period, P_s. It should be noted that any value of $T_{\max}$ could be selected, and should be chosen based on the scenario.

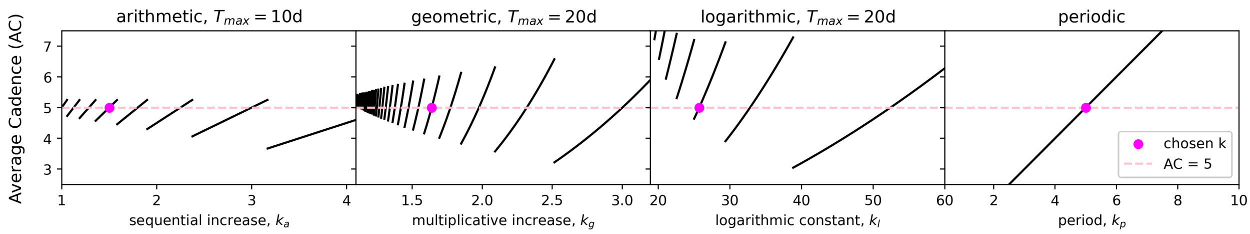

Figure 3.2: Demonstration of average cadence (AC) against the values of k for all of the explored cadence strategies at $T_{\max}$ denoted by Table 3.1. Note that all subplots share a y-axis. Locations in the plot where a horizontal line of constant AC (such as the dashed pink line above) crosses a line from a strategy shows values of $k_{g,l,a,p}$ which share a number of observations/number of TOAs per unit time. Magenta points signify the values of k chosen.

Following from these equations, it is important to ensure that we can pick values of k, $T_{\max}$ and ΔT_{start} for each strategy that on average result in the same average observation cadence. Two strategies working at a roughly similar average cadence (AC) will observe roughly the same number of times over an equal amount of time. Average cadence, or AC, refers to the average number of days between observations. Confusingly, high cadence refers to frequent observations but a high value of AC implies the opposite. Extreme care has been taken to in this report to ensure what is being discussed is clear. Figure 3.2 demonstrates how multiple values of k for the same strategy have an equal AC. The AC can be easily calculated by dividing the observations in one strategy period by the strategy period itself. To fairly compare two strategies, their AC must be equal.

For our project, we decided to try strategies at three different ACs (5d, 15d, 30d) specifically so it could be investigated if some strategies only outperform others at higher/lower AC. Table 3.1 shows the chosen ACs and the corresponding chosen values of constants. Our values of k were picked to ensure that each strategy retrieved 5-10 TOAs in one P_s before restarting.

AC	Strategy	$k_{a,g,l,p}$	$T_{\max}$	ΔT_{start}	Num. obs.	Num. sims.
5	arithmetic	1.5000	10 d	0.5d	593	$10,000$
	geometric	1.6394	20d	0.5d	591	$10,800$
	logarithmic	25.7197	20d	0.5d	594	$10,500$
	periodic	5.0000	N/A	N/A	599	$10,500$
15	arithmetic	4.3333	30d	2d	196	$10,000$
	geometric	1.6693	50d	2d	194	$10,000$
	logarithmic	34.7648	50d	2d	196	$10,000$
	periodic	15.0000	N/A	N/A	199	$10,350$
30	arithmetic	1.8063	58d	2d	102	$10,000$
	geometric	3.5075	90d	2d	95	$10,000$
	logarithmic	35.2264	70d	2d	95	$10,000$
	periodic	30.0000	N/A	N/A	99	$10,000$

Table 3.1: A table to show the specific cadence constants, $k_{a,g,l,p}$, $T_{\max}$, ΔT_{start} , used and the number of observations they gave over the observation time.

Glitch ID	Parameter	Value
A	ν	4 Hz
	$\dot{\nu}$	-1.8×10^{-12}
	$\Delta\nu$	1×10^{-7}
	$\Delta\dot{\nu}$	-4.2×10^{-15}
B	ν	4 Hz
	$\dot{\nu}$	-1.8×10^{-12}
	$\Delta\nu$	1×10^{-7}
	$\Delta\dot{\nu}$	-4.2×10^{-15}
	$\Delta\nu_d$	3.23×10^{-8}
	τ_d	50d
C	ν	4 Hz
	$\dot{\nu}$	-1.8×10^{-12}
	$\Delta\nu$	1×10^{-7}
	$\Delta\dot{\nu}$	-4.2×10^{-15}
	$\Delta\nu_s$	4.70×10^{-8}
	τ_s	5d
	$\Delta\nu_l$	3.23×10^{-8}
	τ_l	100d

Table 3.2: A table depicting each of the simulated glitches and their respective parameters.

3.2 Simulation

In this project, "simulation" refers to the generation of fake pulsar TOAs over a given period of time, fitting glitch parameters using an incomplete model and then investigating how much the fitted results differ from the true values. Performing this multiple times with varying start and end points can provide a better picture of how a method performs on average.

Each subset of simulations acts upon a set of 6000 TOAs, generated by TEMPO3 (unreleased, Patrick Weltevrede [19]) and distributed randomly (with uniform probability) onto the pulses described by the pulsars in Table 3.2. They span 3000 days†, causing an AC for the whole data set of 0.5d. Each TOA was generated with an associated uncertainty in arrival time of 100µs, which is typical [14]. The TOA data was then sampled at a particular strategy, with an algorithm selecting only the TOAs which lie near to the times of observation described by that strategy. Rarely, a high cadence observation would have its closest TOA be a previously

†MJD 58500-61500

selected one; in such scenarios the next closest TOA was chosen instead. This generation process is repeated a number of times per simulation set.

Strategies with non-constant cadences, changing on a cycle, will have regions in their strategy both where the observation frequency is higher and where it is lower than the AC. This is a good motivation behind investigating alternative observation strategies to begin with, as mentioned earlier it can be used to keep phase uncertainties low or remove a degree of ambiguity. However, it does create "good" and "bad" regions for a glitch to occur in an observation strategy. Glitches occurring where the observation separation is highest will be the most difficult to solve correctly as it is much more difficult to resolve a glitch occurring some million(s) of pulses ago versus one which occurred in the previous few days: there is less time for the behaviour to diverge from a descriptive model. It should be noted that the inverse is also true for regions of high observation cadence. Therefore, when generating TOAs, it would be unfair to compare the results of a periodic observation strategy to another which happens to have its highest sampling rate at time $t = t_g$ when the glitch occurs. To counteract this effect, we introduced a random offset to the start of the dataset on an individual simulation basis. As a result, the start point and therefore the high cadence points, which are separated by time P_s, are varied in each simulation by an amount up to P_s (i.e. $0 \leq$ offset $\leq P_s$). A TOA set is sampled at different offsets a number of times less than or equal to $\sim P_s/0.5\text{d}^{\ddagger}$, before a new one is generated. Every sampling of each TOA file is passed through TEMPO2 for fitting and parameter retrieval, and stored for later plotting. TEMPO2 is setup in different ways for the three complexities of glitches which were investigated as detailed in Table 2.2.

The true parameters of the glitches simulated are described in Table 3.2. Simulations ran are different permutations of glitches with differing observation strategies. The parameter values of the glitches were chosen to reflect real values of glitches found in the JBO glitch catalogue [20]. Glitch A in Table 3.2 contains no recovery component, only an instantaneous jump in ν and $\dot{\nu}$. Glitch B introduces a single recovery component of order $\tau_d = 50$d, longer than the average cadence for all explored strategies. Glitch C contains two separate exponential recovery components, of order 5d and 100d. The shorter of the two recovery components, with parameters $\Delta\nu_s$ and τ_s, has a timescale which matches the average cadence of the most frequent observation strategy (5d). This has been done in order to explore the quality of results as observation cadence exceeds the timescale of decay. The longer of the exponential responses has parameters $\Delta\nu_l$ and τ_l.

Time constraints stopped all possible permutations from being trialled, so emphasis was placed on the more complex glitches, B and C, as preliminary results revealed very little difference between quality of results between strategies for simpler glitches. Each permutation of strategy and glitch consisted of a number of TOA files and sampling offsets such that their product was roughly 10,000; this number of simulations was motivated only by the total runtime being approximately 1 hour.

3.3 Assumptions and Caveats

In order to retrieve useful results, we have made a handful of assumptions in our simulations. Each one is listed here with justification.

In order to ensure that convergence algorithms are not working from the value 0, when provided with data for a glitch in generated TOAs we also provide it with an initial model, which perfectly describes the pre-glitch behaviour of the simulated pulsars. In fitting, these two parameters are still free to change, as would be the case in a real post-glitch scenario, but exact pre-glitch values allow any convergence algorithms more iterations in glitch parameter recovery over spin parameters. This assumption is motivated by many real tracked pulsars having extremely well described models before a glitch occurs [21]. It can be expected that with real pulsars, provided there are no systematic errors, known parameters in literature are on average distributed around the true values.

Coherence in pulsar astronomy is the property that the measured post-glitch residuals maintain some amount of identifiable structure, and can be directly compared to figure similar to 2.2b. We make this assumption as it often is true for small and medium sized glitches. Larger glitches are much easier to solve by hand and therefore could be adjusted to re-establish coherence manually. For similar reasons, we assume that the glitch epoch can always be roughly known by a human observer; that is to say that the two TOAs between which the glitch occurred can be easily identified. Another way by which we maintain coherence in our simulations is by allowing TEMPO2 perfect knowledge of the pulse numbers of each pulse: this ensures that the pulsar will not choose a solution which has lost a single rotation or more during the glitch.

The above assumption can be restated as an assumption of human solvable glitch. Glitches occur infrequently enough in pulsars such that it is not unreasonable to assume than an observer would be able to check that algorithmically fitted estimations are sensible. This assumption helps us in simulation as it allows us to throw

‡Ideally we would generate a new set of TOAs for every simulation, but this is computationally inefficient. The maximum number of offsets which can be simulated before we would start repeating results is equal to $P_s \times 1/\text{AC}$. AC here of the whole dataset is 0.5d

out estimations which are clearly erroneous and, in a realistic scenario, would be nudged towards a more feasible model by some astronomer.

Additionally, we are always assuming that there is no timing noise. Espinoza et al. (2014) [22] show that there is good reason to believe that the trigger mechanisms between timing noise and glitches are different, providing some motivation to avoid timing noise in our simulations. Hobbs et al. (2010) [23] show that for pulsars of a young characteristic age ($\tau_c < 10^5$), the recovery component of previous glitches often dominates over what is observed as timing noise. Older pulsars such as millisecond pulsars tend to exhibit very small amounts of timing noise, but also glitch incredibly rarely [24]. Timing noise is an important factor seen to obscure glitch parameters, especially for smaller glitches, adding difficulty that is not present in our simulations. Due to the apparent quasi-periodic nature of timing noise [6], over shorter timescales of order months around a glitch, it can often be ignored or its effects are reduced. Though our simulations span roughly 8 years, our generated TOAs are acting predictably and without timing noise, so conclusions will hold true over timescales where timing noise can be safely ignored.

As detailed before, our simulations are of glitches made up of either no recovery components (glitch A in Table 3.2), a single recovery component (glitch B) or two recovery components (glitch C). For ease of simulation by time constraints, our algorithm always knows the number of exponential components present, but this ended up causing some problems later. Only single sizes of each of these glitches were picked, each sharing many parameters with one another and being on the larger side in magnitude [20]. Ideally, given enough time, we would have trialled additional glitch sizes. Luckily, due to the non-existence of timing noise, it can be assumed that results would be similar for smaller glitches, as they would still be identifiable rather than being lost in noise [22]. Similarly, in the data samples generated, only a single glitch occurs in > 8yrs. In truth, this is unlikely for some pulsars and likely for others. Using the same justification as above, our pulsar predictability means that comparative results would apply over shorter lengths of time (i.e. the length of time before another glitch occurs).

Akin to how we provide a perfect pre-glitch solution, we allow guesses of the exponential timescale equal to the exact timescale itself for fitting in TEMPO2. In realistic scenarios, typically only the order of magnitude can be discerned through visual inspection of the residuals, but this would introduce a bias in a particular direction with our simulations, so convergence is allowed to begin on the true parameters and we let the timescale wander through free fitting of all parameters.

Finally, the errors given on TOAs by TEMPO3 [19] are perfectly 100µs in every arrival time. The value of 100µs is on the higher end of typical for errors given by a common TOA creation method, PGS [14]. This number on all TOAs provides the error on all derived values from TEMPO2, which are found to be significant underestimates in our case, derived parameters are found to deviate significantly from the sample mean frequently, as will be shown in Section 4 below.

4 Results

Each glitch and its various parameters are detailed in respective figures. The information on averages and standard deviations is also presented in three tables for glitches B and C. Naturally, we cannot assume that all the fitted variables are independent and Gaussian, and so we also present a handful of corner plots to ensure that we understand the correlations in the parameters [25]. Additionally we present contour plots tracing the density of scatter plots associated with glitch parameters $\Delta\nu$, $\Delta\dot{\nu}$, and any response magnitudes and timescales.

While it is seen that most strategies at all cadences will on average be able to retrieve the correct parameters*, the spread on these calculated model solutions given by TEMPO2 will be used as an analogue for the quality of a strategy: a strategy able to retrieve the most accurate results most frequently is a better observation strategy than one with an increased likelihood to fit parameters further from their true values. For each glitch we present tables of retrieved values averaged from glitch values fitted by TEMPO2, with errors given by the standard deviation on the averaged values.

As glitch A contains only an instantaneous jump in ν and $\dot{\nu}$, with no recovery components (the final term in Equation 2.3 = 0), it is an unrealistic scenario, as often glitches exhibit at least some additional behaviour, such as that seen in glitch B and C below. As a result and to save space we choose not to present it in the same way. Correlations between parameters are similar to glitch B and any detail we would have presented here would have been restated for glitch B anyway. Simulation of glitch A was done as a stepping stone to increasing complexity to the level found in glitch B and C.

One important process by which we can investigate the validity of our results is to justify the correlations found in corner plots in figures 4.2 and 4.5, showing parameter correlation in glitch B and C respectively. It should be noted that the corner plots presented are collated from all strategies at an AC of 5d for each glitch. In all but a single case (discussed later) the correlations appeared to be consistent between the different strategies.

*Especially as a result of some of our assumptions

The only difference at higher values of AC was some correlations lessening or in some cases being entirely removed. This is expected as the restrictions on a model by infrequent TOAs are less harsh compared to those which are more frequent. In real pulsar astronomy, understanding parameter correlation can help derive errors and uncertainties on found values of real pulsar glitches (see Bayesian techniques in pulsar astronomy [26]). As such, the presentation of our corner plots is a necessary first step in understanding the underlying statistics which may emerge differently in employing more complex observation strategies.

4.1 Glitch B

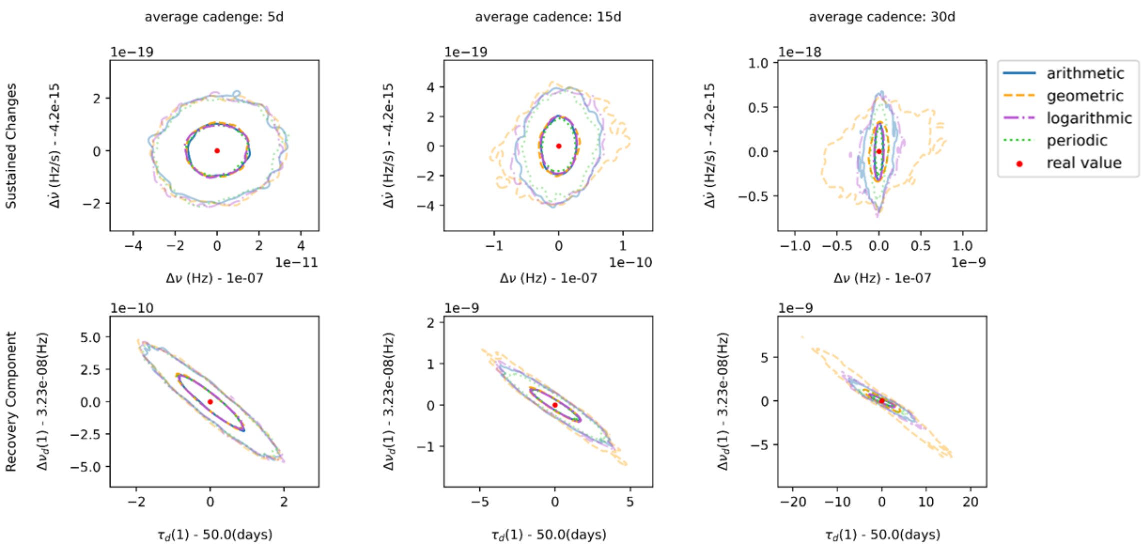

Figure 4.1: Contour plots showing the distribution of retrieved glitch parameters for glitch B, split into individual two dimensional plots separated by AC and retrieved component (i.e. base parameters, recovery parameters). Contour lines differ between observation strategy as shown in the legend. The inner contours contain 67% of points and outer contours contain 99% of points. The shape of the underlying point scatter is maintained.

Parameter	AC	Arithmetic	Geometric	Logarithmic	Periodic
	5	0 ± 10	0 ± 10	0 ± 10	0 ± 10
$\langle \Delta\nu \rangle - \Delta\nu_{\text{true}}$ $(\times 10^{-12})$	15	0 ± 19	0 ± 24	0 ± 20	0 ± 17
	30	1 ± 45	2000 ± 6600	29 ± 810	0 ± 29
	5	0 ± 7	0 ± 7	0 ± 7	0 ± 6
$\langle \Delta\dot{\nu} \rangle - \Delta\dot{\nu}_{\text{true}}$ $(\times 10^{-20})$	15	0 ± 12	0 ± 12	0 ± 12	0 ± 11
	30	0 ± 19	0 ± 18	0 ± 19	0 ± 17
	5	-1 ± 14	0 ± 14	-1 ± 14	0 ± 13
$\langle \Delta\nu_d \rangle - \Delta\nu_{d,\text{true}}$ $(\times 10^{-11})$	15	0 ± 27	0 ± 32	-1 ± 29	0 ± 25
	30	-1 ± 56	700 ± 2400	10 ± 310	-1 ± 40
	5	0 ± 6	0 ± 6	0 ± 6	0 ± 6
$\langle \tau_d \rangle - \tau_{d,\text{true}}$ $(\times 10^{-1})$	15	0 ± 11	0 ± 12	0 ± 12	0 ± 10
	30	0 ± 20	0 ± 39	0 ± 23	0 ± 16

Table 4.1: A table showing the average distance from the true values retrieved by TEMPO2 when fitting glitch B. Errors quoted are standard deviations of each sample of ∼10,000 points. It is expected that the the mean difference is around 0 and the standard deviation represents the spread of the data, varying massively in some cases. True values can be found in Table 3.2. Values are rounded to similar precision, but no more than two significant figures (on the standard deviation).

Glitch B results detail the first set of simulations which may be comparable to some real glitches under our assumptions. There exists both instantaneous jumps and a singular exponential recovery component. We present corner plot Figure 4.2, containing the two new parameters $\Delta\nu_d$ and τ_d. It can be seen that some parameters which were correlated in Figure 4.4 are no longer correlated, or correlated to a lesser degree for all cadence strategies. Again there was little difference between correlation depending on the strategy used and so the data was aggregated for use in the corner plot.

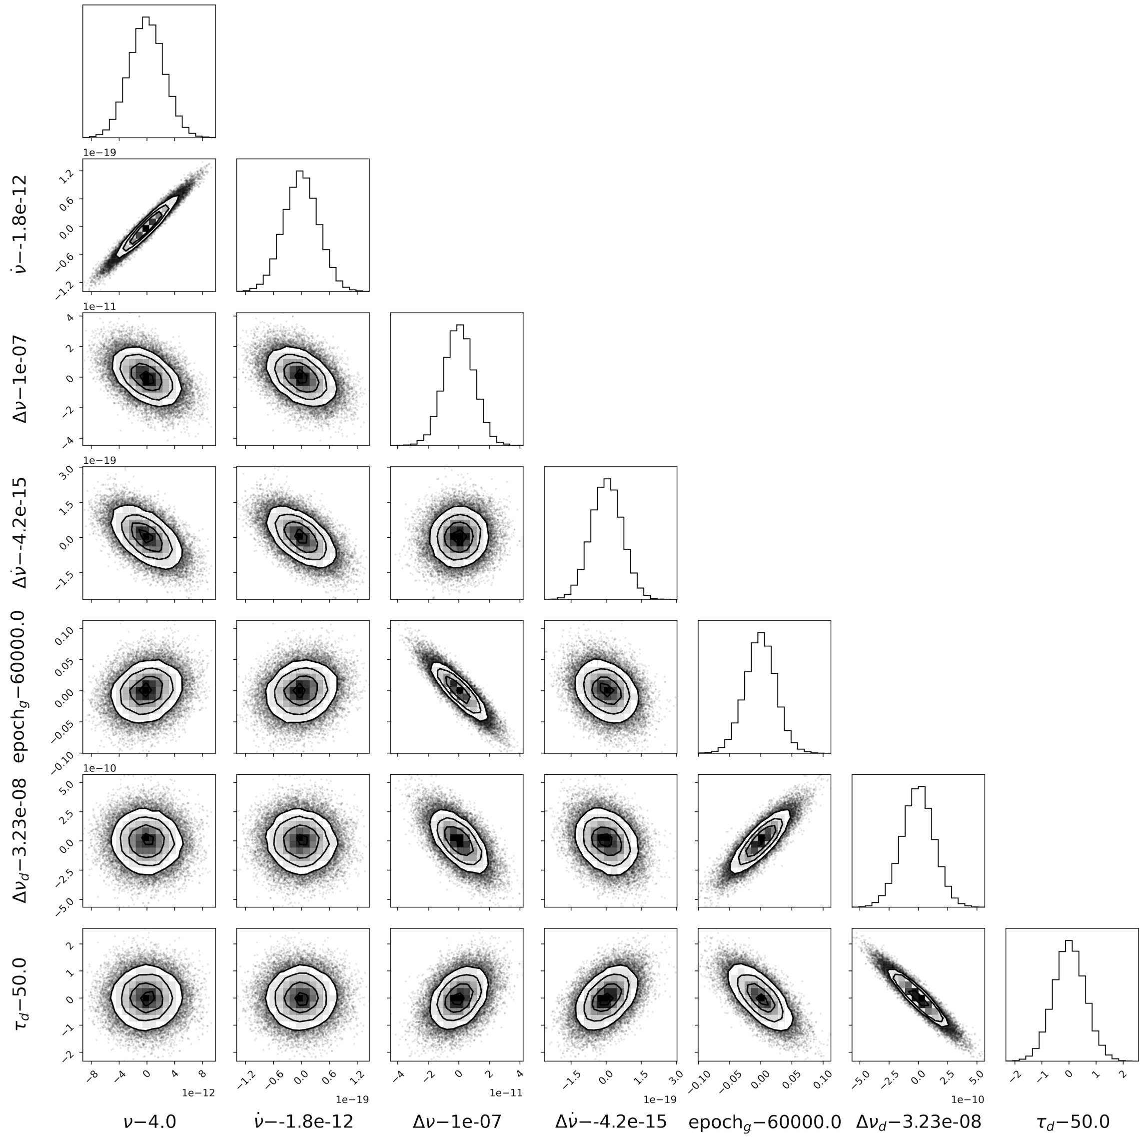

Figure 4.2: A corner plot for for glitch B collated from ∼40,000 data-points measured using all three strategies at an AC of 5. Note strong correlations between parameters ν and $\dot{\nu}$, epoch$_g$ and $\Delta\nu_d$; there are also anti-correlations in parameters $\Delta\nu$ and epoch$_g$ (t_g), $\Delta\nu_d$ and epoch$_g$ (t_g), $\Delta\nu_d$ and τ_d.

Additionally, Figure 4.1 shows the retrieval quality at each of the three trialled cadences. Parameters $\Delta\nu$, $\Delta\dot{\nu}$, $\Delta\nu_d$ and τ_d are seen to follow near identical distributions at high (AC = 5d) cadence in Figure 4.1. Roughly similar results follow in the medium (AC = 15d) cadence contours, though it is clear that geometric is beginning to fall behind: its contours and therefore retrieved parameter distribution no longer coincides with the other strategies and therefore is said to be returning less accurate parameters. This can also be seen in Table 4.1, where the standard deviation (quoted as the error) of geometrically retrieved parameters for $\Delta\nu$ and $\Delta\nu_d$ in all cases is larger than the same for any alternative strategies. It should be mentioned that in $\Delta\dot{\nu}$ and τ_d, the strategy performs roughly the same as its peers.

Advancing to the low cadence (AC = 30d) case, the difficulty geometric observing has in estimating glitch parameters is only amplified, now falling behind in all but the value of $\Delta\dot{\nu}$. Logarithmic observing is also significantly worse at retrieving parameters in the instantaneous changes to ν (i.e. $\Delta\nu$ and $\Delta\nu_d$).

To our knowledge, all of the correlations seen in Figure 4.2 are as expected. Particularly of note is the anti-correlation between parameters $\Delta\nu_d$ and τ_d. The following is best explained in tandem with Figure 4.3. Suppose two near identical models which describe a set of given TOAs, one of which has a higher $\Delta\nu_d$ than the true value, and the other, a lower. For the set of TOAs following a glitch, TEMPO2 fits a supposed model to the data and retrieves a number of parameters; for both models to describe the TOAs within their error bars, a larger instantaneous jump in ν would require a faster (smaller) response timescale, τ_d, such that the value of ν is pulled closer to the value described by the post-glitch TOAs in a timely manner. The model with a lower $\Delta\nu_d$

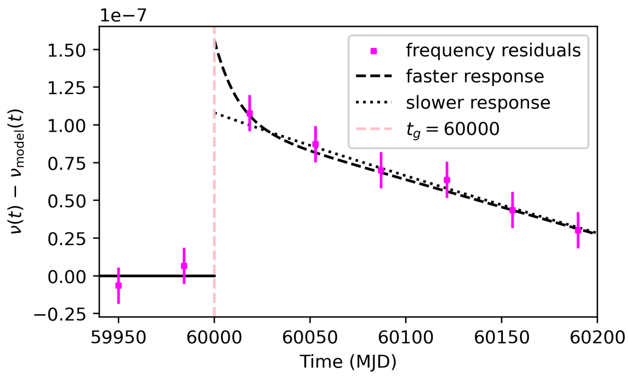

Figure 4.3: Fictitious plot demonstrating how two differing models may be good fits for the same set of TOAs. Two models with differing values of $\Delta\nu_d$ and τ_d describe behaviour exhibited in the TOAs well. Similar plots can be used to describe all correlations. As before, pre-glitch model of ν has been subtracted from data.

is not under time-constraints quite as strict to adhere to post-glitch observation and in fact, a similar response might overshoot the behaviour described by the arrival times.

Similarly, Figure 4.2 shows more easily justifiable correlations in glitch epoch (t_g) and both $\Delta\nu$ and $\Delta\nu_d$. Should the time at which a glitch occurs be earlier than its true time, the model must also have a greater instantaneous shift in ν as a result of the increased time the pulsar is given to spin down to the value of ν described by the post-glitch TOAs. This manifests as an earlier t_g being correlated with greater values of $\Delta\nu$ or $\Delta\nu_d$, as seen in the corner plot. The inverse also applies: if a glitch is thought to occur late, it must also be modelled with a reduced shift in ν to account for the lessened time it has to align with observation, hence the correlation.

It is important to note that even though many of these parameters are correlated, the scale of their spreads must be taken into account. Take for instance the correlation in ν and $\dot{\nu}$: it may be difficult to imagine a solution with which a differing ν may be compensated for with a similarly different $\dot{\nu}$. This is analogue to fitting multiple straight lines to N roughly collinear points. As these are not glitch parameters, TEMPO2 has data from the span of ~8 years to fit to, equivalent to N of order thousands. $\nu_{\text{true}} = 4$ Hz, and the scale of its standard deviation shown in Figure 4.2 is of order 10^{-12}: a profoundly accurate measurement. The small amount that it does vary is compensated for by a similarly small variation in $\dot{\nu}$, but compared to some other correlations, both these parameters are effectively constant. The corner plots must be read with caution and the scales of variation carefully examined before definitive conclusions are made in all cases.

In some scenarios for the geometrically sampled TOAs, TEMPO2 would estimate a solution with a phase offset, $\Delta\phi_g$, of roughly 60 rotations†. This is a curious result as we should expect to know perfectly the number of lost rotations when TOAs contain pulse number information. We decided for our results we would filter out all these solutions, with justification from our assumption of human solvable glitch: no astronomer would accept, without some other motivation, these solutions as realistic.

4.2 Glitch C

Glitch C is also described in Table 3.2. It differs from glitch B by containing two different exponential responses, at two vastly differing timescales, rather than one. The short response, denoted with a magnitude $\Delta\nu_s$, has a timescale $\tau_s = 5$d. The longer, with magnitude $\Delta\nu_l$ has a timescale $\tau_l = 100$d. The shorter timescale is exactly equal to the average cadence of the most frequent observations. Pulsars such as the Vela often contain one or more response components on timescales as short as this and sometimes even of orders seconds to minutes [12]. The averaged retrieved glitch parameters with their sample standard deviations can be found in Table 4.2.

We present corner plot Figure 4.5, where it can be seen that some of the correlations discussed in 4.1 are no longer present or are altered. For instance, correlations between timescales such as $\tau_{s,l}$ and response magnitudes $\Delta\nu_{s,l}$ are no longer as simple. There exists also curved correlations between different timescales such as exhibited between $\Delta\nu_s$ and $\Delta\nu_l$. An exceptionally interesting correlation manifests in parameters $\Delta\nu_s$ and τ_s, having an apparent "cut-off" point manifesting in a harsh line in the scatter plot between these two values.

Much of the justification of correlation holds for glitch C also. Any parameters which are differently correlated are to be expected also, particularly between response parameters as the number of restrictions is loosened as the glitch complexity increases; there is a much larger set of models which can describe any given set of

†As in the glitch is modelled to have lost ~60 rotations before the first post-glitch TOA.

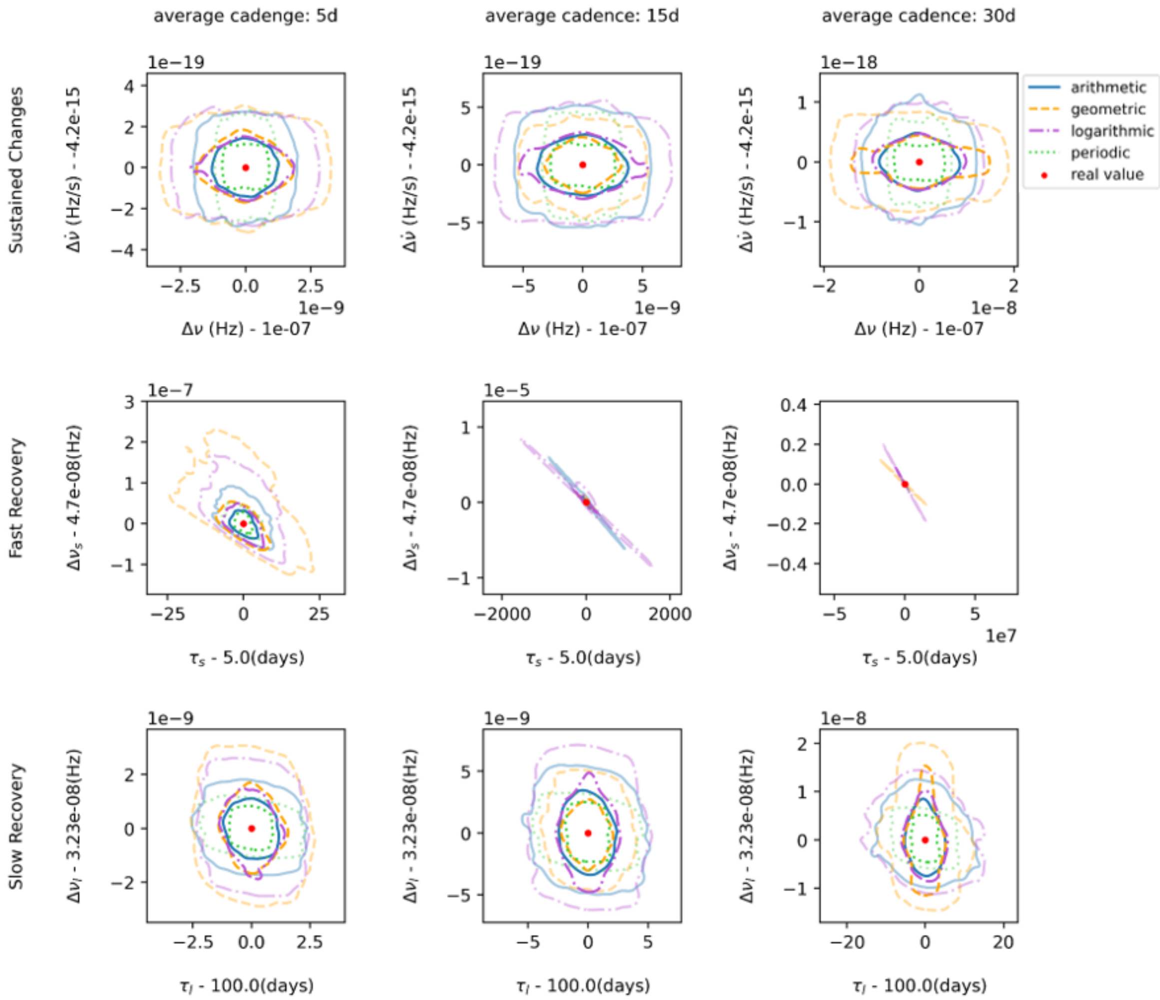

Figure 4.4: Contour plots showing the distribution of retrieved glitch parameters for glitch C, split into individual two dimensional plots separated by AC and retrieved component (i.e. base parameters, recovery parameters for short and long responses). Contour lines differ between observation strategy as shown in the legend. The inner contours contain 67% of points and outer contours contain 99% of points. The shape of the underlying point scatter is maintained.

observations. Specific to glitch C we expect there to be more than one unique solution: exponential responses are non-orthogonal to one another and are therefore inherently correlated.

The correlation between parameters τ_s and $\Delta\nu_s$ exhibits some unexpected structure manifesting as a harsh diagonal. On one side of this line there are no data points indicating some limit on either of these two parameters, having a knock-on effect to those which are most correlated. It is useful to note that the minimum value (i.e. a point which lies on this line) of $\Delta\nu_s$ when $\tau_s = \tau_{s,\mathrm{true}}$ is perfectly equal to 0, suggesting that $\Delta\nu_s$ cannot be negative in software. However, it is seen that for other values of τ_s, negative values of $\Delta\nu_s$ are allowed. The gradient of line is interestingly also equal to the magnitude of $\Delta\nu_s$ (in the units of the axis). Due to time constraints, we simply decide that this emergent property is a result of limits in magnitudes on TEMPO2 since it has no major effects on the results or other correlations, but note that it requires further investigation before future experimentation.

5 Discussion

The most obvious point to mention is that across the simpler glitches, glitch A and glitch B, there is little to no discernible difference between the quality of retrieved parameters and the spread (standard deviation) of the simulations at high cadence. This tells us that in these scenarios, there is little reason not to observe with a non-periodic observation cadence, as to avoid the degeneracy found by Dunn et al. (2021) [4]. The limit by which this finding holds appears to be lie somewhere between an average cadence of 15d and 30d, though the divergent quality of worse strategies has already begun to emerge in the 15d cases, visible in the increased

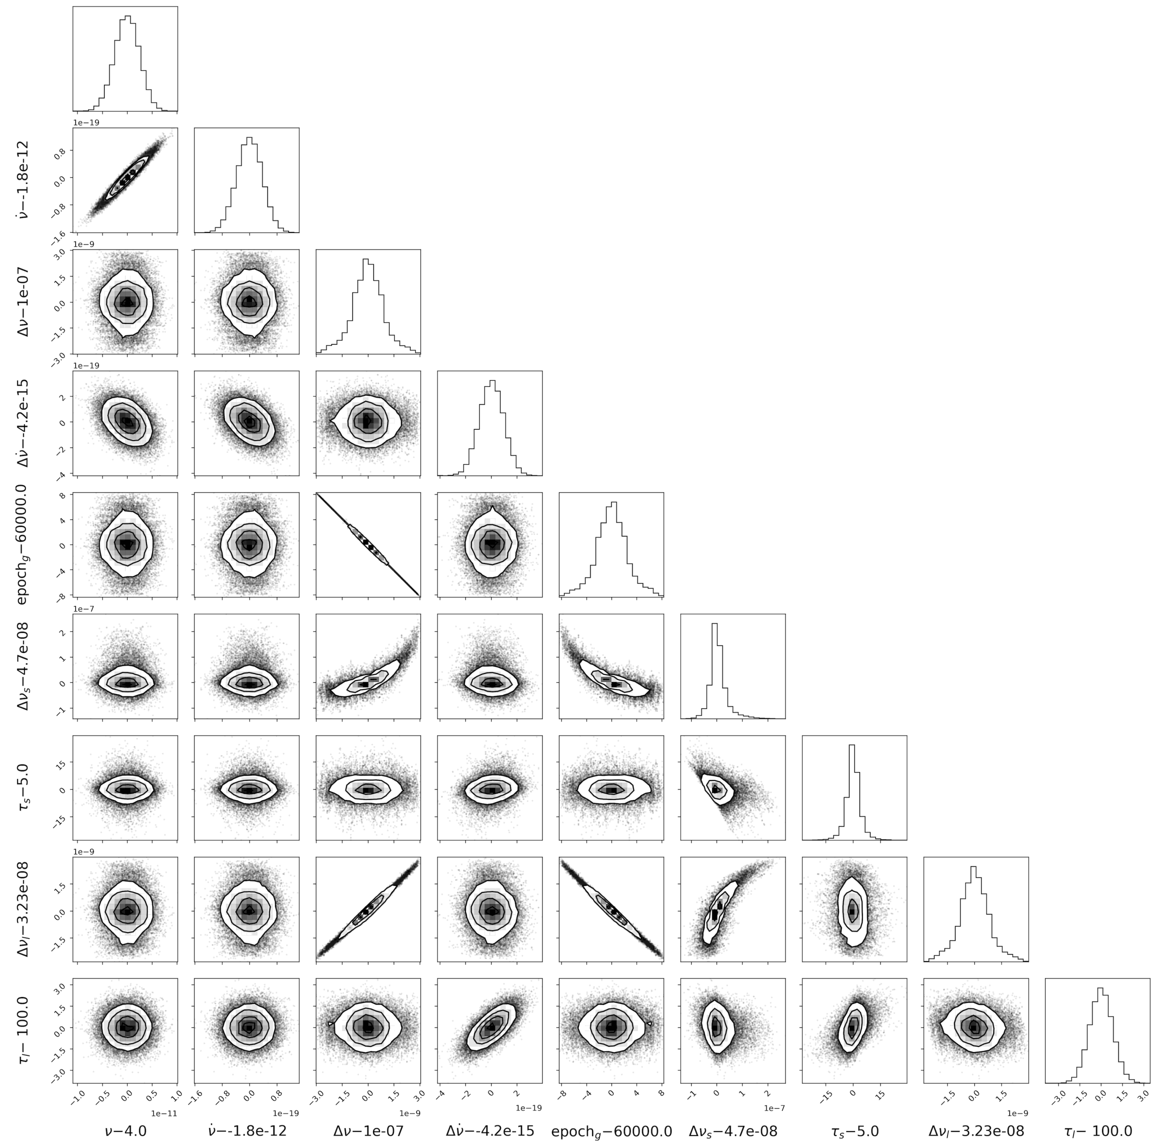

Figure 4.5: A corner plot for for glitch C collated from ∼40,000 data-points measured using all three strategies at an AC of 5. Note strong correlations between parameters $\Delta\nu$ and ν, $\Delta\nu$ and $\Delta\nu_l$; there are also strong anti-correlations in parameters $\Delta\nu$ and epoch$_g$ (t_g), $\Delta\nu_l$ and epoch$_g$ (t_g); the latter of them appearing to be analogue to the similar correlation in glitch B. Correlations with $\Delta\nu_s$ are particularly interesting, appearing to increase in strength at higher $\Delta\nu_s$.

sample standard deviation for all glitch parameters in the geometric case for glitch B. A similar conclusion can be drawn from the complex glitch C, though the strategies diverge in quality much sooner.

In the less frequent observations of glitch C, the results retrieved for the timescale and magnitude of fast recovery as shown in Table 4.2 are profoundly wrong. Understandably, when a glitch has the ability to occur and fully decay between two observations, any fitting algorithm will simply be making guesses. We do note that the results themselves are still centred around the correct value, but we put this up to one main reason: the value we feed to the fitting algorithm is already the correct timescale as mentioned in Section 3.3. This assumption means that even random guesses away from the value would, when averaged, return the correct parameter. The fitting algorithm attempts to find a solution with a smaller χ^2_R by taking steps in increasing or decreasing τ_s; Sometimes as a result of parameter correlations there will be a suitable alternative minima nearby, but on average the guess already lies on the minimum point. There is a case to be made that in further experimentation, giving only an estimate of the timescale to the fitting algorithm, it would reasonable to assume that even in extremely infrequent observation cases, short timescale response activity should still be identifiable some of the time. As mentioned in Section 3.2, in strategies which exhibit periodic behaviour, there are good and bad regions of higher and lower local cadence: a glitch occurring during a higher cadence time may have its shorter exponential parameters resolvable. This being said, a glitch which occurs near a periodic observation is similarly more easily estimated.

Parameter	AC	Arithmetic	Geometric	Logarithmic	Periodic
$\langle\Delta\nu\rangle - \Delta\nu_{\text{true}}$ ($\times10^{-10}$)	5	0 ± 8	0 ± 12	0 ± 11	0 ± 5
	15	0 ± 23	0 ± 20	0 ± 29	0 ± 16
	30	1 ± 47	-1 ± 78	1 ± 54	1 ± 31
$\langle\Delta\dot{\nu}\rangle - \Delta\dot{\nu}_{\text{true}}$ ($\times10^{-20}$)	5	0 ± 10	0 ± 10	0 ± 9	0 ± 9
	15	-1 ± 18	0 ± 14	0 ± 18	0 ± 16
	30	0 ± 32	0 ± 27	0 ± 32	0 ± 26
$\langle\Delta\nu_s\rangle - \Delta\nu_{s,\text{true}}$ ($\times10^{-8}$)	5	0 ± 2	1 ± 5	1 ± 4	0 ± 2
	15	0 ± 110	4 ± 30	0 ± 170	2 ± 11
	30	$\sim 0 \pm 5 \times 10^5$	$\sim 0 \pm 5 \times 10^6$	$\sim 0 \pm 4 \times 10^6$	0 ± 1000
$\langle\tau_s\rangle - \tau_{s,\text{true}}$	5	0 ± 3	0 ± 6	0 ± 4	0 ± 2
	15	-2 ± 170	0 ± 54	3 ± 309	0 ± 21
	30	$\sim 0 \pm 4 \times 10^5$	$\sim 0 \pm 8 \times 10^6$	$\sim 0 \pm 3 \times 10^6$	0 ± 1300
$\langle\Delta\nu_l\rangle - \Delta\nu_{l,\text{true}}$ ($\times10^{-11}$)	5	-1 ± 69	0 ± 110	0 ± 99	1 ± 48
	15	3 ± 202	4 ± 180	10 ± 264	5 ± 14
	30	30 ± 430	70 ± 710	40 ± 500	20 ± 290
$\langle\tau_l\rangle - \tau_{l,\text{true}}$ ($\times10^{-1}$)	5	0 ± 8	0 ± 9	0 ± 8	0 ± 8
	15	0 ± 16	0 ± 13	0 ± 17	0 ± 15
	30	0 ± 39	0 ± 33	0 ± 43	0 ± 30

Table 4.2: A table showing the average distance from the true values retrieved by TEMPO2 when fitting glitch B. Errors quoted are standard deviations of each sample of $\sim 10,000$ points. It is expected that the the mean difference is around 0 and the standard deviation represents the spread of the data, varying massively in some cases. True values can be found in Table 3.2. Note that values are scaled for comparison with others in the same table, for comparison to retrieved parameters in Table 4.2 adjustments must be made as some parameters are magnitudes larger. Values are rounded to similar precision, but no more than two significant figures (on the standard deviation).

Realistically, in the cases where there are two exponential recovery components present in a glitch, but the observation cadence is infrequent enough where in many cases a the short component is not visible, no human would decide to employ a fitting algorithm to fit for more than one component unless they had further knowledge or justification to do so*. Additionally, the inclusion of a quick response in low cadence scenarios can "confuse" various fitting algorithms. TEMPO2 does not apply restrictions on most fitting values (such as ensuring that timescale is always positive) and so it will make spurious claims that are unlikely to be true. In glitch C, this was most commonly exhibited as $\Delta\nu_s$ having a tremendous *negative* magnitude with an exceptionally long timescale. Through correlation, to counteract, $\Delta\nu$ was equally mis-estimated in such scenarios. This caused large deviations above what was expected from the glitch B case. In further experimentation, false solutions out by many magnitudes could be refitted with fewer exponentials, or the faster response parameters may be nudged towards more sensible solutions as a human might in manual glitch fitting.

Armed with this knowledge of unaccounted complexity in glitch C, parameter retrieval on values relating to the shorter timescale should be ignored, and the quality of all other results cannot be compared to their glitch B counterpart. However, as this phenomena emerged to an extent for all strategies, they can and should still be compared to one another, particularly at the frequent observation case. It is clear from the contour plot, Figure 4.4 that the parameter estimation in $\Delta\dot{\nu}$ is better than the same for $\Delta\nu$ when comparing across all strategies. Strategies such as geometric and logarithmic appear to spread much more widely in their distributions on the contours and values of standard deviation, with periodic and arithmetic observing taking the lead. A similar sentiment was echoed in the results of glitch B to a lesser extent, potentially suggesting that the value of $T_{\max}$, i.e. the maximum gap between observations, plays a more important role than expected. We can extend this interpretation to the 30 AC case, where the value of $T_{\max}$ follows in increasing order: periodic, arithmetic, logarithmic and geometric. The data set of average cadence of 30 days is the only set of simulations where the value of $T_{\max}$ is different between all strategies and it follows that for nearly all parameters across glitch B and C in the 30d AC case, the standard deviations also follow in worsening order: periodic, arithmetic, logarithmic and geometric. The only parameters estimated which do not adhere to this observation are $\Delta\dot{\nu}$ for glitch B and C, and τ_l in glitch C. In these outliers all strategies appear to perform somewhat similarly.

It must be said that it is not unjustifiable for $T_{\max}$ to play a larger role in data spread than initially thought. Altering the value of $T_{\max}$ changes the probability of a glitch landing in a good or bad region of observation. With reduced probability for high local cadence observations to coincide with the time of the glitch, t_g, the probability for TEMPO2 to estimate glitch parameters more accurately to the true values is reduced. It is unfortunate that, as a result of the strategy definitions, it is particularly difficult and sometimes impossible to

*Such as prior glitches exhibiting predictable behaviour.

find values of k at constant $T_{\max}$ which simultaneously allow for all strategies to offer equal average observations per unit time and suitably showcase the uniqueness of the strategy[†]. As is seen in Table 3.1, it can be done for AC $= 5, 15$d in the geometric and logarithmic cases. However, in all cases, periodic always has a value of $T_{\max} = k_p$ and arithmetic must always have a reduced $T_{\max}$ such that its AC can be comparable to the other strategies.

Looking again at parameter estimation on $\Delta\dot{\nu}$ and τ_l, this time at medium cadences on glitch C, Table 4.2 shows that geometric observing sometimes even exceeds the abilities of periodic observations in these parameters. This is inverse to the glitch B case, where the geometric strategy was worse in all cases across all cadences, suggesting that there is a link between strategy quality and glitch complexity. In the glitch B case regarding the same parameters, it can be seen that they are the only two for which geometrically spaced observations are able to maintain some degree of accuracy (notice in $\Delta\nu$ and τ_d, the geometric strategy estimations exhibit a spread up to $\sim 10^2$ times larger than periodic or arithmetic. This might suggest a given strategy having preference towards specific parameters, having a knock-on effect to the quality in retrieval of any correlated values as described in the corner plots: Figure 4.2 and 4.5.

6 Conclusion

We can draw several conclusions from this experiment. All of the following must assume that our assumptions hold.

1. At a high average cadence of observations, there is some evidence to suggest that the degeneracy in glitch parameters detailed by Dunn et al. (2021) [4] can be avoided by using an observation strategy alternative to periodic, such as logarithmic, geometric or arithmetic; with minimal to no impact on the quality of retrieved parameters.
2. There is additionally some evidence pointing towards a link between the maximum allowed observation gap and quality of estimated parameters, consistently allowing periodic and geometric, the two strategies with the lowest maximum observation wait time, to retrieve the best glitch parameters most consistently.
3. As the average cadence of observations becomes less frequent, it can be seen that in all cases presented, a periodic observation strategy will retrieve the most accurate glitch parameters from TEMPO2. We conclude that this is as a result of the glitch only being a maximum of half the TOA gap away from the closest TOA, whereas with other strategies this distance could be greatly inflated some of the time.
4. In some scenarios, geometrically distributed observations outperform periodically distributed observations in some parameter estimation, suggesting the existence of potential "preferred parameters" for a given strategy.

Further investigation on all of the above points is being worked on in the coming year (2025) as we begin to try and reduce the number of assumptions we make, allowing the ability to discern if the conclusions we have made thus far are emergent properties of the points stated in Section 3.3, or if they are properties which hold in real pulsar astronomy. For instance, investigation into the potential existence of correlations to $T_{\max}$ could be achieved with the introduction of a new observation strategy: random. Random observations within a given time-frame would average to a specific determined cadence; the maximum distance between two successive observations similarly would be governed by a $T_{\max}$.

Most importantly, we next want to reduce the number of assumptions we make. The harshest assumptions we make in Section 3.3 which limit the significance of any conclusion we might be led to assert are the non-existence of timing noise and allowing perfect pulse numbering. Working with timing noise is a process intrinsic to pulsar astronomy and removing it completely, though justified in some scenarios, severely limits the usefulness of conclusions. Many pulsars exhibit timing noise, particularly those which are seen to glitch somewhat frequently [23]. The assumption of known pulse numbers removes some degree of ambiguity in solutions. A more realistic scenario is where we are not 100% sure if a pulse or number of pulses have been unaccounted for between two observations, and so increases the complexity of fitting and algorithms required to automate fitting across many differently setup pulsars.

Removing constricting assumptions like these means we must remove any assumptions we have made as a result of them. Timing noise being absent allowed us to reasonably believe that bigger and smaller glitches (within reason) would bring us to the same or similar conclusions as those from the glitches we investigated. This would not hold in a scenario where different sizes of glitches have differing probabilities of being completely lost or unidentifiable within the timing noise. There are many such cases where real glitches have gone unnoticed for years.

[†]Some values of k which follow rules set out can be poor representations of a strategy. For instance a sufficiently large geometric constant may provide the desired average cadence, but by only alternating between two values: ΔT_{start} and $k_g \Delta T_{\text{start}}$, for the next value in the sequence, $k_g^2 \Delta T_{\text{start}}$, would exceed $T_{\max}$.

In real pulsar astronomy, it is often the case where cadence is increased in the time following a glitch in order to re-establish an accurate enough model. Simulating similar behaviour would be more accurate and opens up the possibility of differing pre and post-glitch ideal strategies. Similar to following the discovery of a new pulsar, reducing cadence with time may be a useful tool in observing many tracked and glitching pulsars.

If strategies do indeed have preferred parameters for which they can estimate more effectively, there exists motivation to investigate some of the mathematical reasoning behind these links. Strategies with small variations from each other should be trialled at constant AC, such as with different values of k. There may even be some correlation to the direction in which the strategy is sampled: what happens if a strategy halves its observation gap after each measure before resetting at some minimum TOA gap? Similarly, there exist equations not described here for cadences which vary sinusoidally or following some other function. There are many other avenues to explore alternative observation strategies.

Bibliography

[1] F. Weber. "Quark matter in neutron stars". In: *Journal of Physics G: Nuclear and Particle Physics* 25.9 (Sept. 1999), R195. DOI: 10.1088/0954-3899/25/9/201. URL: https://dx.doi.org/10.1088/0954-3899/25/9/201.

[2] A. G. Lyne et al. "45 years of rotation of the Crab pulsar". In: *Monthly Notices of the Royal Astronomical Society* 446.1 (Nov. 2014), pp. 857–864. ISSN: 0035-8711. DOI: 10.1093/mnras/stu2118. URL: http://dx.doi.org/10.1093/mnras/stu2118.

[3] A. Basu et al. "Observed glitches in eight young pulsars". In: *Monthly Notices of the Royal Astronomical Society* 491.3 (Nov. 2019), pp. 3182–3191. ISSN: 0035-8711. DOI: 10.1093/mnras/stz3230. eprint: https://academic.oup.com/mnras/article-pdf/491/3/3182/31490452/stz3230.pdf. URL: https://doi.org/10.1093/mnras/stz3230.

[4] L. Dunn et al. "Effects of periodicity in observation scheduling on parameter estimation of pulsar glitches". In: *Monthly Notices of the Royal Astronomical Society* 504.3 (Apr. 2021), pp. 3399–3411. ISSN: 1365-2966. DOI: 10.1093/mnras/stab1097. URL: http://dx.doi.org/10.1093/mnras/stab1097.

[5] F. E. Marshall et al. "The Big Glitcher: The Rotation History of PSR J0537-6910". In: *The Astrophysical Journal* 603.2 (Mar. 2004), pp. 682–689. DOI: 10.1086/381567.

[6] A. G. Lyne and F. Graham-Smith. *Pulsar astronomy.* eng. 4th edition. Cambridge astrophysics ; 48. Cambridge ; Cambridge University Press, 2012. ISBN: 9781107010147.

[7] R. Wijnands and M. van der Klis. "A millisecond pulsar in an X-ray binary system". In: *Nature* 394.6691 (July 1998), pp. 344–346. ISSN: 1476-4687. DOI: 10.1038/28557. URL: https://doi.org/10.1038/28557.

[8] Q. Liang et al. *A Test of Gravity with Pulsar Timing Arrays.* 2023. arXiv: 2304.02640 [astro-ph.CO]. URL: https://arxiv.org/abs/2304.02640.

[9] R. Ruderman. "Neutron Star Crustal Plate Tectonics. II. Evolution of Radio Pulsar Magnetic Fields". In: *The Astrophysical Journal* 382 (Dec. 1991), p. 576. DOI: 10.1086/170744.

[10] F. Crawford and M. Demiański. "A Comparison of Measured Crab and Vela Glitch Healing Parameters with Predictions of Neutron Star Models". In: *The Astrophysical Journal* 595.2 (Oct. 2003), pp. 1052–1057. DOI: 10.1086/377470. arXiv: astro-ph/0306242 [astro-ph].

[11] S. Zhou et al. *Pulsar Glitches: A Review.* 2022. arXiv: 2211.13885 [astro-ph.HE]. URL: https://arxiv.org/abs/2211.13885.

[12] D. Antonopoulou et al. "Pulsar glitches: observations and physical interpretation". In: *Reports on Progress in Physics* 85.12 (Dec. 2022), p. 126901. DOI: 10.1088/1361-6633/ac9ced. URL: https://dx.doi.org/10.1088/1361-6633/ac9ced.

[13] D. R. Lorimer and M. Kramer. *Handbook of pulsar astronomy.* eng. Cambridge observing handbooks for research astronomers ; 4. Cambridge: Cambridge University Press, 2005. ISBN: 0521828236.

[14] J. Wang et al. "A comparative analysis of pulse time-of-arrival creation methods". In: *Astronomy & Astrophysics* 658 (Feb. 2022), A181. ISSN: 1432-0746. DOI: 10.1051/0004-6361/202141121. URL: http://dx.doi.org/10.1051/0004-6361/202141121.

[15] J. H. Taylor. "Pulsar Timing and Relativistic Gravity". In: *Philosophical Transactions of the Royal Society of London Series A* 341.1660 (Oct. 1992), pp. 117–134. DOI: 10.1098/rsta.1992.0088.

[16] G. B. Hobbs et al. "tempo2, a new pulsar-timing package - I. An overview: tempo2, a new pulsar-timing package - I. Overview". In: *Monthly Notices of the Royal Astronomical Society* 369.2 (May 2006), pp. 655–672. ISSN: 0035-8711. DOI: 10.1111/j.1365-2966.2006.10302.x. URL: http://dx.doi.org/10.1111/j.1365-2966.2006.10302.x.

[17] E. Gügercinoğlu et al. "Glitches in four gamma-ray pulsars and inferences on the neutron star structure". In: *Monthly Notices of the Royal Astronomical Society* 511.1 (Jan. 2022), pp. 425–439. ISSN: 0035-8711. DOI: 10.1093/mnras/stac026. eprint: https://academic.oup.com/mnras/article-pdf/511/1/425/42363150/stac026.pdf. URL: https://doi.org/10.1093/mnras/stac026.

[18] P. V. Padmanabh et al. "The MPIfR–MeerKAT Galactic Plane Survey – I. System set-up and early results". In: *Monthly Notices of the Royal Astronomical Society* 524.1 (June 2023), pp. 1291–1315. ISSN: 1365-2966. DOI: 10.1093/mnras/stad1900. URL: http://dx.doi.org/10.1093/mnras/stad1900.

[19] P. Weltevrede et al. "The glitch-induced identity changes of PSR J1119-6127: The glitch-induced identity changes of PSR J1119-6127". In: *Monthly Notices of the Royal Astronomical Society* 411.3 (Dec. 2010), pp. 1917–1934. ISSN: 0035-8711. DOI: 10.1111/j.1365-2966.2010.17821.x. URL: http://dx.doi.org/10.1111/j.1365-2966.2010.17821.x.

[20] C. M. Espinoza et al. "A study of 315 glitches in the rotation of 102 pulsars: A study of 315 glitches in 102 pulsars". In: *Monthly Notices of the Royal Astronomical Society* 414.2 (Apr. 2011), pp. 1679–1704. ISSN: 0035-8711. DOI: 10.1111/j.1365-2966.2011.18503.x. URL: http://dx.doi.org/10.1111/j.1365-2966.2011.18503.x.

[21] G. Hobbs et al. “Long-term timing observations of 374 pulsars”. In: *Monthly Notices of the Royal Astronomical Society* 353.4 (Oct. 2004), pp. 1311–1344. ISSN: 0035-8711. DOI: 10.1111/j.1365-2966.2004.08157.x. eprint: https://academic.oup.com/mnras/article-pdf/353/4/1311/3000980/353-4-1311.pdf. URL: https://doi.org/10.1111/j.1365-2966.2004.08157.x.

[22] C. M. Espinoza et al. “Neutron star glitches have a substantial minimum size”. In: *Monthly Notices of the Royal Astronomical Society* 440.3 (Apr. 2014), pp. 2755–2762. ISSN: 1365-2966. DOI: 10.1093/mnras/stu395. URL: http://dx.doi.org/10.1093/mnras/stu395.

[23] G. Hobbs et al. “An analysis of the timing irregularities for 366 pulsars”. In: *Monthly Notices of the Royal Astronomical Society* 402.2 (Feb. 2010), pp. 1027–1048. ISSN: 1365-2966. DOI: 10.1111/j.1365-2966.2009.15938.x. URL: http://dx.doi.org/10.1111/j.1365-2966.2009.15938.x.

[24] J. W. McKee et al. “A glitch in the millisecond pulsar J0613-0200”. In: *Monthly Notices of the Royal Astronomical Society* 461.3 (June 2016), pp. 2809–2817. ISSN: 0035-8711. DOI: 10.1093/mnras/stw1442. eprint: https://academic.oup.com/mnras/article-pdf/461/3/2809/8107796/stw1442.pdf. URL: https://doi.org/10.1093/mnras/stw1442.

[25] D. Foreman-Mackey. “corner.py: Scatterplot matrices in Python”. In: *The Journal of Open Source Software* 1.2 (June 2016), p. 24. DOI: 10.21105/joss.00024. URL: https://doi.org/10.21105/joss.00024.

[26] L. Lentati et al. “temponest: a Bayesian approach to pulsar timing analysis”. In: *Monthly Notices of the Royal Astronomical Society* 437.3 (Dec. 2013), pp. 3004–3023. ISSN: 1365-2966. DOI: 10.1093/mnras/stt2122. URL: http://dx.doi.org/10.1093/mnras/stt2122.