The Astrophysical Journal, 850:106 (12pp), 2017 November 20 https://doi.org/10.3847/1538-4357/aa927a

The Spin Distribution of Fast-spinning Neutron Stars in Low-mass X-Ray Binaries: Evidence for Two Subpopulations

A. Patruno[1,2], B. Haskell[3], and N. Andersson[4]
[1] Leiden Observatory, Leiden University, P.O. Box 9513, NL-2300 RA Leiden, The Netherlands
[2] ASTRON, Netherlands Institute for Radio Astronomy, Postbus 2, NL-7990 AA Dwingeloo, The Netherlands
[3] Nicolaus Copernicus Astronomical Center, Polish Academy of Sciences, ul. Bartycka 18, 00-716 Warsaw, Poland
[4] Mathematical Sciences and STAG Research Centre, University of Southampton, Southampton SO17 1BJ, UK
Received 2017 May 22; revised 2017 October 2; accepted 2017 October 7; published 2017 November 21

Abstract

We study the current sample of rapidly rotating neutron stars in both accreting and non-accreting binaries in order to determine whether the spin distribution of accreting neutron stars in low-mass X-ray binaries (LMXBs) can be reconciled with current accretion torque models. We perform a statistical analysis of the spin distributions and show that there is evidence for two subpopulations among LMXBs, one at a relatively low spin frequency, with an average of ≈300 Hz and a broad spread, and a peaked population at higher frequency with an average spin frequency of ≈575 Hz. We show that the two subpopulations are separated by a cut-point at a frequency of ≈540 Hz. We also show that the spin frequency of radio millisecond pulsars (RMSPs) does not follow a log-normal distribution and shows no evidence for the existence of distinct subpopulations. We discuss the uncertainties of different accretion models and speculate that either the accreting neutron star cut-point marks the onset of gravitational waves as an efficient mechanism to remove angular momentum or some of the neutron stars in the fast subpopulation do not evolve into RMSPs.

Key words: binaries: general – gravitational waves – stars: neutron – (stars:) pulsars: general – X-rays: binaries – X-rays: stars

1. Introduction

The fastest-spinning neutron stars are found either as old radio millisecond pulsars (RMSPs) or in their progenitor systems, the accreting neutron stars in low-mass X-ray binaries (LMXBs). Accretion theory predicts that the majority of neutron stars in LMXBs should pulsate and therefore be accreting X-ray pulsars. However, the number of AMXPs is only 19 among ≈200 non-pulsating LMXBs (see Patruno & Watts 2012 for a review). Ten other LMXBs are nuclear-powered pulsars (NXPs), showing short-lived "burst oscillations" during runaway thermonuclear explosions on the neutron star surface (Watts 2012). Tauris (2012) proposed that when an AMXP finally turns into an RMSP, the neutron star strongly decelerates and loses about half of its rotational energy during the detachment process of the companion from its Roche lobe. According to this scenario, radio millisecond (ms) pulsars should therefore be relatively slower than accreting ms pulsars.

However, this does not explain why accreting submillisecond pulsars have not been found (so far). Indeed, a consequence of the scenario proposed by Tauris (2012) is that the fastest RMSPs known today (like the 716 Hz PSR J1748-2446ad, Hessels et al. 2006) should have been spinning at more than 1000 Hz during their accreting phase. This is not observed in the population of accreting neutron stars, where no object is known to spin faster than about 619 Hz (Hartman et al. 2003).

An interesting scenario predicts that gravitational waves act as a "brake" on the neutron star spin and balance the accretion torques, although it is still unclear whether this really is the case (Haskell & Patruno 2011). Bildsten (1998) and Andersson et al. (1999) had anticipated that a speed limit to spinning neutron stars should be expected if gravitational waves are emitted by these systems. As noted by Chakrabarty et al. (2003) and Chakrabarty (2008), at frequencies on the order of 700 Hz or more, the braking torque applied by gravitational waves (the strength of which scales as the fifth power of the spin frequency for deformed rotating stars) might be sufficiently strong to balance accretion torques and prevent a further spin up. However, it is by no means clear that we need gravitational waves to explain these systems. Patruno et al. (2012b) first showed that the presence of a magnetosphere with a minimum strength of $\sim 10^8$ G is in principle sufficient to explain the lack of neutron stars spinning faster than ∼700 Hz. In that work, the condition for the spin equilibrium set by the disk/magnetosphere interaction was calculated for an average mass-accretion rate during an outburst. However, Bhattacharyya & Chakrabarty (2017) and D'Angelo (2016) have recently shown that transient accretion with a varying mass-accretion rate can strongly affect the spin evolution of these systems. In particular, Bhattacharyya & Chakrabarty (2017) have shown that when this effect is taken into account, the existence of a magnetosphere with a strength of 10^8 G is no longer sufficient to explain the observed spin frequency limit. This suggestion is indicative, but several open issues remain, the most important ones being which precise mechanism generates gravitational waves, whether this mechanism is sufficient to balance accretion torques on timescales of hundred millions of years, and to what extent it can reproduce the observed spin evolution of accreting pulsars. Moreover, alternative hypotheses like the possible presence of a trapped disk (which substantially modifies the long-term accretion torque) need to be considered seriously (D'Angelo 2016).

Considering the overall population, Papitto et al. (2014) performed a statistical analysis of different subsets of neutron stars (see also Hessels 2008 for seminal work on the same topic). They found that RMSPs are on average slower than NXPs, a fact that could in principle be ascribed to the loss of

Table 1
Radio, Accreting, and Nuclear-powered Millisecond Pulsar Groups

Group	Nr. Objects	$\langle\nu\rangle$ (Hz)	Std. Dev. (Hz)	Shapiro-Wilk (p-value)	Reject H_0
RMSPs	337	253.6	131.6	1.2e-5	yes
AMXPs	19	367.8	153.0	0.086	no
Spiders	61	366.5	155.2	0.70	no
NXPs	10	502	123.5	0.017	yes
LMXBs	29	414.1	155.5	0.012	yes

Note. H_0 is the null hypothesis that the data follow a normal distribution.

angular momentum during the Roche-lobe decoupling phase (although see D'Angelo 2016 for a critique of this scenario).

In this paper we perform an extended statistical analysis on an updated sample of millisecond pulsars—including radio, accretion and nuclear-powered systems. We find that even though most of the conclusions of Papitto et al. (2014) still hold, there is a significant difference in the behavior of RMSPs and accreting neutron stars, even when we consider the small sample of AMXPs alone (Section 2). We discuss the spin distribution of fast-spinning accreting neutron stars by dividing the observed systems into subpopulations and performing a statistical analysis to search for evidence of subpopulations by statistical inference. We justify this approach in Section 3 and show the existence of a strong peak in the distribution and a separation cut-point above which a significant subgroup of objects seem to cluster. In Section 4 we verify whether there is sufficient observational evidence for the presence of a common underlying population for all different subgroups of neutron stars (accreting and non-accreting). In Section 5 we discuss open problems with the current torque theories, whether they are sufficient to explain the observed distribution, and what role may be played by gravitational waves. We then discuss the properties of all accreting neutron stars that are close to the observed spin frequency limit (Section 6). Finally, we consider recent evidence pointing to the absence of a magnetosphere in some of these systems, and the implications of this for the spin evolution (Section 7).

2. Pulsar Spin Distributions

Papitto et al. (2014) considered several different neutron star groups, specifically the binary RMSPs, the eclipsing radio ms pulsars (which we refer to as "spiders" from now on), accreting and nuclear-powered ms pulsars, and a subgroup of objects that comprised all accreting ms pulsars plus all eclipsing radio pulsars. The latter group was called "transitional millisecond pulsars" under the assumption that all AMXPs and eclipsing radio pulsar binaries show transitions between a radio ms pulsar state and an LMXB state (as seen in three such systems, see Archibald et al. 2009; Papitto et al. 2013; Bassa et al. 2014). In this work we do not consider this last group. Instead, when we refer to transitional ms pulsars (tMSPs), we consider only the three systems for which there is actual evidence of a transition. We also consider both isolated ms pulsars and binary ms pulsars (which form the RMSPs sample after excluding the spiders), since the progenitor of isolated ms pulsars must have gone through episodes of accretion (recycling) in their past history. We do this because we wish to introduce the least amount of selection bias in our samples (see also Section 2.2). Our sample of neutron stars is summarized in Table 1 (see also the catalog https://apatruno.wordpress.com/about/millisecond-pulsar-catalogue/ for a complete list of objects). Another difference in this work is that in contrast to Papitto et al. (2014), we do not interpret our p-values as the probability for the null hypothesis since it is not possible to associate any probability for the null in hypothesis testing. Finally, we note that some of the differences in the results might be ascribed to the increased sample sizes.

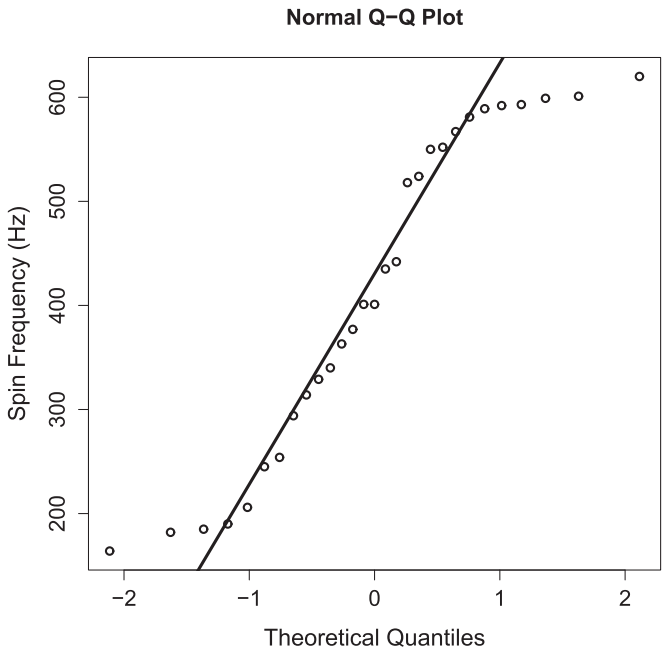

Figure 1. Quantile-quantile plot for the LMXB (AMXPs+NXPs) distribution. The 45° oblique line is the expected theoretical trend for a normal distribution. The deviations are strongest toward the tail of the distribution.

2.1. Do Pulsar Spins Follow a Normal Distribution?

Throughout this work, we use the statistical package R (v 3.3.3) to perform our analysis. The first test we apply is a Shapiro–Wilk (SW) normality test (package *stats*) to check whether the spin frequency of each sample is compatible with a normal distribution (the null hypothesis H_0 being that the data follow a normal distribution). The choice of the Shapiro–Wilk test is motivated by its higher power for a given significance when compared to other tests. We first test the AMXPs+NXPs sample (henceforth referred to as LMXBs) for normality, under the hypothesis that both the AMXPs and NXPs belong to the same underlying population. We use a significance level of $\alpha = 0.05$ throughout this work, which means that we have a 5% chance of false positives for each test performed. The SW test gives a p-value of $p < 0.02$, and thus we reject our null hypothesis at the 5% significance level. A quantile-quantile plot of the sample shows how the data deviate from the expected 45° line of a normal distribution (see Figure 1).

The deviation from normality is evident also from a simple by-eye inspection of the histogram of the spin distribution shown in Figure 2, which displays a prominent peak around the 500–600 Hz bin (see Table 2 for the full list of pulsating LMXBs used in this work).

We then perform the SW test for all other pulsar samples, with the results summarized in Table 1. We cannot reject the null hypothesis only for the spiders and the AMXPs samples.

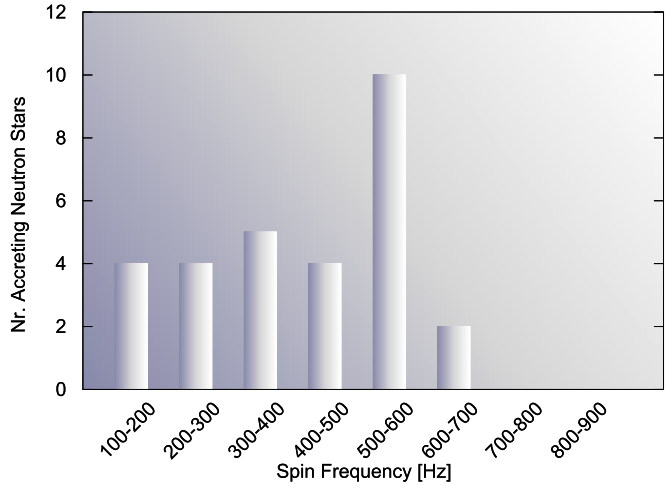

Figure 2. Histogram of accreting neutron stars comprising all known AMXPs and NXPs (as reported in Table 2). No source is known above ≈619 Hz.

Table 2
Accretion and Nuclear-powered Millisecond Pulsars

Source Name	Spin Frequency (Hz)	Type	Orb. Period (hr)
4U 1728–34	363	NXP	4.2
KS 1731–260	524	NXP	N/A
IGR J17191–2821	294	NXP	N/A
4U 1702–429	329	NXP	N/A
SAX J1750.8–2900	601	NXP	N/A
GRS 1741.9–2853	589	NXP	N/A
EXO 0748–676	552	NXP	3.8
4U 1608–52	619	NXP	12.9
4U 1636–536	581	NXP	3.8
MXB 1659–298	567	NXP	7.1
Aql X–1	550	AMXP	18.9
IGR J00291+5934	599	AMXP	2.5
PSR J1023+0038	592	AMXP	4.8
XSS J12270–4859	593	AMXP	6.9
SAX J1808.4–3658	401	AMXP	2.0
XTE J1751–305	435	AMXP	0.7
XTE J0929–314	185	AMXP	0.7
XTE J807–294	190	AMXP	0.7
XTE J1814–338	314	AMXP	4.3
HETE J1900.1–2455	377	AMXP	1.4
Swift J1756.9–258	182	AMXP	0.9
SAX J1748.9–2021	442	AMXP	8.8
NGC6440 X-2	206	AMXP	0.95
IGR J17511-3057	245	AMXP	3.5
Swift J1749.4-2807	518	AMXP	8.8
IGR J17498-2921	401	AMXP	3.8
IGR J18245-245	254	AMXP	11.0
MAXI J0911–655	340	AMXP	0.7
IGR J17602–6143	164	AMXP	N/A

Note. The sources highlighted in bold are discussed in Section 6.

2.2. What is the Spin Distribution of RMSPs?

Lorimer et al. (2015) showed in an analysis of a sample of 56 RMSPs that their distribution was consistent with being log-normal. Tauris (2012) and Papitto et al. (2014) also showed that the distribution of RMSPs is not consistent with a normal

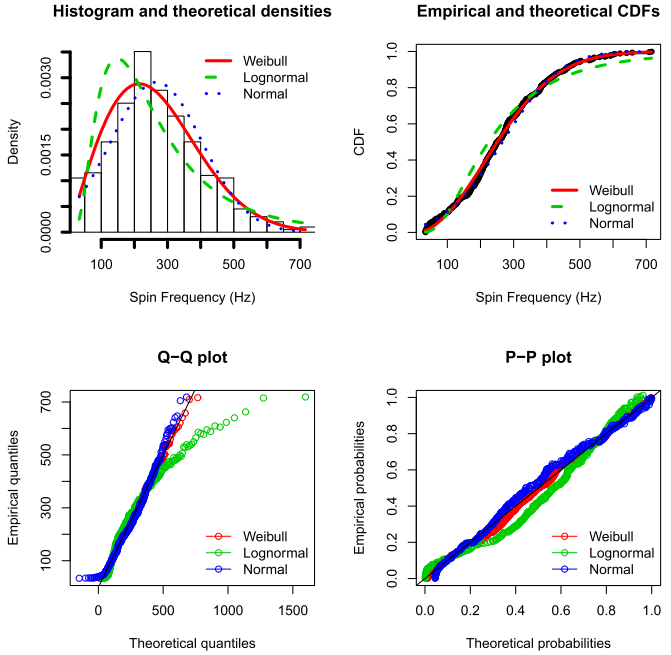

Figure 3. Probability density of RMSPs (top left panel) with three theoretical distributions fitted (Weibull, log-normal, and normal). Top right panel: empirical and theoretical cumulative density functions for RMSPs. Bottom left panel: empirical quantile-quantile plot for RMSPs. Note that the data show the strongest deviations from the theoretical log-normal distribution (oblique solid black line). Bottom right: empirical probability plot of RMSPs.

distribution, in line with our previous findings. Since in this work we use both isolated and binary RMSPs, we now justify this by means of statistical analysis. First we perform an SW test for normality for isolated and for binary millisecond pulsars. The results show that neither is consistent with a normal distribution (p-values of 0.02 and 0.0001 for isolated and binary millisecond pulsars, respectively). A k-sample AD test shows that they they are compatible with being drawn from the same parent population (p-value of 0.14). Furthermore, the mean spin frequency of isolated and binary RMSPs is nearly identical (245 versus 248 Hz, with similar standard deviations). Therefore our decision to include both systems in the sample of RMSPs is justified. We now try to characterize the spin frequency distribution of all RMSPs by considering three theoretical distributions: log-normal, normal, and a Weibull distribution. The latter has the form

$$f(x,\,k,\,\gamma) = \frac{k}{x}\left(\frac{x}{\gamma}\right)^{k} e^{-(x/\gamma)^{k}}, \qquad (1)$$

where k and γ (both assumed positive) are fitting parameters known as shape and scale parameters, respectively. We fit the three distributions and then compare the agreement between the theoretical expectations and the actual data in four plots in Figure 3: a histogram of empirical and theoretical densities, a cumulative distribution function plot, a quantile-quantile plot, and a P–P plot (which compares the empirical cumulative distribution function of our data with the three theoretical cumulative distribution functions). The strongest deviations are seen for the log-normal distribution, followed by the normal distribution. The Weibull distribution with shape

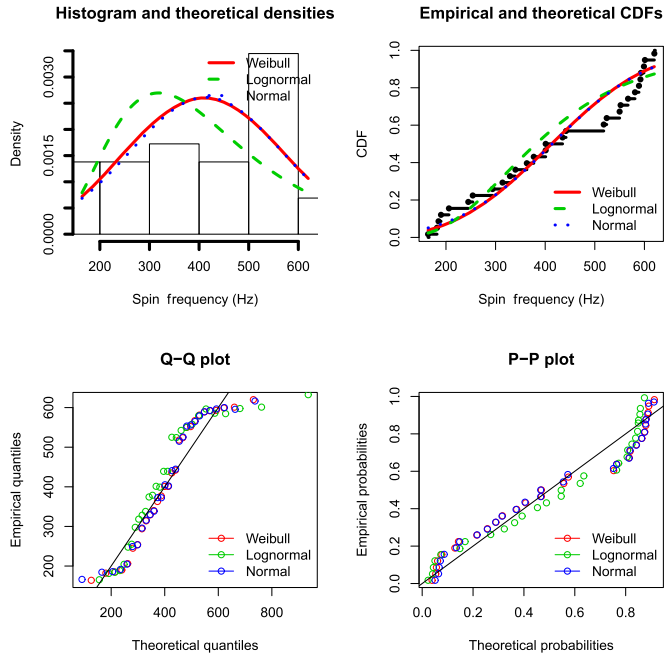

Figure 4. Same as for Figure 3, but for the LMXB sample.

$k = 2.10 \pm 0.08$ and spread $\gamma = 292 \pm 8$ Hz gives the best match with the data.[5]

To determine whether the best match of the Weibull distribution really is superior to the normal and log-normal distributions, we look at the goodness-of-fit statistics (package *fitdistrplus*, with parameter estimation done by maximizing the likelihood function), in particular, we use the Bayesian information criterion (BIC; Schwarz 1978). The BIC is an optimal criterion to select the best model among many since it introduces a penalty for each additional free parameter used in the fit. In this way, the BIC can prevent us from erroneously concluding that a model with many free parameters is necessarily better than a model with fewer free parameters. This is not possible with a simple χ^2 goodness-of-fit test.

The difference between the minimum BIC number (5036.4 for the Weibull distribution) and the other two BIC numbers of the normal and log-normal models is above 30, indicating a very strong evidence against the higher BIC models (Kass & Raftery 1995). A one-sided Kolmogorov–Smirnov test confirms our results, giving a p-value of 0.38 for the Weibull distribution and values much lower than 0.01 for the normal and log-normal distributions. Again, we caution that this does not mean that the Weibull is necessarily the true underlying distribution of RMSPs (we have not considered, for example, the β distribution and other viable options) but that of the models we considered, we reject the normal and the log-normal distribution and suggest a Weibull as having higher likelihood of being the true underlying distribution.

A similar analysis applied to the LMXBs shows deviations from all three theoretical distributions (see Figure 4). Spiders show a good match both for a normal distribution and for the Weibull, with relatively strong deviations seen in the log-normal model (and a BIC number higher than 7 and 9 from the Weibull and normal BIC numbers).

[5] We note that a Weibull distribution with $k = 2$ is equivalent to the Rayleigh distribution, i.e., a distribution that often emerges in physical processes that involve scattering or that involve the sum of two normally distributed and independent vectorial components (Papoulis & Pillai 2002).

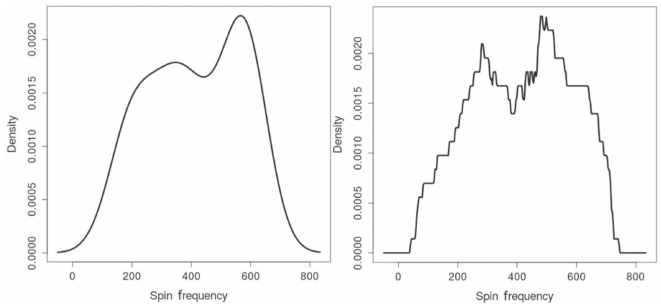

Figure 5. Kernel density estimation for the spin distribution of accreting neutron stars. A Gaussian kernel has been used in the left panel, whereas a rectangular (and thus discrete) kernel has been used in the right panel. The bandwidth of the kernels is 71 Hz in both cases. The KDE is strongly suggestive of a bimodality in the spin distribution.

3. The Spin Distribution of Accreting Neutron Stars

The separation of fast-rotating accreting neutron stars into AMXPs and NXPs might not be the best way to split the population of LMXBs (see Table 2) because the presence of accretion-powered pulses, which distinguishes AMXPs from NXPs, may not be related to their spin frequency.

We thus consider the population of AMXPs and NXPs together and then use a kernel density estimator (KDE) to try to extract features from the spin distribution. Indeed, although histograms are the simplest non-parametric density estimators, their properties heavily depend on the choice of the bin width (in addition to being discrete estimators by definition). We use a Gaussian kernel, which is continuous, and a rectangular discrete kernel. The KDE is performed with the package *density* and the result is shown in Figure 5. The distribution seems to be bimodal, with a clear spike around ∼550 Hz and a broader peak around ∼300 Hz.

In summary, the evidence suggests that

1. our LMXB sample is inconsistent with being drawn from a population following a normal distribution;
2. a by-eye inspection of the spin distribution shows a prominent peak at 500–600 Hz;
3. a KDE of the spin distribution appears bimodal.

Given these observations, we are motivated to model the data to extract information under the hypothesis that we are not dealing with a single population. In particular, we wish to test whether there is evidence for two subpopulations (not necessarily being AMXPs and NXPs) within a parent population without specifying which of the 29 data points of our sample belongs to which subpopulation. To do this, we use a fixed mixture model (e.g., Connolly et al. 2000), which uses an expectation-maximization algorithm (Guoshen et al. 2012) for fitting mixture models. A mixture model is a statistical test designed exactly for this purpose and has the advantage that it can estimate a "cut-point" between the two subpopulations. We use a mixture model with Gaussians, although other choices are possible. The mixture models operate with two parameters: a location parameter μ, where the data are concentrated and that corresponds to the mean value of the distribution (for a Gaussian), and a scale parameter σ, which gives the spread of the data and corresponds to the standard deviation for a

Table 3
Mixture Model Parameter Estimation

	Estimate	Confidence Interval (95%)
μ_1 (Hz)	302	255–348
σ_1 (Hz)	92	68–135
μ_2 (Hz)	574	555–593
σ_2 (Hz)	30	21–48
λ	0.6	0.4–0.8
Cut-point (Hz)	538	526–548

Gaussian. The mixture model then is defined as

$$f(\nu|\lambda, \mu_1, \sigma_1, \mu_2, \sigma_2) = \lambda \cdot \mathcal{D}_1(\nu|\mu_1, \sigma_1) + (1 - \lambda) \cdot \mathcal{D}_2(\nu|\mu_2, \sigma_2), \quad (2)$$

where ν is the spin frequency, λ is the mixture parameter, and $\mathcal{D}_1$ and $\mathcal{D}_2$ are the distributions used (in this case, Gaussians). The mixture parameter λ acts as a weight for the two distributions $\mathcal{D}_1$ and $\mathcal{D}_2$. The function $f(\nu|\lambda, \mu_1, \sigma_1, \mu_2, \sigma_2)$ is then evaluated with the expectation-maximization (EM; package *cutoff*) algorithm that is optimal for parameter estimation in probabilistic models with incomplete data sets. In short, the EM algorithm works in two steps, the first is the "expectation," where each member of the sample is assigned a guessed probability to belong to one of the two subpopulations. In the second step (the "maximization"), the maximum likelihood estimation is calculated. The two steps are then repeated until convergence. The probability that an accreting neutron star belongs to either of the distributions $\mathcal{D}_1$ and $\mathcal{D}_2$ is calculated as

$$p_1 = \frac{\lambda \cdot \mathcal{D}_1(\nu|\mu_1, \sigma_1)}{\lambda \cdot \mathcal{D}_1(\nu|\mu_1, \sigma_1) + (1 - \lambda) \cdot \mathcal{D}_2(\nu|\mu_2, \sigma_2)} \quad (3)$$

$$p_2 = \frac{(1 - \lambda) \cdot \mathcal{D}_2(\nu|\mu_2, \sigma_2)}{\lambda \cdot \mathcal{D}_1(\nu|\mu_1, \sigma_1) + (1 - \lambda) \cdot \mathcal{D}_2(\nu|\mu_2, \sigma_2)}. \quad (4)$$

We then use our significance level of $\alpha = 0.05$ (i.e., the false-alarm probability), along with 20,000 Monte Carlo simulations to calculate the location of the cut-point that separates the two distributions, and we find a value of about 540 Hz. The cut-point is defined as the spin frequency where the probability p_1 is equal to α (and $1 - \alpha$ for p_2). Therefore we expect less than one false positive (i.e., less than one accreting pulsar wrongly assigned in either of the two subpopulations). The results of our test for the parameters μ_1, μ_2, σ_1 and σ_2, λ and the cut-point are reported in Table 3 and Figure 6.

We then calculate the BIC numbers to estimate the significance for the presence of two subpopulations rather than a single one (with the package *mclust*). The BIC number differences are larger than about 7, which implies a strong evidence for the presence of the two subpopulations.[6] Finally, we use a parametric bootstrap method (package *boot*) to test the null hypothesis of a one-population spin distribution versus the two-component one. We produced 10^5 bootstrap realizations of the likelihood ratio statistics and checked the final p-values of the two models. A one-population model gives a p-value lower than 1%, whereas the two-subpopulation model gives $p \gg 5\%$, which confirms our previous findings.

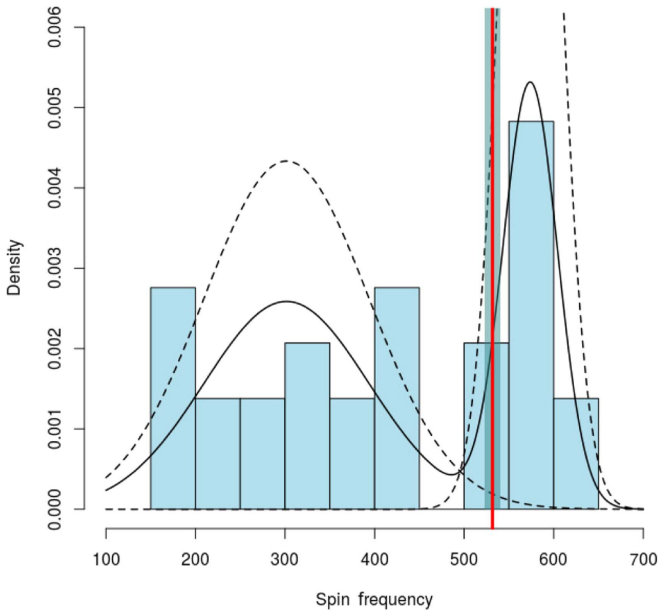

Figure 6. Mixture model for the spin distribution sample in Table 2. The dashed lines represent the initial distributions $\mathcal{D}_1$ and $\mathcal{D}_2$. The solid black line is the final mixture model after convergence. The vertical red line and the shaded gray area around it is the location of the cut-point that divides the sample into two subpopulations. The histogram of the spin distribution is plotted in 10 bins to display the cut-point at 540 Hz more clearly and has no influence on the final results.

Finally, we verified a posteriori our assumption that the two subpopulations both follow a Gaussian distribution. An SW test for normality of the slow and fast subpopulation gives a p-value of 0.36 and 0.56, respectively, thus justifiying our assumption.

4. Is there a Common Underlying Pulsar Population?

If we consider our two subpopulations, then the slower systems have an average spin frequency of about $\langle\nu\rangle \approx 302 \pm 92$ Hz (see Table 3). This is within one standard deviation from the average RMSP spin frequency, which is ≈250 Hz (see also Hessels 2008 and Papitto et al. 2014). Therefore, the suggestion that the slow subpopulation is indeed the progenitor population of RMSPs is consistent with observations.[7]

This leads us to consider whether the samples of AMXPs, NXPs, RMSPs, and spiders are compatible with being drawn from the same underlying population. To check this, we use a k-samples Anderson–Darling test (package *kSamples*), which does not require us to specify the distribution function of the population (which, in our case, is indeed unknown). We first compare LMXBs and RMSPs: the k-samples AD test gives a p-value of about $\approx 10^{-7}$ (H_0: both samples are drawn from the same underlying distribution), and thus we reject the null hypothesis that both samples come from a common population. Similar results are obtained if we compare spiders and RMSPs (p-value $\approx 10^{-8}$) and spiders versus NXPs (p-value 0.005).

[6] The BIC numbers can be transformed into a Bayes factor (Wagenmakers 2007), which, in this case, would correspond to a value of about 40.0.

[7] We note that an additional spin-down during the Roche-lobe decoupling phase has been suggested by Tauris (2012), which could remove ∼50% of the rotational energy of the accreting pulsar. If this is the case, then the two average spin frequencies are not compatible with each other.

Table 4
Anderson–Darling k-sample Test

Samples	p-value	Reject H_0
LMXBs and RMSPs	2.3e-8	yes
Spiders and RMSPs	5.5e-8	yes
Spiders and NXPs	0.006	yes
AMXPs and RMSPs	0.0014	yes
AMXPs and NXPs	0.038	marginal
Spiders and AMXPs	0.73	no

Note. H_0 is the null hypothesis that both samples are drawn from the same underlying distribution.

The comparison between AMXPs and NXPs (p-value 0.038) instead shows a marginal difference. The only comparison where H_0 is clearly not rejected is the one between spiders and AMXPs (p-value 0.73). The results are summarized in Table 4.

Since both the SW and the k-sample Anderson–Darling test show no detectable difference in the distributions of the AMXPs and spiders, we check whether the reason might be due to the low power of our experiment. The power of the experiment is an estimate of the probability to find a real effect and can be defined as one minus the probability of a Type II error (also called a false negative). For this we need to first estimate the effect size, i.e., the quantitative difference between the two samples. We do this using Hedge's g (Hedges 1981), which is an appropriate estimator when the sample sizes are not equal, under the hypothesis that both samples are drawn from a normal distribution. The g value is 0.1, which indicates that the difference between the means of the the two samples is less than 0.1 standard deviations, i.e., the effect is very small (if present at all). The 95% confidence interval of g is $[-0.43, 0.64]$, which means that if there is "an effect," i.e., a difference between the two distributions, then it is only slightly more likely to have a higher mean spin frequency for the AMXPs over the spiders. The power of our experiment is also very small, on the order of 10%, and therefore it is not surprising that we detect no significant difference between the two samples. We stress that this does not necessarily imply that AMXPs and spiders come from a single population. A more conservative conclusion would be that the sample sizes are too small to detect a (presumable relatively small) difference.

4.1. Selection Effect of the Small Sample Size

Once an accreting pulsar turns into an RMSP, the pulsar will spin down due to magnetic dipole radiation. Therefore it might not be surprising that the distribution of RMSPs and LMXBs appears different. Furthermore, in addition to the magnetic dipole spin-down, some other effects contributing to the difference in spin frequency of the two populations might also be present, see e.g., Kiziltan & Thorsett (2009) and Tauris (2012). We therefore wish to verify whether the observed RMSPs population can reproduce the slow LMXB subpopulation after removing the spin-down effect. We therefore carried out a set of Monte Carlo simulations by assuming a simplistic magnetic dipole spin-down recipe for the RMSPs and "rewinding" the spin of RMSPs by evolving it backward in time. We first select all spin-down values of known RMSPs from the ATNF catalog.[8] Then we remove all sources that belong to globular clusters (since their spin-down measurement is affected by the cluster gravitational potential well). We also remove all sources that belong to the population of spiders as defined in the preceding sections because some (and possibly all) of them are tMSPs and thus are still sporadically accreting, and their spin might be altered by accretion torques. We then simulate 10^5 RMSPs with a spin frequency randomly selected from a distribution equal to the one found in this work (Weibull with $k = 2$ and $\gamma = 292$ Hz). We then simulate 10^5 ages ($\tau_{\rm age}$) following a uniform distribution with boundaries 10^8 and 10^{10} year and spin frequency derivatives ($\dot{\nu}$) with the same distribution as the observed ones. We finally randomly select a pulsar, spin-frequency derivative, and age and increase the pulsar frequency as $\nu_{\rm final} = \nu_{\rm initial} + \dot{\nu}\ \tau_{\rm age}$. The value of $\nu_{\rm final}$ represents the spin frequency at the moment the pulsar stops accreting and is born as a radio pulsar. Then we calculate the mean, median, and standard deviation of our (backward) evolved pulsar population. Of course, the results are sensitive to the maximum allowed spin frequency, so we select three cases, one with $\nu_{\rm max} = 730$ Hz (compatible with the cutoff suggested by Chakrabarty et al. 2003), another with $\nu_{\rm max} = 1000$ Hz, and finally one with $\nu_{\rm max} = 1500$ Hz. When $\nu > \nu_{\rm max}$, we remove our pulsar from the population. We then sample our simulated population by extracting 19 randomly selected objects (since this is the number of sources in our slow LMXBs sample) and calculate the mean, median, and standard deviation of this random sample. We repeat the procedure 10^3 times, and for each run we perform a Shapiro–Wilk test. This latter test is made because, as we have verified in Section 3, the slow subpopulation is compatible with being drawn from a normal distribution.

Our results show that when $\nu_{\rm max} \approx 700$ Hz, our final sample is consistent with a normal distribution in approximately $9/10$ of the cases, with a mean spin frequency value of $\approx$360 Hz, whereas increasing the spin-limit to 1000 and 1500 Hz makes the distribution consistent with a Gaussian in $7/10$ and $3/10$ of the cases, respectively (and mean spin frequencies of $\approx$400 and 435 Hz, respectively). Therefore, although we cannot exclude that a spin cutoff is present at subms periods, our simulations favor a lower cutoff. This is, again, compatible with the current observations.

5. Observations versus Theory

Our analysis suggests the presence of two subpopulations of neutron stars in LMXBs, separated at around 540 Hz, with the fast population having a very narrow standard deviation of only 30 Hz. Indeed, a peak due to clustering of several sources at around 500–600 Hz is already prominent with a by-eye inspection of the spin distribution, see Figure 2. Having drawn this conclusion, we are led to the obvious question: what is the mechanism that leads to the clustering of these fastest-spinning objects? The issue of a "speed limit" for accreting neutron stars is not new. However, it is usually discussed in the context of the spin equilibrium associated with accretion and whether additional torques (in particular, due to gravitational waves) are required to explain the data. The usual evidence in favor of such a mechanism is the absence of neutron stars spinning near the break-up limit. Gravitational-wave emission provides a natural explanation as it does not require particular fine-tuning of the accretion flow and also has no direct connection to the star's magnetic field (Papaloizou & Pringle 1978; Wagoner 1984). The evidence of a clustering at the fastest observed spins

[8] http://www.atnf.csiro.au/people/pulsar/psrcat/.

provides additional arguments in this direction. Whatever the mechanism is that causes the clustering, it must set in sharply as soon as the stars reach a given spin rate. Again, gravitational-wave torques, which scale with a high power of the spin frequency (ν^5 for deformed stars), may lead to exactly this behavior. However, this is a phenomenological argument. It is entirely possible that the answer has nothing whatsoever to do with gravitational waves. The accretion torque may simply become much less efficient as soon as the star reaches above 540 Hz.

These issues have been discussed at length in the literature (Chakrabarty et al. 2003; Patruno 2010; Patruno et al. 2012b) and we will not be able to resolve them here. However, we can phrase the questions in a new light. As we are arguing in favor of a distinct population of fast-spinning LMXBs, we can ask whether the systems that belong to this population are somehow different from their slower-spinning relatives. Are there other observed phenomena that make these systems distinct? If so, do these observations provide a clue to the underlying explanation?

Additional evidence in support of the conclusion that the observed peak at 500–600 Hz is connected to some underlying physical phenomenon and not due to chance may come from the thermal emission. If we consider all 20 neutron star LMXBs with both known spin frequency and a measured surface temperature, then all sources falling in the fast subpopulation have a measured T, while the remaining ones have only upper limits (with the exception of one source plus another marginally outside the 540 Hz cut-point confidence interval (Ho et al. 2011; Haskell et al. 2012; Gusakov et al. 2014). This may be a fairly weak argument, as the temperature upper limits are often not very constraining, but the fact that all sources in the fast population have a measured temperature might indicate that they are hotter, on average, than those in the slow population. If this is, indeed, the case, then any scenario that explains the spin distribution must involve some additional heating. As we show in the following, an unstable r-mode high enough to balance the spin-up torque due to accretion would naturally reheat the star to these levels, and could provide an explanation. Nevertheless, several aspects of this picture remain problematic.

5.1. The Accretion Torque

As a starting point for a more detailed discussion, it is natural to consider the accretion torque. All available accretion models build on the idea that the flow of accreting matter will at some point become dominated by the star's magnetic field. The material torque in this interaction region is usually approximated by

$$\dot{J} = \dot{M}\sqrt{GMR_M} = N_{\rm a}, \tag{5}$$

where the magnetosphere radius is given by

$$R_M = \left(\frac{\mu^4}{2GM\dot{M}^2}\right)^{1/7}, \tag{6}$$

with $\mu = BR^3$ the star's magnetic moment, M the neutron star's mass, and R its radius. The actual location of the transition to magnetically dominated flow is obviously not this precise, so it is common to introduce a factor $\xi \approx 0.1 - 1$ to parametrize the unknowns. We omit this factor here in the interest of simplicity, however.

Meanwhile, outside the corotation radius,

$$R_c = \left(\frac{GM}{\Omega^2}\right)^{1/3}, \tag{7}$$

where Ω is the angular spin frequency of the star, the magnetic field lines rotate faster than the local Keplerian speed, resulting in a negative torque. If $R_M > R_c$ the accretion flow will be centrifugally inhibited and matter may be ejected from the system.[9] It is easy to see that this will happen if the spin period becomes very short, or the rate of flux of material onto the magnetosphere drops. This is the propeller regime. In this phase, accreting matter is flung away from the star, leading to a spin-down torque. In order to account for this effect, we alter the material torque (Ho et al. 2014)[10]:

$$\dot{J} = N_{\rm a}\left[1 - \left(\frac{R_M}{R_c}\right)^{3/2}\right] = N_{\rm a}(1 - \omega_s), \tag{8}$$

where we have introduced the so-called fastness parameter ω_s. Within this model, it is easy to see that the system will not spin up beyond $\omega_s = 1 \longrightarrow R_M = R_c$. In essence, we can estimate (for a given stellar magnetic field, accretion rate, etc.) at which point the system reaches spin equilibrium. Broadly speaking, this accretion model will lead to a flat spin distribution unless we fine-tune the magnetic fields. Without adding features to the model, it is certainly difficult to explain the two populations from Figure 6.

If we observe the star to spin-up during accretion, then we can also use these results to constrain the magnetic field. In order to link the observed X-ray flux to the accretion rate, one would typically take

$$L_X = \eta\left(\frac{GM}{R}\right)\dot{M}, \tag{9}$$

where η is an unknown efficiency factor, usually taken to be unity.

As we progress, it is useful to consider the uncertainties of different torque models. The main uncertainties are well known and relate to three factors:

1. the nature of the accretion flow (thin or thick) in the region close to the neutron star magnetosphere,
2. the disk/magnetosphere interaction (typically encoded in the ξ factor),
3. the estimate of the mass-accretion rate due to transient behavior and the presence of outflows (leading to additional torques, which tend to be equally uncertain).

In a recent work, D'Angelo (2016) reviewed these uncertainties and concluded that the spin equilibrium of AMXPs can be explained without the need of additional spin-down mechanisms like gravitational waves. However, these models cannot naturally explain the pile-up of sources between ≈500–600 Hz, and the conclusion only applies to sources with a dynamically important magnetosphere. As we

[9] Spruit & Taam (1993) showed that material can be expelled from the disk only if $R_M \gtrsim 1.3R_c$. However, this has no influence on the arguments presented in our discussion, and we therefore ignore this difference for the sake of simplicity.

[10] This model is somewhat simplistic, but it captures the main features of the problem. Other prescriptions tend to be more complex, but would lead to broadly the same conclusions.

show later, it is not clear that all LMXBs fall in this category. If the magnetosphere is too weak (or even absent) to truncate the accretion disk, then an additional spin-down mechanism will definitely be necessary.

In addition, Spruit & Taam (1993) and then D'Angelo & Spruit (2010, 2012) showed that the presence of a magnetosphere can alter the regions of the accretion flow close to the corotation radius. The plasma interacting with the magnetosphere absorbs part of the neutron star angular momentum and remains temporarily close to the corotation radius, altering the density/temperature profile of the inner disk regions. Whether such a *trapped-disk* configuration really exists in accreting neutron stars is not yet clear, but there is evidence that it might be a viable option (Patruno et al. 2009; Patruno & D'Angelo 2013; Jaodand et al. 2016; van den Eijnden et al. 2017).

The main differences between a trapped disk and the propeller model were discussed in D'Angelo (2016). In essence, the trapped-disk model predicts a stronger spin-down than propeller during the declining portion of an outburst, which can lead to a lower spin frequency limit. The reason is that the strength of the spin-down in the propeller model scales with the amount of mass available in the inner disk regions, which drops substantially during the declining outburst phase. If a trapped disk is formed in most systems (with a magnetosphere), then the net neutron star spin-up during the outburst may be strongly reduced. Moreover, if a residual trapped disk persists even during quiescence, then there is an additional spin-down torque and the discrepancy between the average spin frequency of accreting neutron stars and RMSPs could be reconciled. However, this involves two relatively strong assumptions: the formation of a trapped disk, and the persistence of this disk once the source accretes at very low mass-accretion rates. Although this is certainly possible (e.g., D'Angelo & Spruit 2010 show that once a trapped disk is formed, it is basically impossible to "untrap" it), this scenario awaits further observational evidence. More detailed calculations are also needed to verify whether a trapped disk can truly explain the cutoff at high frequencies and the presence of two subpopulations of LMXBs.

Another potential problem of the trapped disk is the fact that in at least three AMXPs (Patruno 2010; Hartman et al. 2011; Papitto et al. 2011; Riggio et al. 2011; Patruno et al. 2012a) a spin-down has been measured during quiescence, and its value is entirely consistent with the expected value from a magnetic dipole spin-down of a neutron star with a field on the order of 10^8 G. There is no clear evidence for enhanced spin-down (at least during quiescence) in AMXPs. Moreover, it is difficult to observe spin-down during the declining portion of an outburst because of the short duration of this phase and/or the lower quality of the data. The current upper limits on the spin-down are thus still compatible with both a trapped disk and a propeller model (see, for example, Patruno et al. 2016 for a discussion of the system SAX J1808.4–3658).

5.2. Gravitational Waves

Let us now consider the conditions for gravitational waves to be effectively spinning down accreting neutron stars. Most importantly, the asymmetry of the accretion flow, with matter channeled onto the magnetic poles of the star, should induce some level of quadrupole asymmetry in the star's moment of inertia tensor. Quantifying this in terms of an ellipticity, ε, it is easy to show that the associated spin-down torque will balance the accretion spin-up (given by $N_{\rm a}$ for simplicity) if

$$\varepsilon \approx 7 \times 10^{-9}\left(\frac{\dot{M}}{10^{-9}M_{\odot}\ {\rm yr}^{-1}}\right)^{1/2}\left(\frac{600\ {\rm Hz}}{\nu}\right)^{5/2}. \quad (10)$$

Given this estimate, it is worth making two observations. First of all, the required ε is smaller than the deformation required to break the crust (Haskell et al. 2006; Johnson-McDaniel & Owen 2013), so there is no reason to rule this scenario out. Second, the required deformation is broadly in line with estimates for the quadrupolar temperature/composition asymmetry induced by pycnonuclear reactions in the deep crust (Bildsten 1998; Ushomirsky et al. 2000). We should also keep in mind that the accretion torque (Equation (8)) would be weaker, so the actual ellipticity required could well be smaller.

Another possibility is that the gravitational-wave-driven instability of the inertial r-modes (Andersson 1998; Owen et al. 1998) is active in these systems. Based on the current estimates for the damping mechanisms (shear- and bulk viscosity, viscous boundary layers, superfluid vortex mutual friction, etc.), this seems plausible (see Haskell 2015 for a recent review). In fact, the estimated temperatures for the LMXBs would place a number of known systems in the unstable regime (Ho et al. 2011; Haskell et al. 2012). Moreover, the r-mode instability would readily balance the accretion torque. Quantifying this statement in terms of the mode amplitude α, we need

$$\alpha \approx 6 \times 10^{-6}\left(\frac{\nu}{600\ {\rm Hz}}\right)^{-7/2}\left(\frac{\dot{M}}{10^{-8}M_{\odot}\ {\rm yr}^{-1}}\right)^{1/2}. \quad (11)$$

As this is much smaller than the expected saturation amplitude for the instability (due to the nonlinear coupling with shorter wavelength modes), $\alpha_{\rm sat} \sim 10^{-3}$ (Arras et al. 2003; Bondarescu et al. 2007), the gravitational waves from r-modes should be able to prevent these systems from spinning up. However, this explanation is problematic. In particular, the fact that the instability is expected to grow much larger than what is required to overcome the accretion torque means that neutron stars should not be able to venture far into the instability region. This means that the actual instability limit should be such that all known systems are (at least marginally) stable. However, the current theory does not provide a mechanism that predicts this (at least not without fine-tuning, see e.g., Gusakov et al. 2014; Chugunov et al. 2017). The internal physics is, of course, complex, and we may simply be missing something. It may be worth noting that the typical instability curve shifts by a few tens of Hz if we vary the neutron star mass from $1.4M_{\odot}$ to $2M_{\odot}$. This may be a complete coincidence, but there could be a connection with the width of the distribution of fast-spinning LMXBs. Furthermore, an r-mode growing to amplitudes large enough to affect the spin evolution of the system would also reheat the star, leading to an internal temperature of (Haskell et al. 2012)

$$T \approx 10^8\ {\rm K}\left[\left(\frac{\alpha}{10^{-6}}\right)\left(\frac{\nu}{600\ {\rm Hz}}\right)\right]^{1/5}. \quad (12)$$

It is thus possible that the faster and hotter systems could harbor an r-mode large enough to significantly affect the spin evolution and halt the spin-up of the star (Ho et al. 2011; Mahmoodifar & Strohmayer 2013).

As gravitational waves may well be emitted by accreting neutron stars, it is natural to consider the different mechanisms that may be responsible for the generation of the waves and which constraints we can place on these mechanisms from the observed spin distribution. The fact that the pile-up at higher frequencies seen in Figure 2 would be naturally explained by gravitational waves becoming relevant at a critical rotation rate provides strong motivation for renewed efforts in this direction.

6. Individual Sources

Returning to observations, let us ask which type of sources belongs to the peaked subpopulation, i.e., the fastest rotators. Thus, we identify 10 sources, 4 of which are AMXPs, and the other 6 are NXPs (highlighted in bold in Table 2). Although the spin evolution of NXPs is not known, the fact that we have some AMXPs for which the spin evolution is constrained to a certain degree allows us to identify possible anomalies that might hint at the presence of a common effect, which is triggered (or becomes more prominent) once the neutron stars pass the 540 Hz cut-point.

6.1. Aql X–1

Turning to individual systems, it is perhaps natural to start with Aql X–1. This is one the most mysterious accreting neutron stars because it has been observed to pulsate for just 150 seconds out of a total observed time of almost 2 Ms (Casella et al. 2008). Messenger & Patruno (2015) showed that the pulse fractional semiamplitude went from ≈6.5% (during the 150 s of the pulsed episode) down to less than 0.26% in the same outburst, with a similar order of magnitude upper limits for the other 18 outbursts analyzed. Such an incredibly low duty cycle for the presence of pulsations has spurred a number of hypotheses on their origin, with no model currently being able to explain the observations (see Casella et al. 2008 and Messenger & Patruno 2015 for discussion). Messenger & Patruno (2015) proposed that the only plausible model is one in which a mode of oscillation, with azimuthal number m and frequency $\nu_{\rm mod} \sim 10^{-2} - 1$ Hz, is triggered. However, oscillation modes need to be excited by some event, and so far, there is no evidence for a trigger shortly before the beginning of the pulsing episode. In addition, one would have to understand why any excitation mechanism would single out such a specific low-frequency mode.

6.2. IGR J00291+5934

In the case of IGR J00291+5934, a strong spin-up (Falanga et al. 2005; Burderi et al. 2007) during an outburst of accretion was observed to be followed by a slow drop-off in the spin rate (Patruno 2010; Hartman et al. 2011; Papitto et al. 2011). This behavior would be naturally explained in terms of a dramatic rise in the accretion torque during the outburst and standard magnetic dipole spin-down in between outbursts. This is interesting because it allows us to constrain the star's magnetic field in two independent ways. In quiescence, the magnetic field follows from the standard pulsar dipole formula, while the spin-up torque encodes the magnetic field through the magnetosphere radius, as in Equation (8).

At first glance, the observed spin evolution of IGR J00291+5934 is consistent with the theory (Patruno 2010). The magnetic field estimated from the spin down is in line with the maximum magnetic field inferred from the spin-up phase. However, this argument is not quite consistent (Andersson et al. 2014). Basically, the latter estimate is based on setting $\omega_s = 1$. This does indeed lead to an upper limit on the magnetic field consistent with the system spinning up, but as this is also the condition for spin equilibrium, the actual predicted spin-up rate would be low. In order to quantify the discrepancy, without changing the torque prescription, we can ask how the result changes if we vary the parameters ξ and η. The first of these is not useful, as any value of $\xi < 1$ leads to a weaker spin-up torque. Changing η is more promising as we end up with a larger $\dot{M}$, meaning that the torque is enhanced, but this resolution is quantitatively uncomfortable. We need roughly $\eta \approx 0.1$ in order for the magnetic field estimates to be consistent.

It is also important to note, given the context of this discussion, that an additional gravitational-wave spin-down torque would only enhance the problem.

6.3. XTE J1751−305 and 4U 1636−536

XTE J1751−305 is a slower-spinning system (at 435 Hz, well below our cut-point) that exhibits a spin-up/spin-down behavior similar to that of IGR J00291+5934. In this case, the inferred magnetic field is also inconsistent (Andersson et al. 2014), but slightly less so, in the sense that the inconsistency arises only when considering more realistic accretion torques than those described by Equation (5) (e.g., Ghosh & Lamb 1979; Ho et al. 2014). In this case, we need $\eta \approx 0.25$ in order to reconcile the data. This system is interesting for another reason. Strohmayer & Mahmoodifar (2014a) reported a coherent quasi-periodic oscillation during the discovery outburst. They suggested that the frequency of this feature, roughly 0.57ν, would be consistent with the star's quadrupole r-mode. This is an interesting suggestion, given the possibility that these modes may be unstable in fast-spinning stars. However, the idea is also problematic. First of all, one would need to explain why it is the rotating frame frequency that is observed rather than the inertial frame one (as one would expect). An argument in favor of this has been put forward by Lee (2014). Even so, after confirming that the observed frequency would indeed be consistent with an r-mode (after including relativistic effects and the rotational deformation), Andersson et al. (2014) point out that the gravitational waves associated with an r-mode excited to the observed level would spin the star down very efficiently. This would prevent the system from undergoing the observed spin up.

A similar feature was later found in the faster-spinning system 4U 1636−536 (Strohmayer & Mahmoodifar 2014b), which is above our cut-point. In this case, the observed feature was found at 1.44ν. This would be roughly consistent with an r-mode in the inertial frame. However, the same provisos apply to this observation.

6.4. PSR J1023+0038 and XSS J12270–4859

PSR J1023+0038 and XSS J12270–4859 are two of the three known tMSPs (see Archibald et al. 2009 and Bassa et al. 2014). PSR J1023+0038, in particular, has been monitored in radio and X-rays for the past 10 years (e.g., Bogdanov et al. 2014; Patruno et al. 2014 and Stappers et al. 2014). Spin down has been measured both in radio (Archibald et al. 2013) and during the accretion-powered phase (Jaodand et al. 2016). The spin down during the latter phase is particularly difficult to

explain with the current accretion torque models (see Jaodand et al. 2016 for an extended discussion). The problem is that the system spins down faster when it accretes than it does in quiescence. This led Haskell & Patruno (2017) to suggest that gravitational waves are playing a role in spinning down the neutron star. The quadrupole deformation required to explain the observation is in line with theoretical expectations ($\varepsilon \approx 5 \times 10^{-10}$), making the explanation consistent. Moreover, if it is correct, then we would have a first handle on what level of accretion-induced asymmetries to expect in LMXB. Even if the gravitational-wave signal itself would be too weak to be detectable (given current technology), it would help us model the general population better.

6.5. Orbital Evolution

A first glance at the orbital periods of LMXBs in Table 2 shows that the average orbital period of the fast population is longer than that of the slow population. The mean orbital period of the fast population is about twice as long as the slow-population period (7.6 hr versus 3.7 hr). The same is true for the median (5.9 hr versus 2.8 hr). However, the standard deviations of the two groups are very large, on the order of the mean itself. This indicates that the distributions are relatively flat, with the exception of a peak in the slow population that is composed of ultra-compact binaries. The donor types of the fast population are all made by main-sequence stars, with the exception of IGR J00291+5934 (which has a semidegenerate companion). The companion stars in the slow population are instead composed of a mixture of white dwarfs and semidegenerate and main-sequence companions. The observational evidence therefore suggests that there might be a difference in the evolutionary histories of the two groups, although at the moment it is difficult to quantify this difference since some overlap is observed in the orbital parameters of the two populations.

7. Is there a Magnetosphere?

Another piece of observational evidence comes from the fact that some of these systems might, in fact, lack a magnetosphere strong enough to affect the dynamics of the plasma in the accretion disk. Recent searches for pulsations in several LMXBs have found no evidence for continuous accretion-powered pulsations, with upper limits on the pulsed fraction of ≈0.1%–0.3% (Messenger & Patruno 2015, A. Patruno et al. 2017, in preparation) and down to a few % fractional amplitude for brief pulsation episodes of short duration (0.25 s–64 s; H. Algera & A. Patruno 2017, in preparation). The sample analyzed includes both NXPs and other LMXBs. It is difficult to reconcile such exquisite uniformity of the neutron star surface with the presence of a strong magnetosphere.

However, in the context of the present discussion, this is problematic. If no magnetosphere (or a very weak one) is present, then no centrifugal barrier can act to mitigate the spin up that is due to the transfer of angular momentum from the accreting plasma. Basically, we have to use the torque $N_{\rm a}$ with $R_M = R$. Then the change in spin scales linearly with the amount of mass transferred, and therefore the neutron star continues to spin up (not even spinning down during quiescence if no magnetosphere is present). Moreover, if there is no magnetosphere, then the accretion flow does not lead to asymmetries on the star's surface, and therefore one may not expect the system to develop the deformation required for gravitational-wave emission (although mode instabilities can obviously still play a role, and frozen-in compositional asymmetries may still lead to the neutron star being deformed).

If one accepts that the current upper limits imply the lack of a magnetosphere in these systems, then, given that the non-pulsating LMXBs constitute the majority of accreting neutron stars, there should exist very fast-spinning ($\nu \gg 619$ Hz) objects, which have not been observed so far.

There are two fairly natural possibilities. First, it could be the case that the magnetic field of NXPs has decayed (for example, through Ohmic dissipation). If this is the case, then the field will not reemerge when accretion stops. If this is the case, the NXPs cannot be the progenitors of the RMSPs. Instead, there may exist an unseen population of very fast-rotating neutron stars with no significant magnetic field. These would only be visible during the accretion phase, when burst oscillations are produced on their surface. This might explain why AMXPs and NXPs behave differently and why we do not see any excess in the spin distribution of RMSPs at high frequencies (whereas we see a second subpopulation among the accreting systems).

Alternatively, the absence of a magnetosphere may indicate that the magnetic field is buried by the accretion flow. The problem of magnetic field burial has been considered in detail for young neutron stars, following supernova fall-back accretion (e.g., Viganò et al. 2013) and in some simplified form also for accreting neutron stars (Cumming et al. 2001; Cumming 2008). To get a first idea whether this idea is viable for the much lower accretion rates we are interested in, we can adapt the usual argument.

Taking the work by Geppert et al. (1999) as our starting point, we first consider the depth at which the field would be buried. To estimate this, we balance the timescale associated with the inflowing matter to that of the Ohmic dissipation. We first of all have

$$t_{\rm flow} = \frac{L}{v_r} \tag{13}$$

where L is a typical length-scale of the problem, and in the case of accretion we have

$$v_r = \frac{\dot{M}}{4\pi r^2 \rho} \tag{14}$$

second we need

$$t_{\rm Ohm} = \frac{4\pi\sigma L^2}{c^2} \tag{15}$$

where σ is the conductivity. If $t_{\rm flow} < t_{\rm Ohm}$ then the magnetic field is frozen in the inward flowing matter. The matter piles up faster than the field can diffuse out and hence we have burial. If the accretion stops, the field emerges on the $t_{\rm Ohm}$ timescale associated with the burial depth.

As we do not expect the field to be buried deep, we consider the envelope where the ions are liquid (the electrons are degenerate and relativistic). Then we have (Geppert et al. 1999)

$$\sigma \approx 9 \times 10^{21} \left(\frac{\rho_6}{AZ^2}\right)^{1/3} \ {\rm s}^{-1}, \tag{16}$$

with $\rho_6 = \rho/(10^6\,\mathrm{g\,cm^{-3}})$, which leads to

$$\frac{t_{\rm Ohm}}{t_{\rm flow}} \approx 6 \times 10^5 \frac{L_5 \dot{M}/\dot{M}_{\rm Edd}}{r_6^2 \rho_6^{2/3} (AZ^2)^{1/3}}, \quad (17)$$

where $\dot{M}_{\rm Edd} = 10^{-8} M_\odot\,\mathrm{yr}^{-1}$, A is the mass number of the nuclei, and Z is the proton number. As we are interested in the outer region, it makes sense to consider Fe^{56}, with $A = 56$ and $Z = 26$. We can also set $r_6 \approx 1$, as we are near the star's surface. Then we have

$$\frac{t_{\rm Ohm}}{t_{\rm flow}} \approx 2 \times 10^4 \frac{L_5}{\rho_6^{2/3}} \frac{\dot{M}}{\dot{M}_{\rm Edd}}, \quad (18)$$

where $L_5 = L/10^5$ cm.

We also need to know how the density increases with depth. As we only aim for rough estimates, we use the pressure scale height

$$H = \frac{p}{\rho g}, \quad (19)$$

where the gravitational acceleration

$$g = \frac{GM}{R^2}, \quad (20)$$

can be taken as constant. From Brown & Bildsten (1998), we take

$$H \approx 265 \left(\frac{2Z}{A}\right)^{4/3} \rho_6^{1/3}\ \mathrm{cm} \longrightarrow \rho_6^{2/3} \approx 2 \times 10^5 H_5^2, \quad (21)$$

with $H_5 = H/10^5$ cm. Finally, it makes sense to let $L \approx H$, so we are left with

$$\frac{t_{\rm Ohm}}{t_{\rm flow}} \approx 0.1 H_5^{-1} \frac{\dot{M}}{\dot{M}_{\rm Edd}}, \quad (22)$$

and we learn that the magnetic field is buried up to a density

$$\rho_{\rm burial} \approx 7 \times 10^{10} \left(\frac{\dot{M}}{\dot{M}_{\rm Edd}}\right)^3\ \mathrm{g\,cm^{-3}}. \quad (23)$$

This estimate agrees quite well with (extrapolations from) for example Geppert et al. (1999).

We also have

$$L_{\rm burial} \approx H_{\rm burial} \approx 8 \times 10^3 \left(\frac{\dot{M}}{\dot{M}_{\rm Edd}}\right)\ \mathrm{cm}, \quad (24)$$

which means that once accretion stops, the field will reemerge after

$$t_{\rm Ohm} \approx 10^{10} \left(\frac{\dot{M}}{\dot{M}_{\rm Edd}}\right)^2\ \mathrm{s}. \quad (25)$$

That is, the field will emerge a few hundred years after a star accreting at the Eddington accretion rate enters quiescence.

We would also need know how long it takes to bury the field in the first place. Somewhat simplistically, this follows from the accreted mass corresponding to the burial estimated depth. This is roughly given by

$$\Delta M \approx 4\pi \rho_{\rm burial} R^2 H \approx 4 \times 10^{-6} \left(\frac{\dot{M}}{\dot{M}_{\rm Edd}}\right)^{4/3} M_\odot, \quad (26)$$

and we see that it would also take a few hundred years to bury the field at the Eddington accretion rate. Lower accretion rates, such as those of many of the systems we have considered, would lead to a more shallow burial, leading to the field being buried and reemerging on much shorter timescales.

These estimates obviously come with a number of caveats. A number of complicating factors may come into play, like possible plasma instabilities (Mukherjee et al. 2013) and the tension of the internal magnetic field, which can lead to sharp gradients and reduce the typical length-scale L, thus reducing the amount of mass and the timescale needed for burial (Payne & Melatos 2004, 2007). At least at an approximate level, however, we can see no reason why the temporary burial scenario would not work. A modest level of accretion may lead to a shallow field burial, with the magnetic field reemerging shortly after a system enters quiescence. Furthermore, the deformed magnetic field that is due to accretion may lead to quadrupolar asymmetries and gravitational-wave emission, which, depending on the strength of the magnetic field and the degree of burial, may be large enough to contribute to the spin evolution of the system (Melatos & Payne 2005; Priymak et al. 2011).

In order to improve on these rough estimates, we may wish to consider the possibility that the absence of pulsations is due to local field burial in the polar regions. This would perhaps not affect the burial depth very much, but the stars would not need to accrete as much mass as we have estimated. Moreover, within such a scenario, it could be that there is still a magnetosphere-disk interaction, which could influence the accretion torque.

8. Conclusions

In this paper we have studied the spin distribution of accreting neutron stars and RMSPs. Our analysis (Sections 2–4) shows that the spin distribution of accreting neutron stars can be best described by two subpopulations, one at relatively low frequencies with a mean spin frequency $\mu_1 \approx 300$ Hz, and a fast one with $\mu_2 \approx 575$ Hz. The fast population is strongly peaked within a very narrow range of frequencies ($\sigma_2 \approx 30$ Hz). The two subpopulations are split at around 540 Hz. The fast subpopulation is composed of a mixture of neutron stars with a magnetosphere and others that have shown only burst oscillations so far.

We have shown that all objects in the fast LMXB population have a measured surface temperature that suggests they might be hotter, on average, than the slow population (Section 5). We have discussed various accretion torque models (Section 5.1) and argued that even when considering the various uncertainties that plague them, no model can naturally explain the presence of a fast subpopulation. We therefore considered the role that gravitational waves might have on the neutron star spin if they are efficiently emitted above the cut-point at ≈540 Hz (Section 5.2). We find that different lines of evidence suggest that gravitational waves might be playing an important role to regulate the spin of accreting neutron stars. In particular, gravitational waves provide a natural way to explain the narrow width of the fast subpopulation, provide a physical meaning for the cut-point (i.e., the point above which gravitational waves might be triggered), and might be linked with the mechanism that induces the observed hot temperatures in some neutron stars. However, we also caution that open questions remain on the exact emission mechanism. In particular, some sources that

fall in the fast subpopulation have a behavior that cannot be easily reconciled with gravitational-wave scenarios (Section 6).

Finally, we have suggested (in Section 7) that the lack of a strong magnetosphere (able to generate a strong centrifugal barrier) in most NXPs may be related to the existence of a fast subpopulation. If the magnetic field is buried, it is expected to reemerge once accretion stops. However, field burial affects the accretion torque, and one might expect that it would lead to large numbers of very fast radio ms pulsars. Since such systems are not observed, one might speculate that the magnetic field is not only buried, but also decays during the accretion phase. If this is the case, then there may exist a large unseen population of very fast-spinning neutron stars.

We would like to thank R. Wijnands and C. D'Angelo for useful suggestions. A.P. acknowledges support from an NWO (Netherlands Organization for Scientific Research) Vidi Fellowship.

ORCID iDs

A. Patruno https://orcid.org/0000-0002-6459-0674

References

Andersson, N. 1998, ApJ, 502, 708
Andersson, N., Jones, D. I., & Ho, W. C. G. 2014, MNRAS, 442, 1786
Andersson, N., Kokkotas, K. D., & Stergioulas, N. 1999, ApJ, 516, 307
Archibald, A. M., Kaspi, V. M., Hessels, J. W. T., et al. 2013, arXiv:1311.5161
Archibald, A. M., Stairs, I. H., Ransom, S. M., et al. 2009, Sci, 324, 1411
Arras, P., Flanagan, E. E., Morsink, S. M., et al. 2003, ApJ, 591, 1129
Bassa, C. G., Patruno, A., Hessels, J. W. T., et al. 2014, MNRAS, 441, 1825
Bhattacharyya, S., & Chakrabarty, D. 2017, ApJ, 835, 4
Bildsten, L. 1998, ApJ, 501, L89
Bogdanov, S., Patruno, A., Archibald, A. M., et al. 2014, ApJ, 789, 40
Bondarescu, R., Teukolsky, S. A., & Wasserman, I. 2007, PhRvD, 76, 064019
Brown, E. F., & Bildsten, L. 1998, ApJ, 496, 915
Burderi, L., Di Salvo, T., Lavagetto, G., et al. 2007, ApJ, 657, 961
Casella, P., Altamirano, D., Patruno, A., Wijnands, R., & van der Klis, M. 2008, ApJL, 674, L41
Chakrabarty, D. 2008, in AIP Conf. Ser. 1068, A Decade of Accreting Millisecond X-Ray Pulsars, ed. R. Wijnands et al. (Melville, NY: AIP), 67
Chakrabarty, D., Morgan, E. H., Muno, M. P., et al. 2003, Natur, 424, 42
Chugunov, A. I., Gusakov, M. E., & Kantor, E. M. 2017, MNRAS, 468, 291
Connolly, A. J., Genovese, C., Moore, A. W., et al. 2000, arXiv:astro-ph/0008187
Cumming, A. 2008, in AIP Conf. Ser. 1068, A Decade of Accreting Millisecond X-Ray Pulsars, ed. R. Wijnands et al. (Melville, NY: AIP), 152
Cumming, A., Zweibel, E., & Bildsten, L. 2001, ApJ, 557, 958
D'Angelo, C. 2016, arXiv:1609.08654
D'Angelo, C. R., & Spruit, H. C. 2010, MNRAS, 406, 1208
D'Angelo, C. R., & Spruit, H. C. 2012, MNRAS, 420, 416
Falanga, M., Kuiper, L., Poutanen, J., et al. 2005, A&A, 444, 15
Geppert, U., Page, D., & Zannias, T. 1999, A&A, 345, 847
Ghosh, P., & Lamb, F. K. 1979, ApJ, 234, 296
Guoshen, Y., Sapiro, G., & Mallat, S. 2012, ITIP, 21, 2481
Gusakov, M. E., Chugunov, A. I., & Kantor, E. M. 2014, PhRvD, 90, 063001
Hartman, J. M., Chakrabarty, D., Galloway, D. K., et al. 2003, BAAS, 35, 865
Hartman, J. M., Galloway, D. K., & Chakrabarty, D. 2011, ApJ, 726, 26
Haskell, B. 2015, IJMPE, 24, 1541007
Haskell, B. 2017, arXiv:1703.08374
Haskell, B., Andersson, N., & Jones, D. I. 2006, MNRAS, 373, 1423
Haskell, B., Degenaar, N., & Ho, W. C. G. 2012, MNRAS, 424, 93
Haskell, B., & Patruno, A. 2011, ApJL, 738, L14
Haskell, B., & Patruno, A. 2017, PhRvL, 119, 161103
Hedges, L. V. 1981, Journal of Educational Statistics, 6, 107
Hessels, J. W. T. 2008, in AIP Conf. Ser. 1068, A Decade of Accreting Millisecond X-Ray Pulsars, ed. R. Wijnands et al. (Melville, NY: AIP), 130
Hessels, J. W. T., Ransom, S. M., Stairs, I. H., et al. 2006, Sci, 311, 1901
Ho, W. C. G., Andersson, N., & Haskell, B. 2011, PhRvL, 107, 101101
Ho, W. C. G., Klus, H., Coe, M. J., & Andersson, N. 2014, MNRAS, 437, 3664
Jaodand, A., Archibald, A. M., Hessels, J. W. T., et al. 2016, ApJ, 830, 122
Johnson-McDaniel, N. K., & Owen, B. J. 2013, PhRvD, 88, 044004
Kass, R. E., & Raftery, A. E. 1995, J. Am. Stat. Assoc., 90, 773
Kiziltan, B., & Thorsett, S. E. 2009, ApJL, 693, L109
Lee, U. 2014, MNRAS, 442, 3037
Lorimer, D. R., Esposito, P., Manchester, R. N., et al. 2015, MNRAS, 450, 2185
Mahmoodifar, S., & Strohmayer, T. A. 2013, ApJ, 773, 140
Melatos, A., & Payne, D. J. B. 2005, ApJ, 623, 1044
Messenger, C., & Patruno, A. 2015, ApJ, 806, 261
Mukherjee, D., Bhattacharya, D., & Mignone, A. 2013, MNRAS, 435, 718
Owen, B. J., Lindblom, L., Cutler, C., et al. 1998, PhRvD, 58, 084020
Papaloizou, J., & Pringle, J. E. 1978, MNRAS, 184, 501
Papitto, A., Ferrigno, C., Bozzo, E., et al. 2013, Natur, 501, 517
Papitto, A., Riggio, A., Burderi, L., et al. 2011, A&A, 528, A55
Papitto, A., Torres, D. F., Rea, N., & Tauris, T. M. 2014, A&A, 566, A64
Papoulis, A., & Pillai, S. U. 2002, Probability, Random Variables, and Stochastic Processes (4th ed.; New York: McGraw-Hill Higher Education)
Patruno, A. 2010, ApJ, 772, 909
Patruno, A., Archibald, A. M., Hessels, J. W. T., et al. 2014, ApJL, 781, L3
Patruno, A., Bult, P., Gopakumar, A., et al. 2012a, ApJL, 746, L27
Patruno, A., & D'Angelo, C. 2013, ApJ, 771, 94
Patruno, A., Haskell, B., & D'Angelo, C. 2012b, ApJ, 746, 9
Patruno, A., Maitra, D., Curran, P. A., et al. 2016, ApJ, 817, 100
Patruno, A., Watts, A., Klein Wolt, M., Wijnands, R., & van der Klis, M. 2009, ApJ, 707, 1296
Patruno, A., & Watts, A. L. 2012, arXiv:1206.2727
Payne, D. J. B., & Melatos, A. 2004, MNRAS, 351, 569
Payne, D. J. B., & Melatos, A. 2007, MNRAS, 376, 609
Priymak, M., Melatos, A., & Payne, D. J. B. 2011, MNRAS, 417, 2696
Riggio, A., Burderi, L., di Salvo, T., et al. 2011, A&A, 531, A140
Schwarz, G. 1978, AnSta, 6, 461
Spruit, H. C., & Taam, R. E. 1993, ApJ, 402, 593
Stappers, B. W., Archibald, A. M., Hessels, J. W. T., et al. 2014, ApJ, 790, 39
Strohmayer, T., & Mahmoodifar, S. 2014a, ApJ, 784, 72
Strohmayer, T., & Mahmoodifar, S. 2014b, ApJL, 793, L38
Tauris, T. M. 2012, Sci, 335, 561
Ushomirsky, G., Cutler, C., & Bildsten, L. 2000, MNRAS, 319, 902
van den Eijnden, J., Bagnoli, T., Degenaar, N., et al. 2017, MNRAS, 466, L98
Viganò, D., Rea, N., Pons, J. A., et al. 2013, MNRAS, 434, 123
Wagenmakers, E.-J. 2007, Psychonomic Bulletin Review, 14, 779
Wagoner, R. V. 1984, ApJ, 278, 345
Watts, A. L. 2012, ARA&A, 50, 609