The Astrophysical Journal, 864:137 (8pp), 2018 September 10 https://doi.org/10.3847/1538-4357/aad6eb
© 2018. The American Astronomical Society. All rights reserved.

The Enigmatic Spin Evolution of PSR J0537–6910: r-modes, Gravitational Waves, and the Case for Continued Timing

N. Andersson[1], D. Antonopoulou[2], C. M. Espinoza[3], B. Haskell[2], and W. C. G. Ho[4,5]
[1] Mathematical Sciences and STAG Research Centre, University of Southampton, Southampton SO17 1BJ, UK
[2] Nicolaus Copernicus Astronomical Center, Polish Academy of Sciences, ul. Bartycka 18, 00-716 Warsaw, Poland
[3] Departamento de Física, Universidad de Santiago de Chile, Estación Central, Santiago 9170124, Chile
[4] Mathematical Sciences, Physics and Astronomy & STAG Research Centre, University of Southampton, Southampton SO17 1BJ, UK
[5] Department of Physics and Astronomy, Haverford College, 370 Lancaster Avenue, Haverford, PA 19041, USA
Received 2017 November 20; revised 2018 July 27; accepted 2018 July 28; published 2018 September 10

Abstract

We discuss the unique spin evolution of the young X-ray pulsar PSR J0537–6910, a system in which the regular spin down is interrupted by glitches every few months. Drawing on the complete timing data from the *Rossi X-ray Timing Explorer* (from 1999 to 2011), we argue that a trend in the interglitch behavior points to an effective braking index close to $n = 7$, which is much larger than expected. This value is interesting because it would accord with the neutron star spinning down due to gravitational waves from an unstable r-mode. We discuss to what extent this, admittedly speculative, scenario may be consistent and if the associated gravitational-wave signal would be within reach of ground-based detectors. Our estimates suggest that one may, indeed, be able to use future observations to test the idea. Further precision timing would help to enhance the achievable sensitivity, and we advocate a joint observing campaign between the *Neutron Star Interior Composition Explorer* and the LIGO-Virgo network.

Key words: gravitational waves – pulsars: individual (PSR J0537–6910) – stars: neutron – X-rays: stars

1. From Observation to Speculation

The young X-ray pulsar PSR J0537–6910 in the Large Magellanic Cloud (associated with the supernova remnant N157B; Marshall et al. 1998) is an intriguing object. Spinning at a frequency of 62 Hz, this is the fastest spinning and most energetic nonrecycled neutron star. It also exhibits abrupt spin ups (glitches) roughly every 100 days (Middleditch et al. 2006). Recent work analyzed this glitch activity (Ferdman et al. 2018; Antonopoulou et al. 2018), drawing on the complete timing data from the *Rossi X-ray Timing Explorer (RXTE)*. The results highlight the (almost) predictable regularity of the glitches, the overall (glitch-dominated) spin evolution, and the interglitch behavior.

The analyses by Marshall et al. (2004), Antonopoulou et al. (2018), and Ferdman et al. (2018) and also raise the issue of the braking index of this neutron star. Assuming a spin-down rate that scales as $\dot{\nu} = -C\nu^n$, the braking index, n, can be estimated as

$$n = \frac{\nu\ddot{\nu}}{\dot{\nu}^2} \tag{1}$$

where ν is the spin frequency, and dots denote time derivatives. The data set points to a negative value of $n \approx -1.2$ over the observed 13 years (Antonopoulou et al. 2018), which is a somewhat surprising result that may be related to the glitch activity. By focusing on the interglitch evolution, one would infer much larger values for the braking index (see Table 1 in Antonopoulou et al. 2018). It would be natural to assume that the observed phenomenology will ultimately be explained by the detailed nature of the glitch relaxation, involving superfluid vortex dynamics and poorly understood friction/pinning forces. However, there may be additional physics at play.

As we will discuss, the long-term postglitch relaxation hints at the system evolving toward an effective braking index of $n \approx 7$. Such a large value would suggest that the spin down of J0537–6910 is not governed by electromagnetic emission. Taking the result at face value, we are instead led to consider the possibility that the star spins down due to gravitational radiation. However, even in this case, we have to go beyond the "standard" scenario. A spinning deformed star would emit quadrupole radiation at twice the spin frequency, leading to a braking index of $n = 5$. This is not what we are looking for. There is, however, a plausible scenario that would "explain" the timing data. If the gravitational-wave driven instability of the inertial r-mode (Andersson 1998; Lindblom et al. 1998; Owen et al. 1998; Andersson & Kokkotas 2001) were to operate in the star, and the associated emission were to dominate the spin evolution, then theory predicts (for the quadrupole mode, which mainly radiates through the mass current multipoles) a braking index of exactly $n = 7$. While this may be a coincidence, and the true explanation for the spin evolution of J0537–6910 lies elsewhere, the possibility is interesting enough that it warrants a more detailed discussion.

In this paper we take a closer look at the *RXTE* timing data, with the aim of establishing to what extent the suggested braking index of $n \approx 7$ is credible and robust. Not surprisingly, the answer will be inconclusive. Next, we consider whether there is a workable r-mode scenario for this system. Again, we cannot draw definite conclusions, but our discussion highlights the parts of the theory that one might have to negotiate in order to arrive at a consistent model. This naturally leads us to the issue of future observations. We provide simple estimates that suggest that the gravitational waves associated with the r-mode scenario should be within reach of the advanced generation of interferometers. In essence, one may be able to use observations to constrain (and perhaps rule out) the presence of an unstable r-mode in this system.

In order to achieve this, one would ideally carry out a targeted search for r-mode gravitational waves from J0537–6910. This

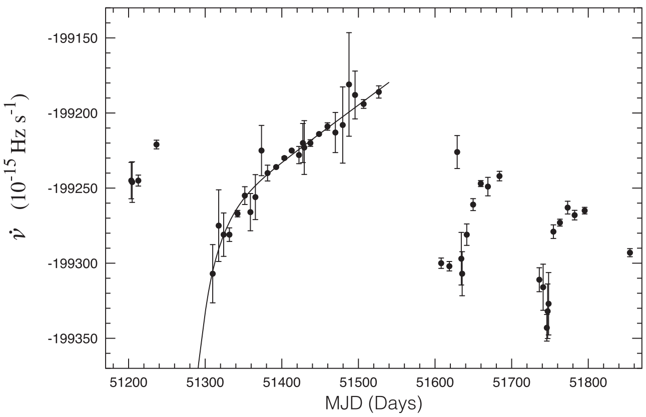

Figure 1. Illustration of the evolution of $\dot{\nu}$ through the first two years of observations, including four glitches (points are calculated as in Figure 5 of Antonopoulou et al. 2018). The solid line shows a model of the form $\dot{\nu}_0 + \ddot{\nu}_0(t - t_g) - (\Delta\nu_d/\tau_d)e^{-(t-t_g)/\tau_d}$, with the best-fitted parameters derived from a fit to all ToAs following the first glitch, which occurred near t_g = MJD 51278, and up to the next glitch (see the text for details). The linear term of this model, $\ddot{\nu}_0 = (4.4 \pm 0.1) \times 10^{-21}$ Hz s^{-2}, implies an underlying braking index of 6.8 ± 0.2.

requires reliable timing information. This is crucial, as the frequent glitches disrupt predictions made by timing models developed prior to each glitch. However, the last *RXTE* timing observations of J0537–6910 took place in 2011. This means that the system has not yet been considered in targeted LIGO searches in the advanced detector era (see Aasi et al. 2014 for the most sensitive search, at twice the spin frequency, and Abbott et al. 2017a for the best current results for other pulsars). Given that simultaneous gravitational-wave and X-ray observations would enable the most effective gravitational-wave search, we argue that there is a strong case for the advanced LIGO-Virgo network to join forces with the Neutron Star Interior Composition Explorer (NICER), an X-ray telescope recently installed on the *International Space Station* (Gendreau et al. 2016). NICER is optimized for detecting pulsations from neutron stars, such as J0537−6910, has twice the collecting area of *XMM-Newton,* and a timing accuracy of 100 ns. With a careful observing strategy, NICER should be able to track the spin evolution of J0537–6910, including detecting glitches, thus allowing gravitational-wave searches to compensate for the complex timing behavior.

2. The Observed Rotational Evolution

PSR J0537–6910 has a unique rotational evolution: it abruptly spins up by a few ppm every few months, which is by far the highest glitch rate observed in any pulsar (Fuentes et al. 2017). The imprint of the frequent glitches in the spin-down rate is dramatic. Most spin ups, $\Delta\nu$, are accompanied by a sharp decrease in $\dot{\nu}$ and a subsequent recovery characterized by a large, positive $\ddot{\nu}$ (see Figure 1 for a sample of the data). The slope of this saw-tooth-like pattern over the course of 13 years returns a long-term negative $\ddot{\nu}$ and an apparent braking index of $n = -1.22 \pm 0.04$ (Antonopoulou et al. 2018).

It is far from trivial to differentiate the effect of glitches from the underlying braking mechanism. This is mainly because the time intervals between glitches are short (only a few months), while the timescales associated with the superfluid response to glitches might be of the order of years (Haskell & Antonopoulou 2014). Therefore, we cannot guarantee that the internal superfluid has reached its equilibrium state at the time of a given $\ddot{\nu}$ measurement.

Between glitches, it is possible to carry out a phase-coherent timing analysis to derive the rotational parameters. The rotational phase of the pulsar, $\phi(t)$, can be described by a truncated Taylor series around an epoch, t_0,

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

where ϕ_0, ν_0, $\dot{\nu}_0$, and $\ddot{\nu}_0$ are the reference phase, spin frequency, and its first two time derivatives, respectively. When this simple timing model is fitted to entire interglitch intervals, the inferred braking indices (calculated as $n = \nu_0\ddot{\nu}_0/\dot{\nu}_0^2$) are typically greater than 10. Such large values of n most likely reflect the early response to the glitch. The analysis of Antonopoulou et al. (2018) shows that $\ddot{\nu}_0$ (or equivalently, n) tends to be smaller for longer interglitch intervals, as data further away from a glitch start to dominate the fits.

In fact, following the first (and largest observed) glitch, a simple analysis gives an "average" braking index of $n = 7.6 \pm 0.1$ (see Table 1 in Antonopoulou et al. 2018). This particular interglitch time interval is much longer than any other (~284 days, while all others are less than 200 days), and an exponential relaxation on a relatively short, 20 day timescale was observed. The estimate of n in Antonopoulou et al. (2018) was derived by fitting only Equation (2) to data after the first 16 days postglitch (to avoid the quickly decaying initial phase) and up to the second glitch. To examine the asymptotic $\ddot{\nu}$ of this interglitch interval in more detail, we need to account for the exponentially decaying term in the timing solution (see Figure 1). We thus fit all available pulse time-of-arrivals (ToAs) between the first and second glitch with a timing model, as in Equation (2), including an exponential term. The parameters of the additional term are consistent with those in Table 3 of Antonopoulou et al. (2018), and we find the underlying braking index to be $n = 6.8 \pm 0.2$.

Given the possible connection between this value for the braking index and a scenario involving gravitational-wave emission through unstable r-modes, we want to further investigate the interglitch braking indices and their apparent softening at late times. Basically, we want to check whether a braking index close to 7 is unique to the evolution after the first glitch, or if it could be a common asymptotic behavior. The details on the observations, the derivation of the ToAs, and pulsar timing tools used in the following can be found in Kuiper & Hermsen (2015) and Antonopoulou et al. (2018).

To calculate the braking index at different moments after each glitch, we fitted Equation (2) to short segments of the data. Because $\ddot{\nu}_0$ is small, of the order of 10^{-20} Hz s^{-2}, its measurement can be significantly affected by noise (ToA uncertainties, but also "timing noise" intrinsic to the pulsar). This compromises the accuracy of the best-fitted values, especially for time intervals shorter than about 50 days or ones that contain too few ToAs. For the shortest interglitch intervals, the entire data had to be used, and only one measurement of the braking index was obtained. For the longer intervals, however, it was possible to get a few measurements at different epochs. We used segments of interglitch data spanning just below 90 days, with the precise length depending on the distribution of the ToAs, and imposed a minimum of at least five ToAs for each fit. Several values of the interglitch n

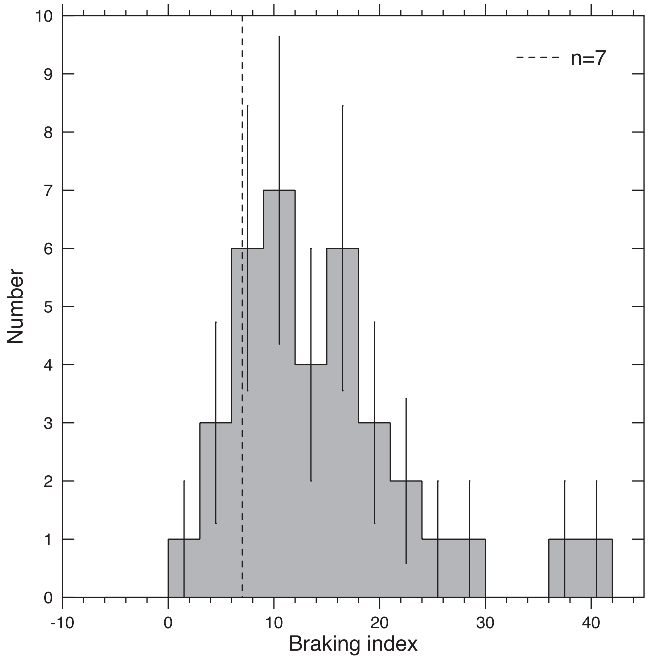

Figure 2. Braking index obtained from direct fits of ν_0, $\dot{\nu}_0$, and $\ddot{\nu}_0$ to interglitch ToAs. Only values from fits centered at least 50 days after each glitch are shown, which excludes some of the shortest interglitch intervals. For the longest interglitch intervals, segments of data $\leqslant$90 days long were used for the fit and were centered approximately 20 days apart (see the text for details and Figure 3 for the evolution of n as a function of time since the glitch). Note that higher braking index values arise mostly from intervals centered at times closer to the glitch epochs (see Figure 3).

for different times after the glitch were calculated by shifting (when possible) the fitting window by 20 days.

A histogram of the results for the braking index is shown in Figure 2. For clarity, we have excluded fits centered on ToAs that are less than 50 days away from the glitch epoch, as these are dominated by the early fast relaxation. Typically, larger values of n correspond to epochs soon after a glitch as demonstrated in Figure 3, which shows the braking index n as a function of time since the preceding glitch, $t_{\rm pg}$. This kind of plot, however, has to be considered with some caution. As reference dates, we used the MJD epoch for the 45 glitches reported in Antonopoulou et al. (2018). The errors from the uncertainty in the glitch epoch are displayed in Figure 2. Using a slightly different set of ToAs (derived from the same *RXTE* data), Ferdman et al. (2018) obtain very similar results for the glitch parameters, but in at least three cases, the glitch epochs of the two data sets are inconsistent within the error bars—errors presented in Figure 2 should thus be viewed as lower limits. Errors in the braking indices are propagated 1-sigma uncertainties on the best-fitted parameters of Equation (2), which are often underestimates. Moreover, although we believe the list of 45 glitches to be complete for spin ups larger than $\sim$1 μHz, some small events may not have been identified. For example, Ferdman et al. (2018) discovered another possible glitch at MJD 52716(1). Missing glitches may introduce an overestimate of $t_{\rm pg}$. The opposite would be true if some of the timing irregularities that we considered as glitches were in reality timing noise features.[6] These are important caveats, but the qualitative picture of the braking index evolution remains the same. There is some tendency toward a value of n close to 7 as $t_{\rm pg}$ progresses, although a robust measurement is possible only for the first glitch (the longest glitch-free interval).

A similar conclusion was reached by Ferdman et al. (2018) by following a rather different approach. They identified 42 glitches and used a subset of them to construct $\dot{\nu}$ as a function of $t_{\rm pg}$. They then combined those values to achieve a set of denser data points that they fit with a single function, $\dot{\nu}(t)$, assuming that the relaxing component of all glitches is described by a single exponential with the same amplitude and characteristic timescale. Using the best-fitted $\ddot{\nu}$ from Ferdman et al. (2018), we infer that the asymptotic value of the braking index for the interglitch time intervals is $n = 7.4 \pm 0.7$. This is consistent with the results shown in Figure 3.

These arguments suggest that an underlying $n \approx 7$, inferred for the evolution after the first glitch, might be accommodated by the entire 13 years of data. Whether this is close to the "real" long-term braking index, and thus probes the dominant braking mechanism, is open to interpretation. In the standard glitch scenario, the recovery is governed by the microphysics of the internal superfluid and can often extend for a very long time. It would thus not be surprising if most glitches were to exhibit a similar relaxation, consistent with a braking index around 7 after some time postglitch, but which would gradually decrease further as long as another glitch does not interrupt the process.

This issue could possibly be resolved by future observations of one or more glitches larger than the first event in the *RXTE* data (or at least, of similar size). Since the size of each J0537–6910 glitch strongly correlates with the time interval to the next one, larger glitches would enable us to populate the region beyond $t_{\rm pg} > 150$ days with more n measurements. This would allow us to assess the significance of the observed tendency. As such big spin ups appear to be at the higher end of the glitch size distribution for this pulsar, they are likely relatively rare. It is, therefore, important if such a glitch is observed that the postglitch relaxation is monitored closely. Frequent observations are not only needed for accurate n calculations, but will moreover provide a way to test the hypothesis in Ferdman et al. (2018) that the parameters of the exponential recovery are common to all glitches.

3. Evidence of an Unstable r-mode?

Turning to the possible explanation for the observed behavior, let us first go through the argument that leads to the braking index for a star that spins down due to an unstable r-mode. In general, the r-mode instability arises through a tug-of-war between gravitational-wave emission, which drives the instability, and various dissipation mechanisms, which saps energy from the mode. Somewhat schematically (adopting the strategy from Owen et al. 1998), the amplitude of the mode, α, evolves according to

$$\dot{\alpha} = \alpha\left(\frac{1}{t_{\rm gw}} - \frac{1}{t_{\rm diss}}\right) - \frac{N}{2I\Omega}. \qquad (3)$$

The evolution depends on the instability growth time, $t_{\rm gw}$, and the dissipation timescale, $t_{\rm diss}$ (we take both timescales to be positive, in contrast with Owen et al. 1998), as well as any external torque, N. Letting I be the star's moment of inertia and

[6] A total of four events were flagged as ambiguous by Antonopoulou et al. (2018). We recalculated $t_{\rm pg}$ under the assumption that these timing features were not real glitches and confirmed that their inclusion does not alter the main features of Figure 3.

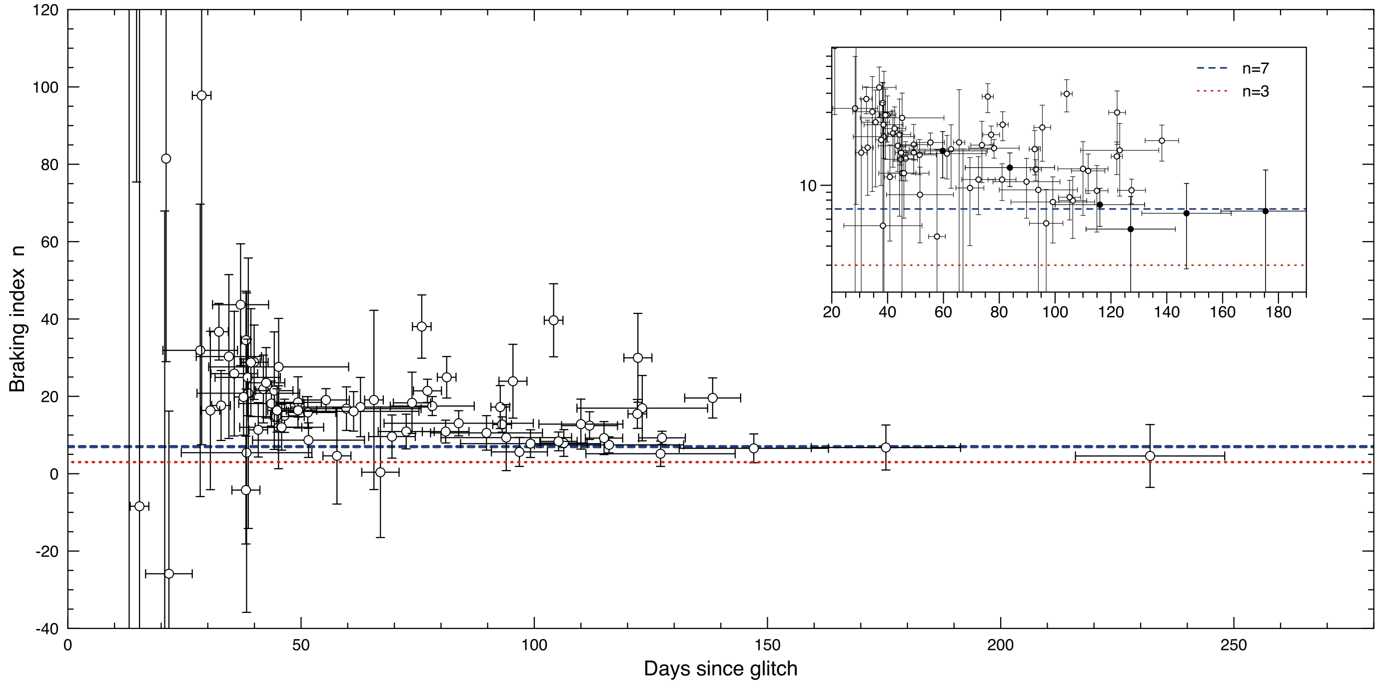

Figure 3. Braking index, n, as a function of time, t_{pg}, since the preceding glitch. The data points are obtained from direct fits of Equation (2) over a sliding window of up to 90 days which was moved forward by 20 days at each step. The dashed horizontal lines indicate values of $n = 3$ (as expected for a spin down dominated by electromagnetic dipole radiation) and $n = 7$ (which would apply in the scenario explored in this paper). The insert presents a zoomed-in plot of the same data on a logarithmic scale, with the data points corresponding to fits of the first interglitch interval highlighted in black.

Ω be the angular spin frequency (we assume uniform rotation for simplicity), we also have

$$\dot{\Omega} = -\frac{2Q\Omega\alpha^2}{t_{\mathrm{diss}}} + \frac{N}{I} \tag{4}$$

where Q is an equation-of-state-dependent quantity. In this phenomenological model, the unstable r-mode is assumed to grow exponentially until it reaches a given saturation amplitude, α_s. From Equation (3), we see that if α_s takes a fixed value, then (assuming that we can ignore external torques) we must have $t_{\mathrm{diss}} \approx t_{\mathrm{gw}}$. It then follows from Equation (4) that the spin frequency $\nu = \Omega/2\pi$ evolves according to

$$\dot{\nu} = -\frac{2Q\alpha_s^2\nu}{t_{\mathrm{gw}}}. \tag{5}$$

If we, in order to keep things simple enough that the scalings with the star's mass (M) and radius (R) are apparent, consider an $n = 1$ polytrope,[7] then the growth time for the $l = m = 2$ r-mode is given by (all r-mode timescale estimates are taken from the review by Andersson & Kokkotas 2001)

$$t_{\mathrm{gw}} \approx 5 \times 10^7 \left(\frac{M}{1.4M_\odot}\right)^{-1} \left(\frac{R}{10\ \mathrm{km}}\right)^{-4} \left(\frac{\nu}{100\ \mathrm{Hz}}\right)^{-6}\ \mathrm{s}. \tag{6}$$

[7] It is worth noting that the different instability timescales are only weakly dependent on the equation of state (Owen 2010; Idrisy et al. 2015), so it seems reasonable to base our discussion on the estimates for polytropes, which have the advantage that the scaling with the star's mass and radius are explicit.

Noting that we have $Q \approx 9.4 \times 10^{-2}$ (Owen et al. 1998), it follows that

$$\dot{\nu} \approx -4 \times 10^{-7} \alpha_s^2 \left(\frac{M}{1.4M_\odot}\right) \left(\frac{R}{10\ \mathrm{km}}\right)^4 \left(\frac{\nu}{100\ \mathrm{Hz}}\right)^7\ \mathrm{s}^{-2}, \tag{7}$$

and we arrive at the braking index $n = 7$, as long as α_s is constant.

The theory prediction thus accords with the behavior inferred from the interglitch evolution of J0537–6910. This is an interesting observation, but it does not mean that we are done. We need to check to what extent this explanation is consistent. This involves considering poorly known aspects of neutron star interior physics, but we should nevertheless be able to make some progress. At each step, we need to be mindful of the assumptions. So far, we have (i) ignored external torques, which means that the gravitational-wave emission dominates the spin down; (ii) assumed that the unstable r-mode has reached saturation; and (iii) that this corresponds to a constant amplitude,[8] α_s.

Given these two assumptions, we can combine Equation (7) with the observed spin parameters for the system. Once we consider a specific neutron star model, for which we can (at least in principle) work out the instability growth time, we can turn the data into a statement about the required saturation amplitude.

[8] The assumption of a constant saturation amplitude is an obvious simplification. More detailed work, building on the idea that the instability saturates due to the nonlinear coupling between the r-mode and the sea on shorter scale inertial modes, suggests a much more complicated (possibly chaotic) evolution of the amplitude (Bondarescu et al. 2007, 2009). Nevertheless, on average, these evolutions stay fairly close to those of Owen et al. (1998), which assumed a fixed α_s. Hence, our assumption is a natural starting point.

First of all, the estimated growth time from Equation (6) immediately tells us that, for a "canonical" neutron star with $R = 10$ km and $M = 1.4\,M_\odot$ spinning at the observed $\nu = 62$ Hz, the instability growth time would be $t_{\rm gw} \approx 30$ years. That is, the mode would not quickly regrow if it were disrupted by the frequent glitches. We also see that we need to keep an eye on the relatively high power of the stellar radius. If we consider the current radius constraint from X-ray observations[9] $R = 10$–14 km (Steiner et al. 2018), then the growth time would be shorter by about a factor of 4 for the largest neutron stars.

Next, making use of the observed $\dot\nu = 1.99 \times 10^{-10}$ Hz s^{-1}, we see that an r-mode-dominated spin down requires

$$\alpha_s \approx 0.12\left(\frac{M}{1.4 M_\odot}\right)^{-1/2}\left(\frac{R}{10\ \mathrm{km}}\right)^{-2}. \quad (8)$$

If we take the radius to be 14 km, then the required amplitude is about a factor of 2 smaller.

If we assume that the star has been spinning down according to Equation (7) throughout most of its history, and that it was initially spinning much faster than it is today, we can estimate the age. This way, we find that the evolution to the current spin rate, ν, takes place on a timescale

$$t_{\rm sd} \approx 4.2 \times 10^9\left(\frac{\alpha_s}{0.1}\right)^{-2}\left(\frac{\nu}{100\ \mathrm{Hz}}\right)^{-6}\ \mathrm{s}. \quad (9)$$

If we let α_s have the predicted value, then it would take 1600–6000 years for a star with a radius in the range 10–14 km to reach the current spin rate. This is in good agreement with the estimated age of the supernova remnant, which is 1–5000 years (Wang & Gotthelf 1998; Chen et al. 2006).

However, given what we think we know about the mechanism that determines the r-mode saturation, the estimated α_s is uncomfortably large. The nonlinear coupling between the large scale r-mode and the sea of shorter wavelength inertial modes is expected to lead to saturation at $\alpha_s < 10^{-2}$ (Arras et al. 2003). The best estimates of the saturation level (involving a huge number of mode couplings) suggest that α_s is unlikely to reach 10^{-3} (Brink et al. 2004a, 2004b, 2005), and the actual level could be quite a bit lower. This expectation is also brought out by evolutions involving the leading three-mode couplings (Bondarescu et al. 2007, 2009; Bondarescu & Wasserman 2013). Is this a fatal objection to the proposed scenario? Possibly, but one can imagine ways of reducing the tension. First, it could be that the growth timescale is shorter than the estimate in Equation (6). We know, for example, that the result for a uniform density star is about a factor of 2 smaller than Equation (6) (Andersson & Kokkotas 2001). But this is not enough to make the results consistent. Instead, we may consider the saturation mechanism. The level of saturation is expected to be close to the threshold where the nonlinear coupling between the r-mode and a pair of inertial daughter modes becomes parametrically unstable (Arras et al. 2003; Bondarescu et al. 2007, 2009; Bondarescu & Wasserman 2013). The threshold amplitude depends on the damping rates of the (supposedly stable) daughter modes and the level of frequency detuning (how close the mode frequencies are to resonance). Focusing on the former of these factors, we note that an increase in the daughter mode damping rate by some factor would affect the r-mode saturation by the same factor. In principle, a more efficient damping of short wavelength daughter modes could bring the theoretical estimate closer to the required saturation amplitude, but the analysis of Bondarescu et al. (2007, 2009) suggests that this is unlikely.

In order to discuss the issue, we need to consider the possible dissipation mechanisms that may act on an unstable r-mode. In the simplest model, the dominant dissipation is the macroscopic shear viscosity due to neutron–neutron scattering. In this case, we have a damping timescale (Andersson & Kokkotas 2001):

$$t_{\rm sv} \approx 6.7 \times 10^5\left(\frac{M}{1.4 M_\odot}\right)^{-5/4}\left(\frac{R}{10\ \mathrm{km}}\right)^{23/4}\left(\frac{T}{10^8\ \mathrm{K}}\right)^{2}\ \mathrm{s}. \quad (10)$$

In order for the star to be in the unstable regime so that the previous arguments hold, we need the damping rate to be slower than the growth rate from Equation (6). This leads to the condition:

$$T > 8.6 \times 10^8\left(\frac{M}{1.4 M_\odot}\right)^{1/8}\left(\frac{R}{10\ \mathrm{km}}\right)^{-39/8}\left(\frac{\nu}{100\ \mathrm{Hz}}\right)^{-3}\ \mathrm{K}. \quad (11)$$

The weak scaling with the star's mass means that the dependence on the radius is dominant. For the observed spin rate, $\nu = 62$ Hz, we would need $T > 3.6 \times 10^9$ K for a 10 km star. This is uncomfortably hot. However, if we take the radius to be 14 km, then we only need need $T > 7 \times 10^8$ K. As a useful comparison, we note that one would expect the star (in absence of any instability) to have cooled to a temperature of roughly $T \approx 2 \times 10^8$ K (Ho et al. 2015a). The relatively high-temperature threshold simply reflects the expectation that we may need a larger-than-anticipated r-mode instability window in order to accommodate this system.

However, the model has an internal "consistency check." The friction associated with the viscosity heats the star. After some time of evolution, one would expect this heating to balance to cooling due to neutrino emission. This balance dictates the star's thermal evolution. The mode heating follows from

$$\dot E_{\rm sv} = \frac{\alpha^2 \tilde J \Omega^2 M R^2}{t_{\rm sv}}. \quad (12)$$

For an $n = 1$ polytrope, we have $\tilde J \approx 1.635 \times 10^{-2}$ (Owen et al. 1998), so

$$\dot E_{\rm sv} \approx 2.7 \times 10^{43}\alpha_s^2\left(\frac{M}{1.4 M_\odot}\right)^{9/4}\left(\frac{R}{10\ \mathrm{km}}\right)^{-15/4} \times\left(\frac{\nu}{100\ \mathrm{Hz}}\right)^{2}\left(\frac{T}{10^8\ \mathrm{K}}\right)^{-2}\ \mathrm{erg\ s^{-1}}. \quad (13)$$

This should be compared to the neutrino luminosity associated with the modified Urca reaction. Using the estimated neutrino luminosity from Shapiro & Teukolsky (1983), we have (for a

[9] Note that the upper limit on the star's radius accords with the constraint from the tidal deformability during the inspiral of the first-detected gravitational-wave signal from merging neutron stars (Abbott et al. 2017b).

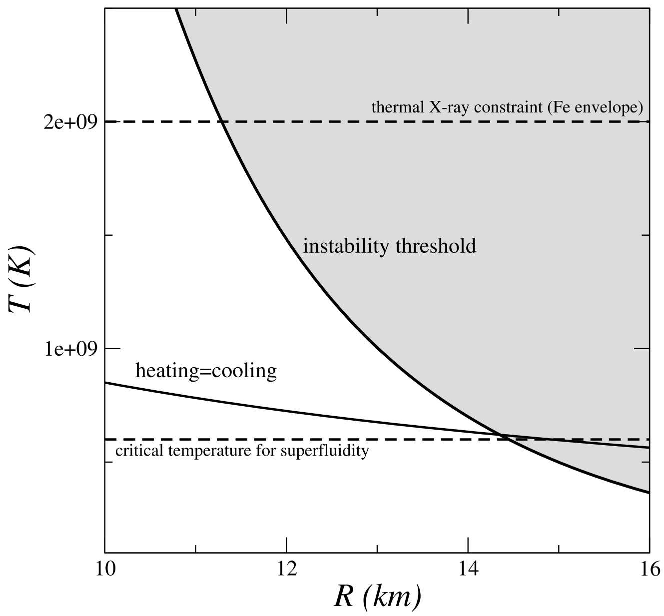

Figure 4. Comparing the two temperatures from Equations (11) and (15). The first indicates the threshold above which the r-mode would be unstable. The second provides the temperature at which r-mode heating balances modified Urca cooling. The results suggest that we need the neutron star radius to be greater than about 14.5 km in order for the system to have reached thermal balance inside the r-mode instability window. However, one would not have to change our estimated growth/damping timescales by much to bring the radius into the suggested range of 10–14 km. We also show the X-ray constraint on the star's core temperature (which sets an upper limit), assuming a heavy element (Fe) envelope, as well as an indicative level for the onset of core superfluidity (with many viable models entering at a lower level than this).

constant density star)

$$\dot{E}_{\rm mU} \approx 5.6 \times 10^{31} \left(\frac{M}{1.4 M_\odot}\right)^{2/3} \left(\frac{R}{10 \text{ km}}\right) \left(\frac{T}{10^8 \text{ K}}\right)^8 \text{ erg s}^{-1}. \tag{14}$$

By equating the two rates, for the observed spin rate, we see that thermal balance implies a core temperature:

$$T \approx 1.3 \times 10^9 \alpha_s^{1/5} \left(\frac{M}{1.4 M_\odot}\right)^{19/120} \left(\frac{R}{10 \text{ km}}\right)^{-19/40} \text{ K}. \tag{15}$$

Making use of the inferred saturation amplitude from the spin down, we see that (for a 10 km star), thermal balance implies that $T \approx 8.6 \times 10^8$ K. If we take the radius to be 14 km, the corresponding temperature is $T \approx 6.4 \times 10^8$ K. In order to see if the model is consistent, we compare the two temperatures from Equations (11) and (15) in Figure 4. The figure shows that we would need the radius of the star to be greater than about 14.5 km in order for the system to have reached thermal balance inside the r-mode instability window. Of course, one would not have to change the estimated growth/damping timescales by much to bring the radius into the suggested radius range (below 14 km). Basically, it may not be unreasonable to suggest that the instability could operate in this system.

The various estimates we used obviously come with a range of caveats, and one should perhaps not read too much into the conclusion that the r-mode scenario would appear to be consistent. We have based the argument on simple Newtonian estimates, which may be adequate for a first attempt, but cannot be used in combination with a more realistic equation of state. If we want to make the model more realistic, then we have to consider the r-mode problem in general relativity (Lockitch et al. 2001, 2003; Idrisy et al. 2015). As an alternative, one could make use of parameterized versions of the different timescales, as advocated by Alford & Schwenzer (2014).

Based on our estimates, the most important issue relates to the saturation amplitude, α_s, which is at least an order of magnitude larger than predicted. It is also legitimate to ask if it is appropriate to rely on the neutron shear viscosity. One important issue is in regards to the expected onset of neutron superfluidity, which would suppress neutron contribution to the shear viscosity. Below the superfluid transition temperature, the main shear viscosity is due to electron–electron scattering. If we had used the corresponding damping timescale (see Andersson & Kokkotas 2001) in our estimates, then the different temperatures would not be consistent. However, it could well be that the outer core of the star has yet to cool below the superfluid transition. From the sample of relevant pairing gaps considered in Ho et al. (2015b, see their Figure 10), we learn that only the models with the largest gaps have a critical temperature above 6×10^8 K. As our estimated temperatures are (just) above this, it does not seem unreasonable to assume that the neutrons are normal and hence that Equation (10) applies. The fact that the composition of the neutron star core is uncertain is also a concern, but the main r-mode damping is associated with the fluid motion at around 70–80% or so of the star's radius.[10] As more exotic phases (and states) of matter (like hyperons; Lindblom & Owen 2002) may not be present in the star's outer core, it seems entirely plausible that their presence (or absence) at higher densities would have little effect on the r-mode damping.

Before we proceed, let us make one further comment on the temperature. The core temperatures we require for the r-mode instability to be active are higher than one would expect if the star was simply cooling in isolation. Hence, it is worth considering whether one may be able to use X-ray observations to constrain the scenario. Taking the observed nonthermal X-ray luminosity of 6×10^{35} erg s^{-1} from Chen et al. (2006) as an upper limit on the surface emission, one would infer a surface temperature of about 5×10^6 K. This converts into a limit on the core temperature of $<2 \times 10^9$ K (Fe envelope) or $<6 \times 10^8$ K (for H). Basically, the core temperatures we inferred would be easily compatible with a heavy element envelope.

Finally, as an alternative, one may consider the possibility that the r-mode is stable (as one might have expected in the first place), but that it is excited by some impulsive mechanism. However, it is not straightforward to make a stable r-mode scenario consistent with the observations. In order to arrive at the suggested braking index of $n = 7$, we need the gravitational-wave emission to dominate the spin down, and the result requires a constant mode amplitude. While one can easily evolve the mode amplitude in the stable regime, the result tends to be very different from what we require. Moreover, the impulsive mode excitation is problematic. One would need to

[10] As long as we ignore the presence of possible viscous boundary layers, e.g., associated with the crust–core transition, which may have strongly localized dissipation (see for example Levin & Ushomirsky 2001).

pump a great deal of energy into the mode; much more than seems allowed by the energy budget associated with the glitches. The unstable r-mode scenario resolves these issues in a seemingly natural way.

While our arguments suggest that the unstable r-mode scenario may accommodate J0537–6910, it is important to keep in mind that this would involve a large instability region at the inferred (relatively high) temperatures. The physics may, in principle, allow for this, but there is obvious tension between this scenario and the many observed fast-spinning accreting neutron stars in low-mass X-ray binaries (Ho et al. 2011; Haskell et al. 2012; Mahmoodifar & Strohmayer 2017; Patruno et al. 2017). In order for those systems not to spin down due the r-mode instability, the threshold must be around 600 Hz at a core temperature a factor of a few lower than Equation (15). This would require the instability to have a sharp feature in a fairly narrow temperature range. This could be problematic, but one can think of scenarios that would predict this behavior. For example, the onset of core superfluidity, which brings vortex-mediated mutual friction into play, may have exactly this effect (see, for example, Figure 6 in Haskell et al. 2009). Other additional physics may also lead to the instability window having a more complex shape (see, for example, Levin & Ushomirsky 2001; Bondarescu et al. 2009; Kantor et al. 2016).

4. Observational Tests

The observational evidence and the theoretical estimates clearly do not settle the issue, but we cannot rule out the notion that the r-mode instability may impact on the spin evolution[11] of J0537–6910. We do not have to bend our understanding of the physics very much to make the observations fit the theory. Given this, let us consider the problem from an observational point of view. To be specific, can we use observations to constrain our ignorance about the theory?

There are two (obvious) ways to address this question. Additional X-ray timing of the pulsar may strengthen (or not) the argument in favor of a braking index close to $n = 7$. This would be further evidence in favor of the r-mode scenario, but it would still be circumstantial. Meanwhile, a dedicated gravitational-wave search may provide a limit on the allowed r-mode amplitude. Since we require the gravitational-wave emission to dictate the observed spin down, one might be able to set a strong enough constraint to rule this out. In reality, the two kinds of observations are linked. In order to achieve the best gravitational-wave sensitivity, one would need a reliable timing solution, e.g., one provided by NICER. In absence of this, one would have to fall back on a less optimal search strategy.

In order to set the stage for a more detailed discussion of the detection problem, we assess the detectability of the emerging gravitational waves in the standard way. First of all, we note that (ignoring relativistic corrections, see below) the frequency of the emerging gravitational waves is (for the main $l = m = 2$ r-mode)

$$f_{\rm gw} = \frac{4\nu}{3} \approx 83 \text{ Hz}. \tag{16}$$

[11] Note that a comprehensive model must also be able to accommodate the long-term behavior of the pulsar, which is governed by an effective negative braking index, which is possibly related to permanent $\dot{\nu}$ offsets associated with the glitches.

We combine this with the gravitational-wave flux to find the strain amplitude for the optimal source and detector orientations (Jaranowski et al. 1998). That is, we use

$$h_0^2 = \frac{10G}{c^3}\left(\frac{1}{2\pi f_{\rm gw} d}\right)^2 \dot{E}, \tag{17}$$

where d is the distance to the source. Combining this with the gravitational-wave luminosity for the r-modes, we arrive at (Owen 2010)

$$h_0 \approx \frac{3\alpha_s}{4d}\left(\frac{10GMR^2\tilde{J}}{c^3 t_{\rm gw}}\right)^{1/2}. \tag{18}$$

Scaling to suitable parameter values, we have

$$h_0 \approx 7.5 \times 10^{-25}\alpha_s\left(\frac{M}{1.4M_\odot}\right)\left(\frac{R}{10 \text{ km}}\right)^3 \times \left(\frac{\nu}{100 \text{ Hz}}\right)^3\left(\frac{50 \text{ kpc}}{d}\right). \tag{19}$$

Assuming that the r-mode amplitude is, indeed, the α_s inferred from the spin down and that the distance to the pulsar is 50 kpc, we have $h_0 \approx (2\text{–}3) \times 10^{-26}$ for a neutron star radius in the range of 10–14 km.

As a rough idea of the detectability of this signal, let us assume that the system does not evolve much during the observation period (the frequent glitches may be a problem). Then, the effective amplitude increases as the square root of the observing time, $t_{\rm obs}$, and we simply assess the detectability by comparing $\sqrt{t_{\rm obs}}h_0(t)$ to $11.4\sqrt{S_n}$, where S_n is the power spectrum of the detector noise.

For a targeted search, the comparison we need is, in fact, straightforward. We can use the most recent targeted search for continuous gravitational waves, based on about 70 days of LIGO data from the first observing run (O1), from Abbott et al. (2017a). Estimating the sensitivity at $f_{\rm gw} \approx 80$ Hz from Figure 1 in that paper, we see that $h_0 \approx 3 \times 10^{-26}$, which is very close to our estimated strain. This is obviously interesting. Of course, we need to do a little bit better to rule out (or in!) the scenario. If we instead consider advanced LIGO operating at design sensitivity, then we would have $\sqrt{S_n} \approx 4 \times 10^{-24}$ at the frequency we are interested in (see figures in Abbott et al. 2016b). From this, we see that the predicted level of signal would be detectable with less than two months worth of data. With a longer integration time, one might be able to put interesting constraints on the model.

Of course, these estimates assume a targeted search, which in turn requires a reliable timing solution. Given that it is not clear that this information will be available, it is worth considering how well one may be able to do with a less optimal strategy. Since we know the location of the source, the natural strategy would be a directed search. In this case, the attainable sensitivity is not at the level we have assumed. Based on the analysis of Wette et al. (2008), demonstrated in a LIGO search for signals from the neutron star in the Cassiopea A supernova remnant (Abadie et al. 2010) (and slightly more sensitive than the semi-coherent analysis used by, for example, Zhu et al. 2016 and Abbott et al. 2016a), one would expect a directed search to be a factor of 3 or so less sensitive than a targeted search. Assuming that one would lose this factor of 3 in sensitivity, one would have to compensate by increasing the effective integration time by a factor of 9–10, leading to a required observation time of 18–20 months for advanced detectors at the design sensitivity. Such integration times would be challenging, but perhaps not prohibitively so.

These, admittedly rough, estimates suggest that we should (eventually) be able to use observations to either confirm or rule out (which may be more likely) the notion that an unstable r-mode is present in J0537–6910. Of course, in order to carry out the suggested gravitational-wave search, one would need to go a couple of steps beyond our simple estimates. Perhaps most importantly, one has to consider the fact that the true r-mode frequency is not going to be $4\nu/3$. In a realistic neutron star model, the r-mode frequency is shifted by a range of effects. For a relatively slowly spinning star (such that one can ignore rotational shape corrections), the largest correction is likely due to relativity (the gravitational redshift and the rotational frame dragging (Lockitch et al. 2001, 2003), see also Andersson et al. 2014). The most detailed investigation of the problem was presented in Idrisy et al. (2015). The results suggest that we should consider the realistic r-mode frequency to lie in the range $1.39\nu < f_{\rm gw} < 1.57\nu$. That is, for J0537–6910, one should search for a signal in the range $f_{\rm gw} \approx 86$–98 Hz (note that the Newtonian result, $f_{\rm gw} \approx 83$ Hz, is not inside this interval).

It is also worth commenting on the challenge of carrying out a search for gravitational waves from J0537–6910 without timing data. Any such effort could be seriously affected by the frequent glitches (Ashton et al. 2017). The glitches would also impact on attempts to stack shorter segments of data to increase the sensitivity. While one may, in principle, be able to stack interglitch data (shorter than the three month or so interval between glitches), this may be practically difficult without an identification of the glitches in the first place. In reality, it may be tricky to carry out the required search in archival O1-O2 LIGO data.

5. Concluding Remarks

The 62 Hz X-ray pulsar PSR J0537–6910 is undoubtedly an intriguing object. As the fastest known young pulsar, it has a complex spin-history that is frequently interrupted by glitches. This makes matching timing observations to theory a challenge. At the same time, one may hope that the enigmatic behavior may shed light on the involved physics, like the superfluid vortex dynamics thought to dictate the relaxation after each glitch event.

We argued that J0537–6910 should be a prime target for a joint observing campaign between NICER and the LIGO-Virgo network. The argument draws on an analysis of the complete timing data from *RXTE* and an observation that a trend in the interglitch behavior of the pulsar may indicate an effective braking index close to $n = 7$. This value would accord with a neutron star spinning down due to gravitational waves from an unstable r-mode. We discussed to what extent this scenario may be consistent and whether the associated gravitational-wave signal would be within reach of ground-based detectors. In essence, our estimates suggest that one may be able to use observations to constrain (or even rule out) the idea. This is an interesting prospect for the future.

We would like to thank Lucien Kuiper for deriving and providing the ToAs used in this study. We are also grateful to Maria Alessandra Papa for helpful discussions. N.A. and W.C.G.H. acknowledge funding from STFC in the UK through grant No. ST/M000931/1. D.A. acknowledges support from the Polish National Science Centre (SONATA BIS 2015/18/E/ST9/00577, P.I.: B. Haskell). C.M.E. acknowledges funding from CONICYT FONDECYT/Regular 1171421 in Chile. B.H. received funding under the European Union's Horizon 2020 research and innovation programme under grant agreement No. 702713. Partial support comes from NewCompStar, COST Action MP1304.

References

Aasi, J., Abadie, J., Abbott, B. P., et al. 2014, ApJ, 785, 119
Abadie, J., Abbott, B. P., Abbott, R., et al. 2010, ApJ, 722, 1504
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016a, PhRvD, 94, 102002
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2016b, LRR, 19, 1
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017a, ApJ, 839, 12
Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017b, PhRvL, 119, 161101
Alford, M. G., & Schwenzer, K. 2014, ApJ, 781, 26
Andersson, N. 1998, ApJ, 502, 708
Andersson, N., Jones, D. I., & Ho, W. C. G. 2014, MNRAS, 442, 1786
Andersson, N., & Kokkotas, K. D. 2001, IJMPD, 10, 381
Antonopoulou, D., Espinoza, C. M., Kuiper, L., & Andersson, N. 2018, MNRAS, 473, 1644
Arras, P., Flanagan, E. E., Morsink, S. M., et al. 2003, ApJ, 591, 1129
Ashton, G., Prix, R., & Jones, D. I. 2017, PhRvD, 96, 063004
Bondarescu, R., Teukolsky, S. A., & Wasserman, I. 2007, PhRvD, 76, 064019
Bondarescu, R., Teukolsky, S. A., & Wasserman, I. 2009, PhRvD, 79, 104003
Bondarescu, R., & Wasserman, I. 2013, ApJ, 778, 9
Brink, J., Teukolsky, S. A., & Wasserman, I. 2004a, PhRvD, 70, 121501
Brink, J., Teukolsky, S. A., & Wasserman, I. 2004b, PhRvD, 70, 124017
Brink, J., Teukolsky, S. A., & Wasserman, I. 2005, PhRvD, 71, 064029
Chen, Y., Wang, Q. D., Gotthelf, E. V., et al. 2006, ApJ, 651, 237
Ferdman, R. D., Archibald, R. F., Gourgouliatos, K. N., & Kaspi, V. M. 2018, ApJ, 852, 123
Fuentes, J. R., Espinoza, C. M., Reisenegger, A., et al. 2017, A&A, 608, A131
Gendreau, K. C., Arzoumanian, Z., Adkins, P. W., et al. 2016, Proc. SPIE, 9905, 99051H
Haskell, B., Andersson, N., & Passamonti, A. 2009, MNRAS, 397, 1464
Haskell, B., & Antonopoulou, D. 2014, MNRAS, 438, L16
Haskell, B., Degenaar, N., & Ho, W. C. G. 2012, MNRAS, 424, 93
Ho, W. C. G., Andersson, N., & Haskell, B. 2011, PhRvL, 107, 101101
Ho, W. C. G., Elshamouty, K. G., Heinke, C. O., & Potekhin, A. Y. 2015b, PhRvC, 91, 015806
Ho, W. C. G., Espinoza, C. M., Antonopoulou, D., & Andersson, N. 2015a, SciA, 1, e1500578
Idrisy, A., Owen, B. J., & Jones, D. I. 2015, PhRvD, 91, 024001
Jaranowski, P., Krolak, A., & Schutz, B. F. 1998, PhRvD, 58, 063001
Kantor, E. M., Gusakov, M. E., & Chugunov, A. I. 2016, MNRAS, 455, 739
Kuiper, L., & Hermsen, W. 2015, MNRAS, 449, 3827
Levin, Y., & Ushomirsky, G. 2001, MNRAS, 324, 917
Lindblom, L., & Owen, B. J. 2002, PhRvD, 65, 063006
Lindblom, L., Owen, B. J., & Morsink, S. M. 1998, PhRvL, 80, 4843
Lockitch, K. H., Andersson, N., & Friedman, J. L. 2001, PhRvD, 63, 024019
Lockitch, K. H., Friedman, J. L., & Andersson, N. 2003, PhRvD, 68, 124010
Mahmoodifar, S., & Strohmayer, T. 2017, ApJ, 840, 94
Marshall, F. E., Gotthelf, E. V., Middleditch, J., Wang, Q. D., & Zhang, W. 2004, ApJ, 603, 682
Marshall, F. E., Gotthelf, E. V., Zhang, W., Middleditch, J., & Wang, Q. D. 1998, ApJL, 499, L179
Middleditch, J., Marshall, F. E., Wang, Q. D., Gotthelf, E. V., & Zhang, W. 2006, ApJ, 652, 1531
Owen, B. J. 2010, PhRvD, 82, 104002
Owen, B. J., Lindblom, L., Cutler, C., et al. 1998, PhRvD, 58, 084020
Patruno, A., Haskell, B., & Andersson, N. 2017, ApJ, 850, 106
Shapiro, S. L., & Teukolsky, S. A. 1983, Black Holes, White Dwarfs and Neutron Stars: The Physics of Compact Objects (New York: Wiley)
Steiner, A. W., Heinke, C. O., Bogdanov, S., et al. 2018, MNRAS, 476, 421
Wang, Q. D., & Gotthelf, E. V. 1998, ApJ, 494, 623
Watts, A. L., Krishnan, B., Bildsten, L., & Schutz, B. F. 2008, MNRAS, 389, 839
Wette, K., Owen, B. J., Allen, B., et al. 2008, CQGra, 25, 235011
Zhu, S. J., Papa, M. A., Eggenstein, H.-B., et al. 2016, PhRvD, 94, 082008