

QUASI-PERIODIC OSCILLATIONS AND BROADBAND VARIABILITY IN SHORT MAGNETAR BURSTS

DANIELA HUPPENKOTHEN¹, ANNA L. WATTS¹, PHIL UTTLEY¹, ALEXANDER J. VAN DER HORST¹, MICHIEL VAN DER KLIS¹,
CHRYSSA KOUVELIOTOU², ERSIN GÖĞÜŞ³, JONATHAN GRANOT⁴, SIMON VAUGHAN⁵, AND MARK H. FINGER⁶

¹ Astronomical Institute “Anton Pannekoek,” University of Amsterdam, Postbus 94249, 1090-GE Amsterdam, The Netherlands; D.Huppenkoth@uva.nl

² Office of Science and Technology, ZP12, NASA Marshall Space Flight Center, Huntsville, AL 35812, USA

³ Sabanci University, Orhanlı-Tuzla, Istanbul 34956, Turkey

⁴ The Open University of Israel, 1 University Road, P.O. Box 808, Ra’anana 43537, Israel

⁵ X-Ray and Observational Astronomy Group, University of Leicester, Leicester LE1 7RH, UK

⁶ Universities Space Research Association, Huntsville, AL 35805, USA

Received 2012 December 5; accepted 2013 March 6; published 2013 April 17

ABSTRACT

The discovery of quasi-periodic oscillations (QPOs) in magnetar giant flares has opened up prospects for neutron star asteroseismology. However, with only three giant flares ever recorded, and only two with data of sufficient quality to search for QPOs, such analysis is seriously data limited. We set out a procedure for doing QPO searches in the far more numerous, short, less energetic magnetar bursts. The short, transient nature of these bursts requires the implementation of sophisticated statistical techniques to make reliable inferences. Using Bayesian statistics, we model the periodogram as a combination of red noise at low frequencies and white noise at high frequencies, which we show is a conservative approach to the problem. We use empirical models to make inferences about the potential signature of periodic and QPOs at these frequencies. We compare our method with previously used techniques and find that although it is on the whole more conservative, it is also more reliable in ruling out false positives. We illustrate our Bayesian method by applying it to a sample of 27 bursts from the magnetar SGR J0501+4516 observed by the *Fermi* Gamma-ray Burst Monitor, and we find no evidence for the presence of QPOs in any of the bursts in the unbinned spectra, but do find a candidate detection in the binned spectra of one burst. However, whether this signal is due to a genuine quasi-periodic process, or can be attributed to unmodeled effects in the noise is at this point a matter of interpretation.

Key words: methods: data analysis – methods: statistical – pulsars: individual (SGR 0501+4516) – stars: magnetic field – stars: neutron – X-rays: bursts

Online-only material: color figures

1. INTRODUCTION

Neutron stars present the best test cases for extreme physics in the high-density regime. A long-standing problem in neutron star physics is our lack of understanding of the neutron star interior, in particular, the dense matter equation of state (Lattimer & Prakash 2007). The conditions in both a neutron star’s core and crust, the composition of its matter and the topology and strength of the magnetic fields remain mysteries that are very difficult to tackle with most conventional methods. The detection of quasi-periodic oscillations (QPOs) in the tails of giant flares from soft gamma repeaters (SGRs) has opened up the possibility of studying neutron star interiors using asteroseismology (see Watts 2012 for a review).

SGRs exhibit regular bursts in the hard X-rays and soft γ -rays ($\lesssim 100$ keV), and very rare giant flares with extremely high isotropic equivalent radiated energy of up to 10^{46} erg (see, e.g., Palmer et al. 2005). Observations of persistent soft X-ray counterparts showing coherent pulsations with large periods of 5–8 s (Kouveliotou et al. 1998, 1999), and the detection of the same periodicities in the tails of the giant flares (Hurley et al. 1999; Palmer et al. 2005), suggested that SGRs are neutron stars. Their behavior is understood within the context of the magnetar model (Thompson & Duncan 1995): in this paradigm the SGRs are isolated neutron stars with exceptionally strong external dipole magnetic fields, largely above the quantum critical limit $B_c = 2\pi m_e^2 c^3 / h e = 4.4 \times 10^{13}$ G (where m_e is the mass of the electron, c the speed of light, h Planck’s constant, and e the absolute value of the electron charge), with internal fields

that may be as high as⁷ 10^{16} G. Giant flares are powered by a catastrophic reordering of the magnetic field (Woods et al. 2001). Since this field is coupled to the solid crust, Duncan (1998) suggested that such large-scale reconfiguration might rupture the crust, creating global seismic vibrations that would be visible as periodic modulations of the X-ray and γ -ray flux. This idea was confirmed by the detection of QPOs in the tails of giant flares from two different magnetars (Israel et al. 2005; Strohmayer & Watts 2005, 2006; Watts & Strohmayer 2006).

If the QPO frequencies can be reliably identified with particular global seismic modes of the neutron star, then they can in principle be used to constrain both the equation of state and the interior magnetic field. This exciting possibility has driven a major effort to develop theoretical models of the vibrations. One major issue is the effect of the strong magnetic field, which threads the crust and the core, giving rise to a spectrum of magneto-elastic oscillation frequencies that includes both continua (which give rise to unusual dynamical responses; Levin 2007) and discrete modes. At present there is some disagreement about the nature and effects of the continua on the resulting frequencies and their longevity (see, for example, van Hoven & Levin 2011; Gabler et al. 2011; Colaiuda & Kokkotas 2012). Uncertainties in the composition of the neutron star crust, and the role of superfluidity, will also have an effect (Watts &

⁷ Supported by period and period derivative measurements; see <http://www.physics.mcgill.ca/~pulsar/magnetar/main.html> for an up-to-date reference list on magnetar spin-down properties.

Reddy 2007; van Hoven & Levin 2008; Steiner & Watts 2009; Andersson et al. 2009). How stellar vibrations cause high amplitude variations in X-ray emission is also not clear, and processes in the magnetosphere may play an active role (Timokhin et al. 2008; D’Angelo & Watts 2012).

A major obstruction to this field of research is the sparsity of data. Since the launch of the first X-ray and γ -ray instruments, only three giant flares have been observed, with just two having data with a sufficient time resolution to detect QPOs. In trying to overcome this lack of observational constraints, it is therefore a reasonable approach to turn to the much more numerous short SGR bursts with lengths of usually less than a second and luminosities around 10^{40} erg s $^{-1}$. Hundreds of SGR bursts have now been observed from many magnetars.⁸ Additionally, several intermediate flares have been detected, with observational properties somewhere between those of the SGR bursts and those shown by giant flares (Ibrahim et al. 2001; Lenters et al. 2003; Guidorzi et al. 2004; Israel et al. 2008). At present the nature of the trigger mechanism for both the giant flares and the shorter bursts is an open question (Thompson & Duncan 1995; Lyutikov 2003; Duncan 2004; Woods et al. 2005; Gill & Heyl 2010; Perna & Pons 2011; Watts 2012; Levin & Lyutikov 2012), but it is certainly possible that the oscillations detected in giant flares may be excited in the smaller events as well. The detection of periodic signals in SGR bursts is, however, restricted by their length: giant flares can last up to hundreds of seconds, whereas a typical SGR burst is shorter than one second, restricting the minimum frequency that can be searched.

To date there has been no systematic search for periodic features in the light curves of the SGR bursts. A search for QPOs in a period of enhanced emission with multiple bursts (a “burst storm”), from the magnetar SGR J1550–5418, carried out using data from the *Fermi* Gamma-ray Burst Monitor (GBM), found no significant signals (Kaneko et al. 2010). El-Mezzeini & Ibrahim (2010) searched a subset of *Rossi X-Ray Timing Explorer* (*RXTE*) data from SGR 1806–20 for periodic features and found some tentative signals; however, there are several points of concern with regard to their methodology, which we address conceptually in Section 2.1 and in detail in Appendix A.

Searching for QPOs in transient light curves is a non-trivial task. Standard methods involving Fourier analysis are defined for infinitely long, stationary processes, owing to the periodic nature of the sine functions used in the Fourier transform. The very nature of a transient event—it has a start, one or more peaks, and an end—complicates the analysis procedure and introduces additional sources of uncertainty. For transient events where the shape of the burst envelope is known, many problems arising from the non-stationarity can be solved either analytically (Guidorzi 2011) or via Monte Carlo simulations (Fox et al. 2001). However, many astrophysical transients do not show a well-behaved burst light curve that is easily reproducible by a simple function. Magnetar bursts in particular exhibit a variety of shapes in the burst envelope, translating into different power spectral shapes in the Fourier domain, which need to be taken into account when deriving significances from the periodogram. This in itself can be interesting, aside from searching for QPOs: the different burst envelope shapes must be created by a physical process in the source, either in the form of noise processes or non-stochastic emission processes, and characterizing the

differences may tell us more about the emission processes at work.

This paper presents the application of a Bayesian method, first derived for long-duration time series data of active galactic nuclei in Vaughan (2010), to timing analysis of magnetar bursts. We choose this method for its capabilities in finding periodicities and QPOs in red-noise dominated periodograms. However, we attempt to answer not only the question of whether there are indeed QPOs, but also to characterize the aperiodic timing behavior of the bursts. Given the uncertainty that exists over the trigger and emission mechanisms for magnetar bursts, such an additional diagnostic will be useful. The method that we develop is general, in the sense that it may be applied to other transients of similar light curve morphology such as gamma-ray bursts (GRBs).

In this paper we illustrate the power of this new method by applying it to timing analysis of a sample of SGR bursts recorded by *Fermi* GBM. Specifically, we search observations taken during an intense flaring episode of the SGR J0501+4516 for periodic and quasi-periodic signals, and characterize the broadband noise processes seen in the bursts. This SGR was discovered on 2008 August 22, when a burst triggered the *Swift* Burst Alert Monitor (Barthelmy et al. 2008; Holland et al. 2008). The same burst subsequently triggered the *Fermi*/GBM, which then recorded high time-resolution data of a total of 29 bursts over 13 days (Lin et al. 2011). An *RXTE* Target of Opportunity pointing revealed a spin period of 5.76 s (Gögüş et al. 2008). With a spin-down period of $(1.5 \pm 0.5) \times 10^{-11}$ s s $^{-1}$, the dipole magnetic field was estimated as 2.0×10^{14} G (Woods et al. 2008; Rea et al. 2009; Gögüş et al. 2010).

In Section 2, we give an overview of the general Bayesian method of searching for periodicities and QPOs in burst data, including a comparison to previous methods. Section 3 presents details of the instrument and the data reduction process for this burst sample. We then characterize both the method’s power and limitations by applying the method first to a large number of simulated observations with and without artificially introduced periodic signals in Section 4. Subsequently, we apply the method to the *Fermi* GBM burst sample from SGR 0501+5416. In Section 5 we first outline the method using one particular burst as an example, before giving results for the whole sample. In Section 6, we discuss the significance of our results, and put them in context with current theoretical models. The purpose of this paper is to lay out the method and test it thoroughly on simulated data before applying it to a small burst sample to demonstrate its power on real data. In future work, it will be applied to a larger sample of short as well as intermediate bursts and giant flares.

2. VARIABILITY ANALYSIS METHODS

Our goals are to search for periodic and quasi-periodic signals in light curves of SGR bursts as well as to characterize the broadband variability behavior of the bursts. To this end, we employ Fourier techniques (van der Klis 1989), extending them for the special case of transient light curves and the presence of broadband variability in our burst light curves. Note here that, following Vaughan (2010), we use the expression *periodogram* for the squared Fourier transform of the data. We assume that it is the sampling of the burst envelope as well as one or several noise processes. We use the expression *power spectrum* to denote the underlying physical process of which the periodogram is a sample, i.e., a realization.

⁸ See, e.g., Woods & Thompson (2006), Mereghetti (2008) for overviews or <http://f64.nsstc.nasa.gov/gbm/science/magnetars/> for a collection of SGR bursts observed with the *Fermi* Gamma-Ray Burst Monitor (*Fermi*/GBM).




Figure 1. *Fermi* GBM observation of burst 080823478⁹ from SGR J0501+4516; left: light curve with a time resolution of 0.001 s. Structure in the burst profile and tail is clearly visible. Right: periodogram of the burst light curve shows flat Poisson noise at high frequencies, and an excess of power over the Poisson level at low frequencies, owing to the complex shape of the light curve.

(A color version of this figure is available in the online journal.)

2.1. Monte Carlo Simulations of Light Curves: Advantages and Shortcomings

Monte Carlo simulations of light curves are a standard tool in timing analysis (see, for example, Fox et al. 2001). The underlying idea is simple: one fits an empirically derived (or physically motivated) function to the burst profile. One then generates a large number of realizations of that burst profile, including appropriate sources of (usually white) noise, such as Poisson photon counting noise. The periodograms computed from these fake light curves form a basis against which to compare the periodogram of the real data. For each frequency bin, a distribution of powers is produced, with a mean that depends both on the Fourier-transformed burst envelope shape and the noise processes introduced into the light curve, while the scatter around that mean follows the noise processes only (a χ^2 distribution with two degrees of freedom—denoted χ_2^2 —for a wide range of noise processes, as long as the central limit theorem holds).

Comparing the observed power in each bin with the distribution of simulated powers in the same bin allows us to make a statement about the probability of the observed power in a particular bin being due to a noise process: if the observed power in a particular bin is a high outlier compared with the distribution of simulated powers in that bin, then the probability of observing the data under the (null) hypothesis of a noise process is $1/N$, where N is the number of simulations performed. If N is large, the observed outlier is unlikely to be produced by the noise process alone.

It should be noted, however, that this test only rejects the null hypothesis, it does not directly give evidence for the alternative hypothesis, i.e., the hypothesis we test the null hypothesis against, to be true. As we will explain in more detail in this section, a faulty assumption for the noise model may well produce significant detections that are, in fact, due to a noise process we have not taken into account appropriately. Conversely, a power that does not exceed the maximum simulated power may still be a significant signal, if the maximum simulated power is a rare event.

Note that the probability derived from the Monte Carlo simulations must be subjected to a correction for the number of frequencies and bursts searched (the number of trials, also called *Bonferroni* correction or “look-elsewhere effect”), since for a large number of frequencies and light curves searched, we would expect a number of outliers that would otherwise be counted as (spurious) detections.

The Monte Carlo method outlined above is versatile and powerful, but it has limitations. The most important limitation comes from our lack of knowledge of the underpinning physical processes producing the observed light curve. Only if the null hypothesis accurately reflects the data—apart from the (quasi-) periodic signal for which we would like to test—is the test meaningful. If important effects that distort either shape or distribution of the powers are missed, then the predictions made will not be accurate, leading to either spurious detections or real signals not being found.

More often than not, especially in the case of short magnetar bursts, we do not have complete information about the emission mechanism. Short magnetar bursts are extremely diverse, varying in light curve shape as well as burst intensity and duration (see, for example, Göğüş et al. 1999, 2000). Unlike for thermonuclear X-ray bursts, where the Monte Carlo technique is widely employed (see, for example, Fox et al. 2001 and Watts et al. 2005), we do not know the underlying aperiodic shape of the light curve (see, e.g., Figure 1 for an illustration of a complicated SGR light curve and periodogram). In order to apply this method, we therefore have to fit the light curve using a parametric approach involving, e.g., higher-order polynomials or splines, and then use this as a template for the Monte Carlo simulations. There is an essential degeneracy in that problem originating from our lack of knowledge: Which features do we fit? Which do we consider to be part of the burst envelope, or potential candidates for a periodic signal on top? This is an arbitrary decision, however one that greatly influences the probability of detecting spurious signals.

The situation is further complicated by the potential presence of so-called red noise: random processes that produce broadband, aperiodic variability, frequently with power-law type shapes in the Fourier domain. Red noise supplies large powers at low frequencies and little at high frequencies, and a realization of a red noise process can appear to the naked eye in the light

⁹ *Fermi*/GBM bursts are numbered by date in the format YYMMDDFFF with YY, the year minus 2000; MM, the two-digit month; DD, the two-digit day of the month; and FFF, the fraction of the day.

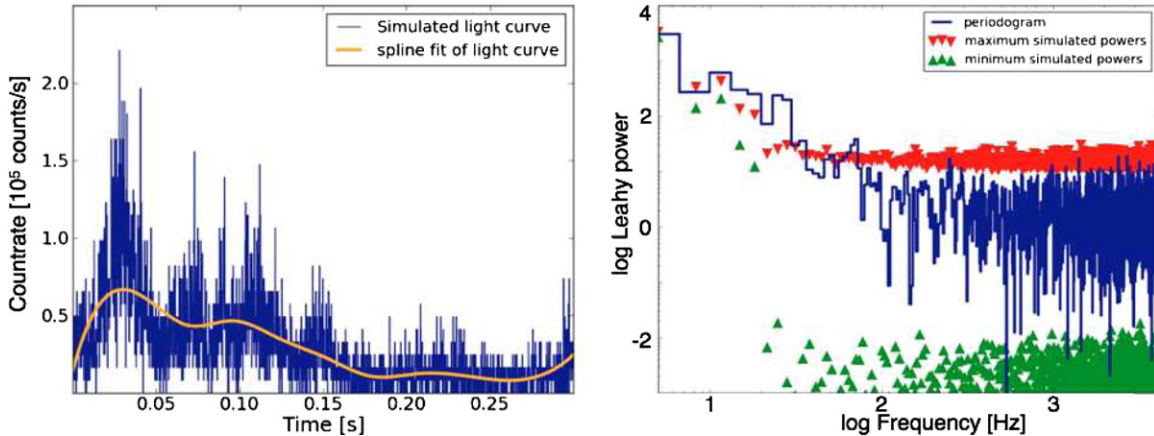


Figure 2. This figure illustrates the effect that inadequate fitting of a light curve containing red noise can have on estimating the significance of potential QPO signals from Monte Carlo simulations. Here, we followed the routine laid out in Timmer & Koenig (1995) to create a light curve from a red noise power spectrum with power-law index of ≈ -2 with Poisson noise added (left panel, solid dark blue line). The input model contains no periodic or quasi-periodic signal. We then binned the resulting light curve to a very coarse time resolution (10 data points) and fit the resulting binned function with a spline (left figure, orange curve). Note that the choice of resolution for the fit is arbitrary: we cannot know a priori from the light curve which features are created by a broadband noise process or a quasi-periodic signal. The binned light curve was used as the basis for Monte Carlo simulations. The figure on the right side shows the periodogram of the fake light curve (dark blue line), with the maximum (red downward triangles) and minimum (green upward triangles) simulated power in each bin. The maximum and minimum powers at each frequency were found by sampling the distribution of powers at that frequency in 1000 Monte Carlo simulations of the light curve fit (orange line in the left panel) with added Poisson noise, i.e., not taking into account red noise, and taking the minimum and maximum samples. Note the spurious detections at 25 Hz and 70 Hz, where peaks in the periodogram clearly stick above the distribution of expected noise powers, even though there is no QPO feature at this frequency: it is entirely created by red noise.

(A color version of this figure is available in the online journal.)

curve like a possibly periodic process (the four peaks in the light curve of Figure 1, for example, seem to mimic periodic behavior, but there is only a very broad bump in the periodogram, and it is impossible to distinguish between a broken power law and a very broad quasi-periodicity in this case). The presence of red noise will significantly alter the distribution of powers in each frequency bin from what it would be if the light curve consisted of a purely deterministic envelope and Poisson-distributed detector noise. Thus, even when the shape of the burst envelope can be adequately modeled by a single, deterministic function, there may nevertheless be false positive detection of single-bin QPO features that are purely due to scatter in low-frequency bins owing to the presence of red noise that has not been taken into account.

In Figure 2, we set up a model that contains only red noise, and find a significant detection despite the fact that there was no QPO introduced into the light curve. This illustrates the fundamental problem with simulating light curves that have an unknown underlying shape. Features may be due to broadband noise features in the light curve, which are not accurately represented by the coarse light curve, and thus not adequately modeled by our null hypothesis. Hence, these features are flagged as a significant detection (with a single-trial probability of 10^{-3}), despite not being due to an underlying (quasi-) periodic process. In the absence of a well-known underlying burst envelope shape or physically motivated models of both noise and burst envelope, it is thus not advisable to apply this method to magnetar bursts (or any transient events with complex light curve shape) in order to derive meaningful conclusions about the presence of QPOs in the light curve. This is one of several shortcomings of the procedure presented in El-Mezeini & Ibrahim (2010). We comment on this paper in the context of our new procedure in greater detail in Appendix A.

We note that the distinction between QPOs and some forms of noise is not clear cut. By convention, one usually defines an upper boundary for the full width at half-maximum (FWHM) $< \nu_0/2$ (where ν_0 is the QPO's centroid frequency; van der

Klis 2006) for the feature to be called a QPO, however, this is a somewhat arbitrary decision. In this work, we consider aperiodic noise only in the form of power laws or broken power laws, as opposed to QPO features, which we assume to be fairly narrow features (following the convention for the FWHM mentioned above) on top of this process. It should be noticed that in principle, one could fit a broadband feature with a wide Lorentzian, thus there is some degeneracy in the modeling. In Section 2.2, we explore whether other simple, plausible models can fit the data, and will describe an alternative, conservative method, based on the assumption that red noise dominates the power spectrum. This is unlikely to be completely true, although many bursts seem to have a red noise component of varying strength, but as we will lay out in the following sections, this assumption is less prone to producing false positive detections, at the cost of increased risk of false non-detections.

2.2. Modeling the Periodogram

Another approach to the problem uses models of the observed periodogram rather than the light curve and assumes any low-frequency broadband variability to be due to a noise process. In a way, this is the other extreme to the approach of using Monte Carlo simulations of light curves: the former is based on the null hypothesis that any power in the periodogram is due to a combination of a deterministic burst envelope, photon counting (white) noise, and a putative (quasi-) periodic signal that is to be detected. When modeling the periodogram, we instead assume that there is no deterministic contribution from the burst envelope and the entire observed emission is due to a noise process. Unless the emission process itself is a noise process, this may not be a valid assumption either. In effect, assuming pure red noise while the light curve has both a noise component and a non-stochastic envelope will cause us to miss weak signals at low frequencies, because they are buried in the much higher variance at low frequencies in a broadband noise process compared with a deterministic burst envelope with only white noise on top.

For the power spectral modeling, we closely follow the Bayesian approach developed by Vaughan (2010). One advantage of the Bayesian framework is the inclusion of our uncertainty in the model parameters (of the assumed low-frequency noise process) in the error estimate of any final quantity, although this still assumes that the functional shape of the spectrum is known; this must be determined separately, as we will lay out below. In addition it provides a statistically rigorous framework to test whether additional model components (such as Lorentzian QPOs) are required by the data (Protassov et al. 2002). In the following, we give only a short outline of the method, and refer the reader to Vaughan (2010) for a thorough discussion.

Following Bayes' rule, the posterior probability of a set of model parameters θ of interest, given the observed data \mathbf{I} and a model H , is defined as

$$p(\theta|\mathbf{I}, H) = \frac{p(\mathbf{I}|\theta, H)p(\theta|H)}{p(\mathbf{I}|H)}. \quad (1)$$

Here, $p(\mathbf{I}|\theta, H)$ is called the *likelihood*, $p(\theta|H)$ the *prior* and $p(\mathbf{I}|H)$ the *marginal distribution* of the data. Note that the latter is often difficult to compute in practice, and only depends on the data. For ratios of posterior probabilities utilizing the same data, the marginal distribution of the data will drop out of the equation, and will consequently be ignored in the following.

We use the Bayesian analog to maximum likelihood estimation (MLE) to fit models to the observed periodogram data and obtain the maximum a posteriori (MAP) estimates of the model parameters. The MAP estimate of a parameter set is defined as

$$\theta_{\text{MAP}} = \arg \max_{\theta} p(\theta|\mathbf{I}, H), \quad (2)$$

where \arg is the argument of the maximized posterior probability, i.e., θ_{max} . The MLE of a given model $S(\theta)$ is computed by maximizing the probability of a data set \mathbf{I} given parameters θ and a model H :

$$p(\mathbf{I}|\theta, H) = \prod_{j=1}^{N/2} p(I_j|S_j), \quad (3)$$

where I_j are the individual powers in the observed power spectrum, and S_j are the powers in the chosen models for a parameter set θ . This is equivalent to minimizing the following function:

$$D(\mathbf{I}, \theta, H) = -2 \log p(\mathbf{I}|\theta, H) = 2 \sum_{j=1}^{N/2} \left\{ \frac{I_j}{S_j} + \log S_j \right\}, \quad (4)$$

sometimes called the deviance (Gelman et al. 2004). Note that Equation (4) is only a valid form of Equation (3) if the data are χ^2 -distributed.

Similarly, we can compute the logarithmic MAP as a combination of prior and likelihood, using Equations (1), (2), and (4):

$$\theta_{\text{MAP}} = \arg \max_{\theta} (p(\mathbf{I}|\theta, H)p(\theta|H)) \quad (5)$$

$$= \arg \min_{\theta} (-\log p(\theta|H) + D(\mathbf{I}, \theta, H)/2). \quad (6)$$

Equation (6) computes the mode of the posterior distribution over parameter space, i.e., the most likely parameters given the

observed data and the model. We use the formalism above for any analysis done in this work.

Without a physically motivated burst emission mechanism, we cannot know what shape the analytic part of the burst envelope takes, or the existence and characteristics of a potential red noise component in the data. Since both the burst envelope and any red noise processes supply power over a large range of frequencies (unlike QPOs, which are confined to narrow regions of frequency space), there is an essential degeneracy in any model we attempt to fit, adding a number of assumptions about the burst shape and red noise properties to whatever inference we try to make. In the absence of any physical motivation, we make a simple, yet probably overly conservative assumption: all broadband power in the periodogram is supplied by a red noise process. The limitations on our inferences that come from this assumption will be further discussed in Section 4, but its main disadvantage is the fact that weak signals in the low-frequency part of the spectrum are likely to be missed. We see this as an acceptable trade-off in return for having a very low false positive detection rate. The advantage of this assumption is that we can treat the entire broadband variability seen in the periodogram as a realization of a noise process and follow an entirely empirical approach to modeling the periodogram: if we find a function that describes the underlying power spectrum well, we can use this as a basis to compute many realizations of this power spectrum and compare these to our observed data. A very broad class of power spectral shapes well represented in nature are power laws:

$$P(\nu) = \beta \nu^{-\alpha} + \gamma, \quad (7)$$

where α is the power-law index, and broken power laws, which can be reduced to Equation (7) by setting $\rho = 0$:

$$P(\nu) = \beta \nu^{-\alpha_1} \left(1 + \left\{ \frac{\nu}{\delta} \right\}^{(\alpha_2 - \alpha_1)/\rho} \right)^{-\rho} + \gamma, \quad (8)$$

where α_1 and α_2 are the power-law indices at low and high frequencies, respectively, and we require $\alpha_2 < \alpha_1$. δ is the break frequency at which the power-law index changes. In both models, β is a normalization term and γ is a constant to account for the presence of white (Poisson) noise in the periodogram. Note that the broken power law is a more general expression of the bent power law used in Vaughan (2010), including the additional smoothness parameter ρ to account for the smoothness of transition between the two power-law components.

Our lack of knowledge of the emission processes in magnetar bursts leads us to choose uninformative prior probability distributions for all model parameters: a $p(\theta) = 1/\theta$ dependence for all scale parameters β , γ , and δ (a Jeffreys prior; see Vaughan 2010, and references therein), and flat priors $p(\theta) = \text{constant}$ for all other parameters. Together, these two classes describe a large range of broadband variability, and are likely to be sufficient in describing the low-frequency behavior of most magnetar burst periodograms. In what follows, we choose our broadband noise model from one of these two. However we include an overall goodness-of-fit test and comment where the model fails this test for individual bursts.

There are several ways to distinguish between models. One often used statistic for nested models, i.e., models where one is a special case of the other, is the likelihood ratio test (LRT). The LRT statistic is based on the ratio of the likelihood values for the two models, the *null hypothesis* H_0 and the *alternative*

hypothesis H_1 :

$$\begin{aligned} T_{\text{LRT}} &= -2 \log \frac{p(\mathbf{I}|\hat{\theta}_{\text{MAP}}^0, H_0)}{p(\mathbf{I}|\hat{\theta}_{\text{MAP}}^1, H_1)} \\ &= D_{\min}(H_0) - D_{\min}(H_1). \end{aligned} \quad (9)$$

In order to decide whether a data set is adequately described by the null hypothesis or not, one often resorts to Monte Carlo simulations of the null hypothesis. In the Bayesian framework, one can employ a certain type of Monte Carlo simulations, so-called Markov chain Monte Carlo simulations (MCMCs), to draw parameter sets from the posterior distribution of possible parameters and generate predictive (fake) periodogram data this way. MCMCs have the advantage that for a stable chain that has converged, the samples generated in that chain will always approximate the posterior distribution of parameters, i.e., the distribution for each parameter that summarizes our entire knowledge of the problem. The posterior distribution for each parameter is obtained by marginalizing (i.e., integrating) over all other model components. In the case where the parameter distributions are non-Gaussian, this allows for far more accurate modes and errors on the individual parameters than standard methods like the covariance matrix. Probably the most widely known and employed MCMC algorithm is the Metropolis–Hastings algorithm (Metropolis et al. 1953; Hastings 1970). However, in many situations, convergence of this algorithm is slow and hence computationally expensive. In this work, we employ the so-called stretch-move algorithm as implemented in python by Foreman-Mackey et al. (2013) in the module *emcee*. *emcee* uses so-called ensemble walkers: a set of Markov chains that is split in two, where each half is evolved using the state of the other half as an input, thereby increasing efficiency in converging toward the posterior distribution of parameters.

The MCMC produces a sample of parameter values (of the null hypothesis, e.g., a continuum model) drawn from the posterior distribution of the data. From this sample we randomly draw parameter vectors and use these to generate fake periodograms. We can then compute a distribution for a statistic T to compare with the same statistic derived from the observed data, T_{obs} . In the case of model selection, the faked data are fit with both a simple and a more complex model (e.g., a power law and a broken power law), H_0 and H_1 , identical to the procedure performed on the observed data. This generates a distribution of T_{LRT} s, which is then used to calculate the corresponding tail area probability (i.e., the probability of obtaining a value of the test statistic that is at least as extreme as the one observed under the assumption of the null hypothesis, also called p -value) for the observed value of $T_{\text{LRT}}^{\text{obs}}$. If this probability is very small (the actual detection level is subject to choice, for example, $p < 0.05$), then the observed reduction in D_{\min} between H_0 and H_1 is larger than can be expected by chance if H_0 were true. More clearly, we reject the null hypothesis in this case, although this test cannot be seen as direct evidence that the alternative hypothesis is true. Conversely, if the probability is not very small, then H_0 is sufficient to describe the data.

Just as data were simulated for assessing the probability of $T_{\text{LRT}}^{\text{obs}}$, we can generate fake data in the form of MCMCs to calculate the distribution of any test statistic we choose. One is particularly sensitive to the kind of model features we are interested in detecting, namely, breaks/bends in the smooth continuum, in that it indicates whether the model provides a good overall fit to the data, or whether additional model components may be needed. This simple statistic for goodness

of fit of aperiodic features, based on the traditional χ^2 statistic, i.e., the sum of the squared standard errors (Anderson et al. 1990; Vaughan 2010), is

$$T_{\text{SSE}} = \chi^2(\mathbf{I}, \hat{\theta}), \quad (10)$$

where

$$\chi^2(\mathbf{I}, \theta) = \sum_{j=1}^{N/2} \frac{(I_j - E[I_j|\theta])^2}{E[I_j|\theta]} = \sum_{j=1}^{N/2} \left(\frac{I_j - S_j(\theta)}{S_j(\theta)} \right)^2$$

and $E[\]$ indicates expectation. This is a good test of overall goodness of fit which will be sensitive to inadequacies in the continuum modeling since all data points are included (as opposed to the T_R statistic, i.e., the biggest outlier in the data, which we will present below in Section 2.2.1).

We have now characterized the broadband noise properties. This information will be the basis for any modeling of the data done in the remainder of this section. In the following, we define a test statistic for outliers in the data, show how to compute posterior predictive p -values for this statistic, and lay out a method to find broader signals, i.e., QPOs, in the data.

2.2.1. Searching for (Quasi-)Periodicities

The procedure for searching for periodicities and QPOs in the data follows the same basic logic applied to the selection of a broadband noise model above. We compute a statistic from the periodogram, then generate a large number (e.g., 1000) of simulated periodograms from an MCMC sample, compute the desired statistic from each simulated periodogram in turn, and finally compare the observed value of the statistic to the distribution generated from the sample of simulations.

In what follows, we have to distinguish very narrow features (with scale parameter, i.e., half-width at half-maximum (HWHM) smaller than or close to the frequency resolution of the periodogram) from broader QPO signals with HWHM that are significantly larger than the periodogram's frequency resolution.

In order to investigate narrow features, a sensible statistic to use is the maximum ratio of observed to model power, or

$$T_R = \max_j(\hat{R}_j), \quad (11)$$

where

$$\hat{R}_j = 2I_j/S_j$$

and I_j and S_j are observed and model powers as defined above. The factor of two normalizes the residuals in such a way that \hat{R}_j will be distributed as χ^2_2 . Drawing from many MCMC simulations, we can compute the tail area probability of T_R from its distribution, or the probability that the observed power $I_{j,\text{max}}$ was produced purely by noise generated by a broadband model. This probability need not be corrected for the number of frequencies searched, as this is already taken into account by the fact that we search the entire frequency range for each simulated periodogram, but if several bursts are searched, it is necessary to correct for the number of bursts searched.

Using the posterior distribution of T_R from the simulations, we can also easily compute the sensitivity to a periodic signal that could have been present in the data, but would have been missed. Sensitivities will be independent of frequency in the white noise range, but strongly depend on frequency in the red noise range, for a simple reason: a signal that would be

highly significant in the white noise range could be buried under strong red noise of equal or larger strength in the low-frequency part of the spectrum, rendering it invisible to our detection methods. We compute sensitivities for the amplitude of a potentially missed periodic signal by finding that value of T_R in our simulated posterior distribution that corresponds to a posterior predictive p -value of 5% or lower. We then compute the corresponding signal powers $I_j = R_j S_j / 2$ and convert these to fractional rms amplitudes at four representative frequencies—40 Hz, 70 Hz, 100 Hz, and 500 Hz—two of which are, for typical magnetar bursts, in the red noise dominated part of the spectrum, one of which is right on the boundary to white noise, and the last of which is safely in the white noise dominated part of the spectrum. It is, in principle, possible to compute sensitivities for every frequency in the periodogram, however for brevity we decided to restrict ourselves to four frequencies where QPOs may be found as an indicator of the rms amplitude a signal would have to have in the different parts of the spectrum in order to be detectable. Note that the sensitivities computed here are different from an upper limit in the sense that they do not require the actual observation of the highest power in the spectrum: the quantity is derived entirely from the simulations, and thus presents a theoretical upper limit to what we could have measured, independent of what we have actually measured in the observed burst itself (see Kashyap et al. 2010 for a discussion on the real meaning of upper limits).

One shortcoming of the T_R statistic is that it optimally detects periodic signals confined to one frequency bin, i.e., either strictly sinusoidal signals or QPOs with a width that is smaller than the frequency resolution of the periodogram. It should be noted that even a strictly sinusoidal signal will distribute power in more than one bin, unless its frequency is exactly the Fourier frequency. This redistribution of the power can account for an average loss of roughly 30% (assuming random distribution of the sinusoidal frequency within a bin) in the frequency bin containing the sinusoid. Broader signals may well be detected, if they are strong enough, but since the power is spread over several bins, this is not an optimal way of detecting broad signals. There are several ways around this restriction. One is to bin (or smooth) the data in some way, and compute T_R for the binned data, assuming that any tentative signal power will now be concentrated in each bin. If we bin the simulated periodograms in the same way, then the test statistic T_R for the binned data is comparable to the distribution approximated by our simulations, and the latter can be used to derive posterior predictive p -values. One can either bin the periodogram with several frequency resolutions and search for QPOs in each, assuming that for a QPO of a given width, all its power will be confined to the central one or two frequency bins if the frequency resolution is coarse enough. Alternatively, one can bin the periodogram geometrically, where the bin size grows with frequency. This way, using the (fairly arbitrary) definition that a QPO must have a FWHM $\Delta\nu$ narrower than $\nu_0/2$, with ν_0 the centroid frequency of the QPO (see, e.g., van der Klis 2006), one accounts for the fact that QPOs at higher frequencies can have a larger range of widths.

An entirely different approach to the problem, which we also include in our analysis, starts out from a model selection point of view, addressing it in a similar fashion to the way one chooses between broadband noise models. Assuming that a quasi-periodicity is simply another type of random process, one may fit the periodogram simultaneously with a broadband noise process as well as a Lorentzian representing a QPO and

compare the resulting fit with that of the broadband noise model only. Following Protassov et al. (2002), we can utilize the likelihood ratio in this case if we compare it to the distribution of likelihood ratios as approximated by MCMC simulations. It is important to note that fitting narrow features with a Lorentzian is statistically challenging (Park et al. 2008; Barret & Vaughan 2012). For quasi-periodic features broader than a single bin, but only distributed over a few bins, we smooth the periodogram with a Wiener filter over 3, 5, and 11 frequency bins and compare the maximum power in each of the resulting smoothed spectra via the same method used for searching for single-bin periodicities presented above. Subsequently, we use the method of fitting Lorentzians only for features broader than five times the frequency resolution of the periodogram. Additionally, we cross-check our detection method for QPOs by searching binned spectra of lower frequency resolution than the original periodogram.

We begin by fitting a Lorentzian plus a constant to the residuals of the data divided by the preferred broadband noise model. At each frequency, we fix the centroid of the Lorentzian to that frequency and let the code fit the scale parameter (HWHM) and the normalization of the Lorentzian. This way, we generate an estimate of the deviance at every frequency. We define the frequency where the MAP estimate is largest as our tentative QPO identification. Note that we use a Jeffrey’s prior ($p(x) \propto 1/x$) on the QPO normalization, and a flat prior on the QPO HWHM that rules out widths outside the range $5\Delta\nu$ to $\nu_0/2$, where $\Delta\nu$ is the frequency resolution of the periodogram and ν_0 the frequency at which we are currently fitting (the centroid frequency). This should help us avoid some of the problems with fitting narrow signals as laid out in Park et al. (2008). Restricting ourselves to HWHM larger than $5\Delta\nu$ is consistent with our choice of fitting the smoothed data: we do not expect the HWHM to be narrower than the binning we have chosen, thus we do not allow optimization to smaller widths. Additionally, we exclude the first and last three frequency bins in order to avoid effects introduced by trying to fit a Lorentzian to one of the edges of the periodogram.

Subsequently, we combine the results from the broadband model fitting and the residual fit that yielded the highest estimate for the deviance, and use both as a starting point for a mixed model to the observed data. We use the same priors as before, but use the best-fit parameter sets of the broadband model fitting and the residual fitting as inputs to the optimization routine. We expect that this will put us fairly close to the global minimum and help us avoid some of the problems associated with trying to minimize a multimodal likelihood function.

Finally, we form a likelihood ratio between the broadband plus QPO model and the broadband model alone. The above procedure is repeated in exactly the same way on a large sample of MCMC generated fake periodograms in order to produce a distribution of likelihood ratios from the broadband model alone. Comparison of the observed likelihood ratio then allows the derivation of a tail area probability that the observed tentative QPO could be generated from the broadband noise model alone. It should be noted that because the model fitting of Lorentzians on the residuals on many simulated periodograms is computationally expensive, we restrict the analysis to a smaller number of simulations (usually $N \sim 500$). The resulting distribution of likelihood ratios will be less reliable, but reliable enough to rule out most cases where there is no QPO present. If the fraction of simulations exceeding our criterion follows a binomial distribution, we can compute the error on the

p -value from the standard deviation of our p -value estimate: for $p < 0.05$ and 500 simulations, the error on the p -value is $\Delta p = \sqrt{p * (1 - p)/N} = 0.0097$. For all borderline cases where the posterior predictive p -value drops below ~ 0.1 , we repeat the analysis with a larger number of simulations (~ 1000 , decreasing the error on the p -value to $\Delta p = \sqrt{p * (1 - p)/N} = 0.0069$) to make our estimate of the posterior p -value on the likelihood ratio more reliable. Since the error on the p -value is not high enough to bring a signal at the 5% level up to $p = 0.1$, we should be able to catch all significant QPOs in this way.

2.3. Summary of Procedure

The Bayesian procedure laid out above has three parts: (1) find the preferred broadband noise model to represent the low-frequency part of the periodogram, (2) search the periodogram for the highest outlier and compare this outlier to those distributed by pure broadband noise to find narrow features, and (3) search for QPOs in the data, using binned data as well as an identical approach for the model selection in the first step. A step-by-step description can be found in Appendix B.

Every step in the analysis follows the same logic: assume a null hypothesis and an alternative hypothesis, compute statistics to summarize the data-model fits for the two different models, generate a sample from this null hypothesis using MCMC, and then compare the distribution of the relevant statistic derived from the sample generated from the null hypothesis to the observed value of that statistic. If the observed value lies in the high-end tail of the distribution, then it is an outlier with respect to the null hypothesis.

Since the entire procedure rests on the correct choice of broadband model, this is the first step of the analysis. The data are fitted with two continuum noise models, which, by definition of the LRT, are required to be nested. The likelihood ratio is the statistic we use to decide which model is preferred by the data. We simulate a large number of fake periodograms from parameter sets drawn from the posterior distribution of parameters, as approximated by a large number of MCMC simulations. Then these fake periodograms are fit with both models again to build a distribution of likelihood ratios from the simple model. We can compute the tail-area probability (p -value) of the observed likelihood ratio to be typical of the distribution (equivalent to asking whether the observed data is sufficiently described by the simpler model) by integrating over the tail of the distribution. If this probability is lower than a chosen significance threshold, then the data is more likely to be drawn from the more complex model hypothesis, which should then be adopted for the rest of the analysis.

Finding periodicities with widths equal to or smaller than a single bin (both smoothed or unsmoothed) follows the same principle. We find the highest outlier in the residuals of the data divided by the best-fit broadband model, and compare this to a distribution of outliers computed in the same way from a large number of periodograms created from an MCMC sample. These fake periodograms do not have a periodic signal (our null hypothesis), thus if the observed outlier were far away from the simulated distribution of outliers, it is unlikely that the outlier has come from this distribution, and we hence favor the alternative hypothesis: that the outlier was indeed produced by a separate physical process.

Finally, we approach the QPO search as a model selection problem. In order not to bias ourselves to a frequency, we fit a Lorentzian at every frequency to the periodogram residuals smoothed over five bins, keeping the centroid frequency of

the Lorentzian fixed while allowing the other parameters to vary. This will give us the MAP estimate for that model at each frequency. We pick the frequency with the highest MAP estimate, and fit a combined broadband model plus Lorentzian to the actual data set. In this case, however, we leave all parameters free, although we use the best-fit parameters from the residual fit as input to the simulations. This minimizes the risk of getting stuck at a local maximum of a multimodal likelihood function. This way, we can compute a likelihood ratio between the broadband noise model with an added Lorentzian component to the pure broadband model fit. We repeat this procedure on a large number of fake periodograms without a QPO and compare the distribution of likelihood ratios to the observed likelihood ratio. Again, if that probability is very small, the observed data are unlikely under the null hypothesis, and the observed feature is unlikely the result of a chance fluctuation from an aperiodic noise spectrum alone.

3. DATA REDUCTION

We now turn to a sample of magnetar bursts and illustrate our method on simulations as well as a small data set as described below.

3.1. *Fermi*/GBM

The GBM is one of two instruments on board the *Fermi* Gamma-ray Space Telescope, launched in 2008 June (Meegan et al. 2009). With its wide field of view and continuous broadband energy coverage between 8 keV and 40 MeV, *Fermi* GBM is well-suited for observing magnetar bursts. The instrument triggers on magnetar bursts, providing high time-resolution data for 30 s before and up to 300 s after the trigger. Three data types were routinely output: CTIME data provide a higher time resolution (64 ms), but low energy resolution (8 channels), whereas CSPEC data provide high energy resolution (128 channels) at low time resolution (1024 ms). Note that CTIME and CSPEC data are available in lower resolution continuously; the quoted numbers are valid for trigger mode only. In this paper, only data of the third type, so-called time-tagged event (TTE) data, were used, since they provide the high time resolution (2 μ s) required for timing analyses, while retaining full spectral resolution as well. For a detailed description of the available data modes and their properties, see Meegan et al. (2009).

3.2. Observations

Fermi/GBM triggered 26 times on SGR J0501+4516 between 2008 August 22 and 2008 September 3, observing 29 bursts. Two of these (080824054 and 080825200) had saturated parts, and were therefore excluded from the analysis due to the rather complicated effects saturation can have on periodograms. Following Lin et al. (2011), we used only Na I detectors with an angle to the source smaller than 50° for each of the 24 triggered and three untriggered bursts. The data were barycentered and channels converted to the mid-energy of each energy bin. The observations were then energy-selected to include only counts between 8 and 100 keV. The lower limit to the energy is set by the detector response (Meegan et al. 2009), the upper limit was found by inspecting energy-resolved light curves and finding no source counts above 100 keV (as indicated by the counts being consistent with the Poisson distribution expected from counting noise). Burst start times and lengths (T90 durations) were taken from Lin et al. (2011), and are summarized in Table 1 of that paper. We added 20% of the burst duration to both ends of the




Figure 3. Light curves of six example bursts from the magnetar SGR J0501+4516 recorded by *Fermi*/GBM. We combined data from all NaI detectors with source angles smaller than 50° to the source. The time resolution corresponds to 0.005 s. Note the strong component of aperiodic variability after the main burst in 080823478 and the differences in peak count rate by almost one order of magnitude between the upper three bursts and the lower three.

(A color version of this figure is available in the online journal.)

burst in order to ensure that we caught the entire burst, and all burst start times and durations in the remainder of this article are to be understood this way. A selection of six bursts is shown in Figure 3, to emphasize the diversity of burst morphologies we encounter.

4. DETECTABILITY SIMULATIONS

We test the power of our detection method on a large number of fake observations: light curves with or without a burst envelope, one or several noise processes, and a periodic signal. We restrict ourselves in the following to detecting periodic or narrow quasi-periodic signals for reasons of computation time, since the QPO detection method introduced in Section 2.2.1 is computationally expensive and hence unfeasible to run on the large number (of the order of several thousand) fake light curves employed here to understand the effect of different components in the light curve on detectability of periodic signals.

While in the previous section we laid out the general principles of the method, our main goal in this section is to characterize how our assumption of pure red noise influences detection rates when this assumption is not true, e.g., in light curves with a strong burst envelope, or when the assumption holds, i.e., in light curves that contain only red noise. We start out with a simple estimate for the importance of the burst envelope on the statistical distribution of the observed powers in our burst sample, and then use one burst from our sample as a template for extensive simulations of light curves into which we artificially inject a periodic signal of varying fractional rms amplitude.

4.1. A Simple Estimate

The method laid out in Section 2.1 is based on the assumption that an observed light curve consists of a deterministic burst envelope—a window function of some kind—and Poisson noise originating in the quantum nature of light when photons impinge on the detector. One can view the deterministic envelope as a

physical process giving rise to the overall shape of the burst, following the same or at least similar functional dependencies for potentially all bursts, and, more importantly, not a realization of a noise process that would alter the general shape of the burst significantly in a stochastic way. This sets it apart from other processes we consider, which contribute to the light curve in a stochastic way. Note, however, that the characteristics given above do not imply that the burst envelope itself may be a realization of a stochastic process, with variable parameters between bursts. The combination of burst envelope and Poisson noise is the null hypothesis against which one wishes to test. One must then ask which part of the light curve is supplied by the burst envelope, and what could be due to a potential periodic process. The presence of red noise clearly renders the fundamental assumption of this method invalid. Assuming pure red noise, on the other hand, lets us avoid a question we cannot easily answer: how much of the observed light curve can be attributed to the burst envelope, and how much to a potential noise process. We do not know a priori what the shape of the burst envelope might be, nor what the power spectral density of the noise process looks like. To first order, we already impose a window function on the periodogram simply by having a short burst: the light curve we Fourier transform is short, equaling a window function that is one between start and end times of the burst and zero everywhere else.

To make a first rough estimate of the effect of the relative strength of the burst envelope, we take all 27 bursts from the sample described in Section 3 and stretch each light curve to have the same total length of 0.2 s in order to make the timescales comparable. We then computed the Leahy-normalized periodogram for each light curve, gathering all powers in 5 Hz bins for all 27 bursts. This yields distributions of powers for each 5 Hz bin. For a noise process, the observed powers should follow a χ^2_2 distribution scaled by the mean power in each bin. Thus, computing the standard deviation for each bin and dividing by the mean power should yield a value close to 1,




Figure 4. Variation between Leahy-normalized periodograms of 27 magnetar bursts. We stretched each light curve to normalize all to the same burst duration by multiplying the photon arrival times by a scaling factor, and then computed periodograms for each of the 27 bursts in our sample. We then gathered powers in 5 Hz bins to yield distributions of powers from all bursts for each of the 5 Hz bins. For the pure red noise assumption to be valid, the resulting distributions should follow a χ^2_2 distribution scaled to the mean of the powers in each bin. Here, we plot the standard deviation in each of the 5 Hz bins divided by the mean of the distribution for each bin. This quantity should be close to 1 for pure red noise. The assumption holds above about 30 Hz, but becomes invalid below. Thus, statements for QPO features below 30 Hz should be interpreted with caution.

(A color version of this figure is available in the online journal.)

if the powers are truly χ^2_2 distributed. This assumption may be broken in two possible ways. First, for steep power laws, the mean in a 5 Hz bin may drop significantly, yielding powers that do not follow a χ^2_2 distribution. Second, for bursts that vary significantly in brightness, the low-frequency red noise component may vary between bursts, and again, the distribution will be altered from our expectation. In Figure 4, we show exactly the dependence on frequency of this quantity. Above 30 Hz, the data seems to follow the noise distributions fairly well, while below 30 Hz, it deviates significantly upward. There are multiple possible reasons for this. While we have corrected for the differences in burst durations, we have not normalized for the differences in fluence. Since burst fluences vary by over an order of magnitude within the sample, this may significantly increase the variation in burst periodograms. Alternatively, differences in burst envelope may account for some of this variation as well. It should be noted that one fundamental assumption underlying this test is the idea that all bursts are governed by the same kind of red noise spectrum. This may not necessarily be true, especially for bursts varying by over an order of magnitude in fluence, and a larger sample of bursts would be needed to draw any solid conclusions about the burst envelope from this kind of analysis. We conclude, for the purpose of our analysis, that the burst envelope seems to become largely unimportant above 30 Hz, and thus above this threshold our assumption of pure red noise is reasonable. Below, one should regard any conclusions drawn about QPOs with caution. However, this simple estimate is only provided to give an idea of where the burst envelope might be important. In the following, we perform detailed simulations of various kinds of light curves, both including and excluding a burst envelope, red noise and periodicities, in order to probe the effect the different components may have on the detectability of QPOs under the assumption of pure red noise.

4.2. White Noise Simulations

In order to test the detectability of (quasi-)periodic signals in complex burst light curves, we simulated a large number of fake observations of bursts and injected a periodic signal with varying frequency and fractional rms amplitude in order to cover a large range of possible signals. The phase of the injected periodic signal was randomized for all simulations to avoid correlations between simulations. For each combination of frequency and fractional rms amplitude, we simulated 100 light curves, which we then ran through our analysis method as if they were real observations. While this is not enough to draw solid statistical conclusions about detectability rates, it gives a qualitative idea of what can be detected and what cannot.

It is important to note that the amplitude of the signal we quote in all of this section is the fraction by which the underlying emission will be modulated. If the underlying signal is flat, then this amplitude corresponds to the fractional rms amplitude as measured from the periodogram. However, if the burst and the periodic signal vary together, such that the fractional amplitude at each point in the light curve is constant, this is not longer true. The reason for this discrepancy lies in the fact that a multiplicative process such as the one described here corresponds to a convolution in the Fourier domain, which will keep the product of the power in both processes—the burst and the periodic signal—constant, but redistributes power toward frequencies close to that of the periodic signal. The result will be a broadened peak in the power spectrum instead of a delta function. While the power in the two central bins will be the constant fractional rms amplitude corresponding to that we would have measured for a flat light curve with a periodicity, the side wings due to the convolution supply power that may be, in practice, indistinguishable from an intrinsically broad QPO. Hence, one would include these side bands into the calculation for the fractional rms amplitude, and in practice measure an amplitude that is larger than that we put in. A characterization of this effect is beyond the scope of this paper; we merely wish to remind the reader that they must take these effects into account when considering the fractional rms amplitudes quoted in this section.

In a first step, we tested the simplest case: (flat) white noise. We created flat light curves with a constant count rate, a periodic signal at 100 Hz of varying fractional rms amplitude, randomized phase, and Poisson noise. In this limit, our method should match standard Fourier analysis techniques and follow the predictions of Groth (1975). In Figure 5, we show the theoretical predictions for white noise together with the results of our simulations. The observed detection rates match the white noise predictions fairly well, and our method deviates from expectations only for a fractional rms amplitude of 5%, but remains within the uncertainty (a 5% error on a detection rate of ~ 0.6 is expected based on 100 simulations). Thus, in the limit of white noise, our method is equivalent to standard Fourier techniques. For any signal at higher frequencies, where slowly varying processes do not distort the power spectrum, we will be as sensitive to a QPO as standard techniques. It should be stressed that the probabilities we quote include a *Bonferroni* correction for the number of frequencies, and are thus not directly comparable to single-trial detection probabilities.

4.3. Pure Burst Envelope Simulations

In order to test the effect of a burst envelope on detectability, we started with the extreme assumption: the burst is dominated




Figure 5. Detection rates for simulated light curves of pure white noise (a light curve of constant count rate), a strictly periodic signal at 100 Hz and Poisson noise. We varied the fractional rms amplitude between 1% and 20% and compared with theoretical predictions calculated using the formalism in Groth (1975). Different symbols and colors indicate (total) detection rates for the unsmoothed periodogram and three smoothing factors included in our analysis. Hollow markers and dashed lines correspond to total detection rates, filled markers and solid lines true positive indicate the number of true positive detections, defined here as the detections at the frequency where the signal was injected (as opposed to the total detection rate, which includes both true positive and false positive detections). Note that for white noise simulations, total and true positive detection rates practically lie on top of each other. True positive detection rates for our simulated light curves match white noise predictions (in green) within the uncertainties (not shown), indicating that our method performs equivalently well to standard Fourier techniques in the white noise regime. Since the rates of true positive detections trace the total detection rates fairly closely, we conclude that our method is not hampered by excessive numbers of false positive detections. For transient phenomena, this regime includes all frequencies above which slowly varying features in the light curve, e.g., red noise, do not dominate the power spectrum.

(A color version of this figure is available in the online journal.)

by a complex burst envelope and Poisson statistics, lacking any red noise. In order to generate the complex burst envelope, we smoothed the light curve of 080823478 (see Figure 1) to an arbitrary cut-off frequency, in our case roughly 35 Hz, creating a smooth light curve with several broad peaks. We included Poisson noise in each simulation, and a periodic signal at 40, 70, and 100 Hz with an absolute amplitude that varied with the flux of the burst envelope such that the fractional root-mean-square amplitude remained constant. We varied the fractional rms amplitude between 1% and 20%, and ran simulations without a periodic signal in order to quantify false positive detection rates.

In Figure 6 we present a selection of periodograms of the resulting combined light curves without Poisson noise. Most notably, the multiplication of a complex burst envelope with a periodic signal in the light curve leads to a significant broadening of the periodic signal in the Fourier domain, including wings and side-lobes. In this scenario, signals below 40 Hz should be undetectable, whereas at higher frequencies detection rates should approach what we would predict for pure white noise. The combination of the envelope from the fit solution and a periodic feature additionally changes the slope of what our method will interpret as broadband noise in the case where a periodic signal is located just at the break where the noise powers drop toward the white noise level. We predict that this will lead to decreased detection rates as well.

Figure 7 presents total detection rates as well as true positive detection rates (both out of 100 simulated bursts) at 40, 70, and




Figure 6. Periodograms of simulated light curves including a complex burst envelope shape and a periodic signal, without Poisson noise. We varied the frequency from 20 Hz to 100 Hz, and kept a constant fractional rms amplitude of 20%. The signal at 20 Hz (green) is almost invisible, and likely impossible to distinguish from the underlying burst envelope, hence we do not consider signals this low in the following analysis. For a signal at 40 Hz (orange), our method is unlikely to be able to distinguish between the broadened periodic signal and the burst envelope, causing low detection rates for even high rms amplitudes. As the signal moves to higher frequencies (magenta: 70 Hz, dark blue: 100 Hz), detection rates converge toward the detection rates predicted for white noise (black dashed line).

(A color version of this figure is available in the online journal.)

100 Hz for five different fractional rms amplitudes and both unsmoothed and smoothed periodograms. True positive detection rates are measured as detections at the frequency of the injected periodic signal. As predicted, detection rates increase toward higher frequencies, where the envelope becomes unimportant. While there are no true positive detections for the unsmoothed periodograms at low frequencies, the number of true positives for the smoothed periodograms closely traces the total detection rates. This points toward a significant broadening of the periodic signal as a result of its convolution with the burst envelope. In general, at frequencies below 100 Hz, detection rates for periodograms smoothed to three or five bins are higher than for the unsmoothed periodograms or those smoothed over 11 bins. At higher frequencies, the detection rates for unsmoothed periodograms approach those of the smoothed periodograms. At 5% fractional rms amplitude or below, there may be a significant contribution from false positive detections, which vanishes for higher fractional rms amplitudes and higher frequencies.

4.4. Envelope Plus Red Noise Simulations

We used light curves composed of both a burst envelope with a simple functional form as well as a power-law red noise component, combined with Poisson statistics and a periodic signal of varying frequency, amplitude, and randomized phase. Although both the burst envelope and the red noise shape are guesses and to some degree degenerate—a different choice of burst envelope shape may lead to a different estimate of the red noise power spectrum—we believe that this type of light curve is likely to be more realistic than the pure burst envelope model, based on the observation that bursts consistently rise faster than they decay (indicating the presence of a deterministic component) and the large variety in burst shapes otherwise

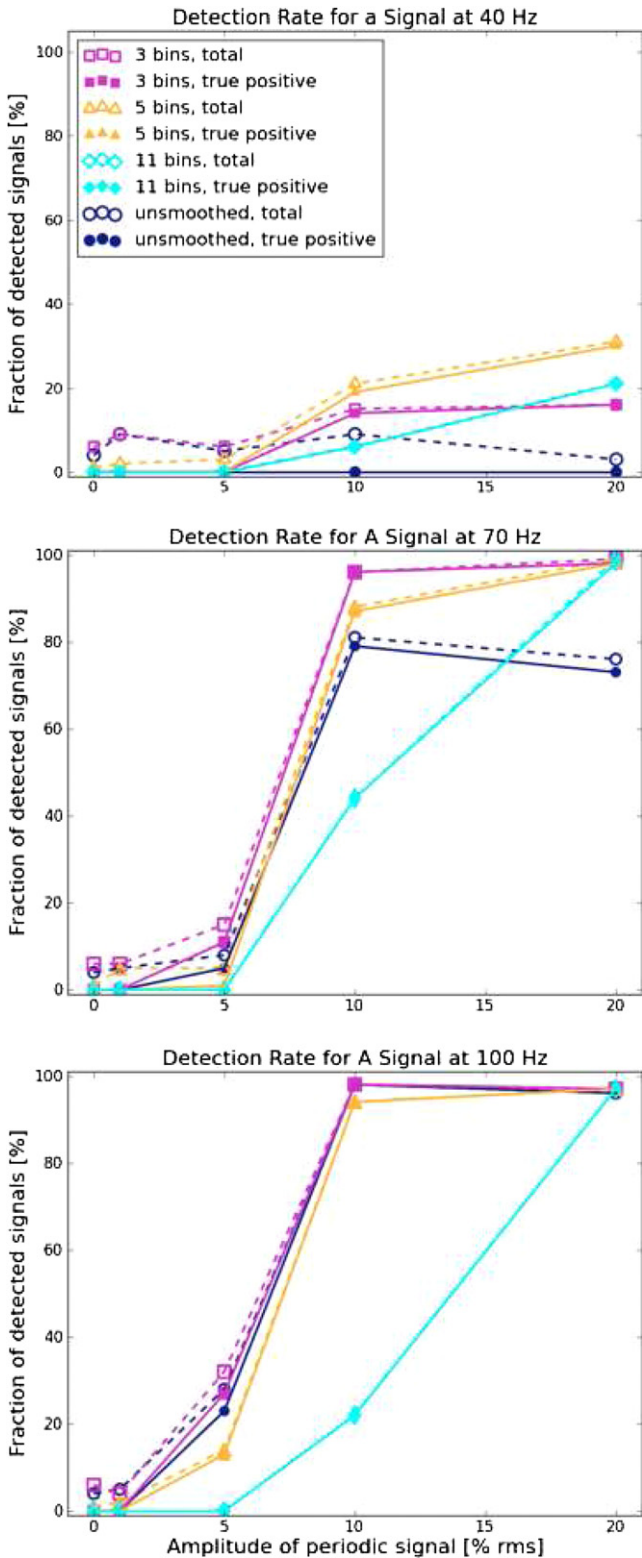


Figure 7. Detection rates for periodic signals at various frequencies on top of a smoothed burst envelope. The envelope is generated by smoothing the light curve of burst 080823478. This introduces a sharp cut-off at around 35 Hz (corresponding to the timescale of smoothing) that is unlikely to be seen in real light curves. Detection rates at 40 Hz are strongly suppressed, indicating that at low frequencies, chances of detecting even a very strong signal are small. Note that while there are no true positive detections (as defined in the caption of Figure 5) at 40 Hz for the unsmoothed periodograms, the number of true positives for the binned periodograms closely traces the total detection rates. Detection rates increase with frequency and fractional rms amplitude as expected.

(A color version of this figure is available in the online journal.)

(indicating the presence of some form of variability on many timescales that is typical for red noise).

For the largely qualitative conclusions we wish to draw here, neither the exact shape of the burst envelope, nor the exact parameters of the red noise are important, although both are interesting questions in their own right and beyond the scope of this paper. Instead, we wish to give a representative example of the general behavior one may expect when applying the method presented here to light curves of transient events with complex light curves.

As before, we use burst 080823478 as a template burst on which to base our simulations. This burst presents an interesting profile, with a main spike and several features that are reminiscent of red noise (see Figure 8). We cannot exclude the possibility that the latter is actually due to a more complicated emission mechanism, and can only state that its timing properties are consistent with red noise. We fit the entire light curve with several models, all based on a fast-rise, exponential decay (FRED) profile of the type

$$f(x) = A\lambda \exp\left(\frac{-\tau_1}{(t-t_s)} - \frac{(t-t_s)}{\tau_2}\right). \quad (12)$$

Here, τ_1 and τ_2 are the rise and decay timescales, respectively, t_s is the burst start time, A is a normalization constant (or burst amplitude), and $\lambda = \exp(2(\tau_1/\tau_2))^{1/2}$ (Peng et al. 2010). This model has been successfully applied to GRBs in the past and appears to be a reasonable first assumption for magnetar bursts with their exponential-like tails and shorter burst rise times compared with the decay timescale. Because the burst has a sharp initial spike and then a long, relatively flat, but very variable tail, a single FRED profile has trouble fitting the entire light curve well: it can either fit the main spike with its sharp decay, or the long tail, but not both together. Hence, we implement two more complex hypotheses: a FRED profile to account for the initial spike, on top of a linear function modeling the slow decay, as well as a model with two FRED components. The former does not fit the beginning and end of the burst well, as it does not drop to background noise level as it should at the start and end of the burst. The latter provides the best fit of the three, but is the model with the largest number of parameters and requires an explanation for the origin of the additional FRED component. The periodogram presented in the right panel of Figure 8 shows how important the choice of burst envelope is for disentangling red noise and deterministic envelope at low frequencies: if the burst envelope could be modeled with a single FRED profile, the low-frequency part of the power spectrum would be entirely dominated by red noise, and the assumptions we make in Section 2.2 hold to a fairly high degree. If the model has additional components, however, either in form of a straight line, another FRED profile or another type, this component will dominate the periodogram up to about 60 or 70 Hz. As a consequence, assuming pure red noise in this part, if the more complex hypotheses were true, our assumption of pure red noise might be a poor one in this frequency range. For what follows, we choose the combination of FRED and linear model, to keep our model as simple and the number of free parameters as low as possible.

Having chosen a model for the overall burst morphology, we make an estimate of the red noise part of the power spectrum: we de-trend the light curve by dividing the light curve by the lightcurve model fit and compute the periodogram of the residuals. De-trending in this way will give us a light curve fluctuating around a mean of 1, and in line with our assumption,




Figure 8. This figure shows the effect that different choices of burst envelope have on conclusions for the relative strengths of the red noise component and envelope component in the low-frequency part of the periodogram. Left: light curve of burst 080823478 with single-component fast-rise exponential decays (FRED) fit (magenta), a FRED profile with a straight line to account for the long tail (orange), and a combination of two FRED profiles (green); right: periodogram of the same burst and all three fits from 4 to 600 Hz. There are clear differences in how strong the envelope is at low frequencies. The Poisson noise level is shown in a light blue dashed line for comparison.

(A color version of this figure is available in the online journal.)

we consider the variance around that mean to be red and white noise only. We fit a power law (for the red noise) plus a constant (accounting for Poisson noise) to the periodogram of the residuals, and take the resulting power-law fit as a template power spectrum to simulate red noise light curves from. Using the method from Timmer & Koenig (1995), we simulate 100 light curve realizations from red noise power spectra only. Note that light curves simulated according to Timmer & Koenig (1995) will have entirely uncorrelated phases, which may introduce a bias into the light curves if this does not accurately reflect reality. More importantly, light curves generated this way are distributed around a mean of zero. In reality, light curves with negative count rates are unphysical, however, any transformations applied to the simulated light curve will result in correlations between phases in the periodogram. We choose a method following Uttley et al. (2005) to generate log-normally distributed light curves that have no negative data points, accepting that the assumption of log-normally distributed light curves introduces a potential bias into our simulations via the correlations it introduces between the phases in the periodogram. Each red noise light curve is combined with the template assumed for the burst envelope. This will provide us with 100 light curves including both an envelope and a red noise component that we can use as fake observations to be analyzed through our method. We add periodic signals in the same way as in Section 4.3, however, since the red noise we included in the simulations does not drop off as sharply as the burst envelope in the previous section, the periodogram at 100 Hz is still contaminated by red noise. Thus, here and in the following section, we also include simulations with a periodic signal at 150 Hz to probe the white-noise dominated region, and run each light curve through our Bayesian detection method.

The detection rates for various frequencies are shown in Figure 9. Signals at 40 Hz are not detectable, with either no or very few true positive detections. Detection rates rise for higher frequencies toward the white noise limit, although even for 150 Hz, a signal at a fractional rms amplitude of 5% or 10% is still somewhat suppressed. The effect of red noise on detectability is more wide-spread in frequency compared with the case of a pure burst envelope, where the power due to broadband variability drops sharply around 40 Hz. Binned or smoothed periodograms are generally better at detecting

periodic signals combined with a burst envelope and red noise, with detection rates in the unsmoothed periodogram only approaching the performance of the smoothed spectra (and at the same time the white noise limit) for high frequencies. This again is due to the broadening of the periodic signal in the convolution with the burst envelope and red noise. As for the pure burst envelope simulations, total detection rates contain a significant contribution from false positive detections, and true positive detection rates approach total detection rates for high frequencies and large fractional rms amplitudes. We note that the false positive detection rate at low fractional rms amplitude seems uncharacteristically high for the three-bin periodogram at all frequencies in these simulations. At present, we do not understand the reason for this. It is possible that the broadband fitting is not entirely reliable in some of the simulations. In practice, the results from the broadband fitting of the periodograms of real bursts are checked to ensure minimization to a global minimum. Additionally, we believe the number of false positive detections is easily corrected for by requiring that, in practice, signals at low fractional rms amplitude need to be significant in at least two different smoothed or binned periodograms.

4.5. Pure Red Noise

As a last example, we test detectability under the hypothesis that our light curve has no deterministic element at all, and consists purely of red noise. We generate light curves using the method from Timmer & Koenig (1995), using the fit to the periodogram of burst 080823478 as a template to achieve comparable burst length, fluence, and rms variability. Again, we introduced a periodic signal of constant rms amplitude and randomized phase into each light curve, changing the fractional rms amplitude of the signal for different simulations. Figure 10 presents the detection rates for the simulations of pure red noise. Detection rates for the pure red noise case are comparable to the case where there is an envelope component, indicating that the presence of a burst envelope does not significantly alter detectability of a periodic signal if there is a red noise component present. Low-amplitude signals are suppressed in the case where the light curve consists only of red noise compared with simulations including a burst envelope. As for the combined light curves, detection rates for one or more




Figure 9. Total (hollow markers and dashed lines) as well as true positive (filled markers and solid lines; for a definition see caption of Figure 5) detection rates for a periodic signal on top of a combined burst envelope and red noise light curve. The different panels correspond to different frequencies of the periodic signals, from 40 to 150 Hz. As before, we include detection rates of both the unsmoothed and smoothed periodograms. Note that for the case where a periodic signal is combined with a burst envelope and red noise, the periodic signal will be significantly broadened, and thus one can more successfully detect these signals in binned or smoothed periodograms. This is especially true at 70 Hz. For higher frequencies, the difference becomes smaller, but is still appreciable. (A color version of this figure is available in the online journal.)

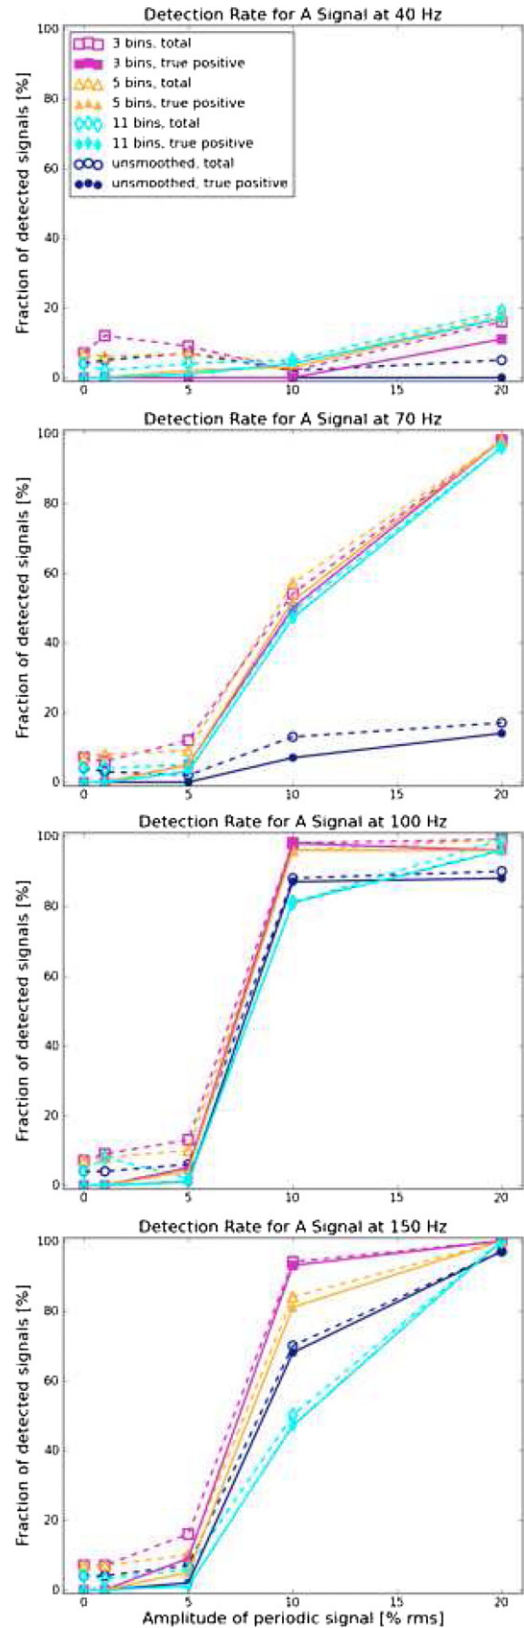


Figure 10. Detection rates for a periodic signal on top of a red noise light curve. The different panels correspond to different frequencies of the periodic signals, from 40 to 150 Hz. As before, we include total (hollow markers and dashed lines) and true positive (filled markers and solid lines, definition in caption of Figure 5) detection rates of both the unsmoothed and smoothed periodograms. Note that signals at 5% rms amplitude or below are generally suppressed, as compared to the white noise detection limit accessible with standard Fourier techniques. (A color version of this figure is available in the online journal.)

of the binned periodograms are always equal or higher than for the unbinned periodograms, and below 70 Hz, even strong signals become nearly undetectable. For high frequencies, the detection rates approach the white noise predictions and are slightly higher than predictions for simulations including both a burst envelope and a red noise component, although detection rates for a fractional rms amplitude of 5% are suppressed even for a 150 Hz periodic signal. This indicates that the red noise component is still significant at these frequencies, for the given input spectrum. In reality, the frequency at which the white noise detection limit holds will depend on the specific burst variability for each observation.

4.6. Conclusions from the Simulations

The different types of simulations allow us to draw two important conclusions for QPO detection: (1) in the limit of a flat light curve, translating to a pure white noise power spectrum, our method does equally well compared to standard Fourier techniques, and (2) detection rates for more complex light curves depend on the underlying emission mechanism. Even a periodic signal may be significantly altered (i.e., broadened) by the presence of a burst envelope and/or red noise, if the periodic signal is modulated by these aperiodic processes, and this broadening will depend specifically on the shape of burst envelope and the red noise parameters, as well as the relative strengths between the two. A significant broadening may in turn affect detectability when it alters what our method interprets as broadband noise, decreasing detection rates even for high fractional rms amplitudes. For the small sample of bursts from SGR J0501+4516, a simple, crude estimate comparing the standard deviation to the mean in small frequency bins across all bursts in the sample reveals that the assumption of red noise holds reasonably well for frequencies above 30 Hz. Below, the assumption may either be broken by the presence of a burst envelope, or alternatively the bursts may be sufficiently varied in red noise properties to produce the observed increase in standard deviation about the mean. This need not mean that our assumption of red noise is invalid in this regime, simply that we do not know this to be true or false. Hence, we caution the reader to interpret results at frequencies this low with care and with the conclusions of the burst simulations in mind.

In general, signals below 70 Hz or so will be very difficult to detect, unless they have fractional rms amplitudes of above 10%. This is not impossible, given the high fractional rms amplitudes observed from the 2004 giant flare (see Watts 2012, Table 1 for an overview), however, even for high amplitudes false non-detections may still occur. We recognize these issues as a shortcoming of the presented method, however, in the absence of any physical model or empirical evidence for a consistent burst envelope structure, we opt for the conservative approach presented here. Thus, we caution the reader to keep the effects described above in mind when interpreting the posterior p -values and sensitivities quoted in Section 5 below.

On the other hand, we have also shown that while the sensitivity of our method depends on the type of light curve analyzed, detection rates for both light curves combining an envelope and red noise—the case we consider most likely for SGR bursts—and for pure red noise light curves, detection rates are quite similar, within the uncertainties, indicating that the additional envelope component does not significantly alter our chances of detecting a signal. Hence, unless light curves are purely deterministic, our method will yield fairly reliable results. At the same time, the false positive detection rate is

generally low in most simulations, which is one of the key goals of developing this technique for transients. False positive detections can be dealt with by requiring detection in more than one smoothed or binned periodogram.

Finally, we would like to make two notes. First, the findings above are based on the assumption that a periodic signal will vary with the underlying light curve, that is, that the fractional rms amplitude, rather than the absolute amplitude, remains constant. This assumption, of course, need not be true. Instead, the absolute amplitude may be constant, in which case a periodic signal would truly remain confined to two frequency bins, and none of the broadening would occur. Second, we also note that the false positive detection rate is low, as expected for a conservative approach. For 100 fake observations, and a detection threshold of $p < 0.05$ for each observation, we find roughly five false positive detections in most runs, exactly as expected. The sole exception is the run that combines a burst envelope and red noise. At present, it is not clear what causes this increase in false positive detection rate. The false positive detection rate can be lowered by tightening the detection threshold to a smaller probability, at the cost of increasing the upper limit to the amplitude of a signal we might have missed, or, in other words, increasing the risk of false non-detections.

5. RESULTS

We computed light curves and periodograms for all 27 bursts (Sections 5.2 and 5.3) as well as time segments before and after each burst (Section 5.1). In each case, we produced a light curve by binning the TTE data to a time resolution of $1/2\nu_{\text{Nyquist}} = 1.22 \times 10^{-4}$ s, corresponding to a Nyquist frequency of $\nu_{\text{Nyquist}} = 4096$ Hz. The time resolution may be arbitrarily chosen, as long as it remains poorer than the time resolution of the detector itself, i.e., $2 \mu\text{s}$ for *Fermi* GBM, although searches with high frequency resolution up to large Nyquist frequencies quickly become computationally expensive. We chose the time resolution based on the Nyquist frequency of interest: we do not expect any signals above 4000 Hz from neutron star seismic oscillations (McDermott et al. 1988). We search both the unbinned periodogram as well as the same periodogram binned to integer multiples (3, 5, 7, 10, 15, 20, 30, 70, 100, 200, 300, 500, and 700) of the frequency resolution of that burst, i.e., the actual frequency resolution of the binned periodograms depends on the frequency resolution of the unbinned periodogram. Additionally, we smooth the spectra with a Wiener filter with different smoothing factors (3, 5, and 11) and compare results of the search of binned periodograms with searches across the smoothed periodograms. Note that while computing sensitivities for binned periodograms is statistically straightforward, doing so for a convolution of the periodogram and a smoothing function is not, hence all sensitivities quoted refer to either the unbinned or binned periodogram, but never the smoothed one.

5.1. Checking for Spurious Timing Signals

Fermi GBM sees the entire unocculted sky at any given point in time. Therefore, the γ -ray background can be rather complex, and one must exclude that a background source supplies significant variability to the burst periodogram. To this end, we performed timing analysis on 1 s and 10 s long segments before and after each burst as well as on the bursts themselves. The light curves constructed out of these segments were Fourier transformed and normalized in order to produce periodograms

Table 1
Posterior Summary of Results for the Broadband Modeling for All Bursts in the Sample

Burst ID	Length (ms)	Fluence (erg cm ⁻²)	Model	$p(\text{LRT})$	$p(\text{SSE})$	α_1	α_2
080822529	86	7.05	flat	0.31 ± 0.01	0.82 ± 0.01
080822647	216	19.3	PL	0.10 ± 0.01	0.70 ± 0.01	$2.41^{+0.39}_{-0.35}$...
080822981	30	4.41	PL	0.16 ± 0.01	0.84 ± 0.01	$2.42^{+1.44}_{-1.23}$...
080823020	66	25.02	PL	0.99 ± 0.002	0.55 ± 0.02	$2.72_{-0.65}$...
080823091	676	82.84	flat	0.59 ± 0.01	0.02 ± 0.005
080823174	447	14.3	PL	0.09 ± 0.008	0.82 ± 0.01	$1.93^{+0.91}_{-0.71}$...
080823248	272	22.18	PL	0.29 ± 0.01	0.85 ± 0.01	$4.19^{+1.95}_{-1.50}$...
080823293a	189	20.10	PL	0.11 ± 0.01	0.75 ± 0.01	$2.65^{+0.61}_{-0.60}$...
080823293b	38	9.54	flat	0.09 ± 0.009	0.95 ± 0.006
080823319	142	19.42	PL	0.16 ± 0.01	0.78 ± 0.01	$2.79^{+1.04}_{-0.70}$...
080823330	192	67.05	PL	0.47 ± 0.02	0.18 ± 0.01	$2.71^{+0.36}_{-0.34}$...
080823354	96	8.62	PL	0.51 ± 0.01	0.89 ± 0.01	$3.35^{+1.37}_{-1.06}$...
080823429	94	14.24	PL	0.09 ± 0.009	0.97 ± 0.005	$4.17^{+1.56}_{-1.28}$...
080823478	264	512.6	BPL	0.003 ± 0.002	0.13 ± 0.01	$2.16^{+2.09}_{-0.84}$	$5.21^{+2.41}_{-3.25}$
080823623	220	21.12	PL	0.30 ± 0.01	0.23 ± 0.01	$1.97^{+0.55}_{-0.46}$...
080823714	406	33.04	PL	0.58 ± 0.02	0.57 ± 0.02	$1.77^{+0.34}_{-0.31}$...
080823847a	264	78.61	PL	0.10 ± 0.009	0.63 ± 0.02	$2.55^{+0.33}_{-0.30}$...
080823847b	108	33.09	PL	0.92 ± 0.008	0.96 ± 0.005	$2.48^{+0.55}_{-0.48}$...
080823986	60	4.37	flat	0.22 ± 0.01
080824346	34	5.70	PL	0.99 ± 0.003	0.78 ± 0.01	$3.02^{+3.45}_{-1.58}$...
080824828	82	6.39	flat	0.42 ± 0.02	0.86 ± 0.01
080825401	128	104.8	PL	0.14 ± 0.01	0.75 ± 0.01	$2.25^{+0.24}_{-0.22}$...
080826136	160	507.3	BPL	0.026 ± 0.005	0.99 ± 0.001	$2.02^{+0.89}_{-1.41}$	$4.86^{+2.82}_{-3.00}$
080826236	88	17.08	PL	0.99 ± 0.003	0.44 ± 0.02	$2.27^{+0.69}_{-0.57}$...
080828875	72	5.28	PL	0.93 ± 0.008	0.92 ± 0.008	$3.42^{+3.09}_{-1.51}$...
080903421	50	10.96	PL	0.96 ± 0.006	0.76 ± 0.01	$5.20^{+2.52}_{-2.81}$...
080903787	100	13.88	PL	0.06 ± 0.007	0.66 ± 0.02	$2.44^{+0.81}_{-0.61}$...

Notes. Burst lengths and fluences are taken from Lin et al. (2011). The posterior probability for the likelihood ratio is always for the simpler model tested (i.e., either power law or constant). α_1 is the power-law index in the simple power law, and the low-frequency power-law index in the broken power-law case. α_2 is the high-frequency power-law index in the broken power-law case. We quote the fifth and ninety-fifth percentiles for each quantity derived from a MCMC sample of 50,000 individual parameter sets.

with a Leahy normalization (noise powers averaging to 2; Leahy et al. 1983).

For none of the segments before and after each of the 27 bursts in our sample did we find significant detections of periodicities or QPOs. All segments present white-noise dominated periodograms consistent with a Poisson noise χ^2 distribution, indicating that the burst emission is not contaminated by a background source with significant timing behavior or instrumental effects on the relevant timescales. This includes any potential signal from the source itself. Any additional background source contributing emission would have to have switched on at the same time as the burst occurred, and switched off equally fast. This is highly unlikely. Note, however, that some instrumental effects, particularly dead time, scale with the source flux, and will not be recognizable in the background periodograms. Dead time in particular has an intricate effect on the burst periodogram, and led us to exclude the brightest bursts, which were also saturated.

5.2. An Example: Timing Analysis of Burst 080823478

In the following, we illustrate the analysis procedure with one specific burst, bn080823478, before giving results for the whole sample. This burst had a duration of $T_{90} = 264$ ms and the

highest fluence of the sample (see Table 1). The periodogram for this burst is presented in Figure 11.

5.2.1. Choosing a Noise Model

After fitting both a simple power law and the more complex broken power law, we computed the likelihood ratio between the two models, $\text{LRT} = 10.69$. Note that here, as well as in the analysis of the remaining sample, we set the smoothness parameter of the broken power law to -1 as in Vaughan (2010). The resulting function should more correctly be called a bending power law in this case, since it turns over in a smooth bend rather than a sharp break. Setting the smoothness parameter to -1 introduces a potential bias into the determination of the low-frequency power-law index for this model, however, including the smoothness parameter in the MCMCs, we found that the posterior distributions of this parameter for the bursts in our sample are very broad, indicating that the parameter is unconstrained. At the same time, it is correlated with the low-frequency power-law index, and thus the quoted values for this parameter should be read with caution. Additionally, we show below that the overall goodness of fit of the model to the data is good, indicating that another component is not needed. The fits to the periodogram and the residuals (data/fit) are presented




Figure 11. *Fermi* GBM observation of burst bn080823478 from SGR J0501 + 4516: periodogram and residuals I_j/S_j (for the light curve, see Figures 1 and 8). Upper panel: unsmoothed (black) and smoothed (orange; Wiener filter, $5\Delta\nu$) periodogram, power-law fit (blue) and broken power-law fit (red). Middle panel and lower panel show the residuals of the power-law fit and broken power-law fit, respectively. The broken power law presents a significantly better fit to the data.

(A color version of this figure is available in the online journal.)

in Figure 11. We use the Gaussian approximation to the covariance and the best-fit model parameters for the power-law model (H_0) as input to 500 MCMC ensemble walkers (see Section 2.2 or Foreman-Mackey et al. 2013 for details) with 100 samples each, after a burn-in phase with 100 samples for each walker. Figure 12 presents the posterior distributions of all five parameters and their correlations with each other. With 1000 randomly picked parameter sets from this sample of 50,000 parameter sets, we create 1000 fake periodograms, and compute the posterior predictive p -value for the LRT of the observed data using the formalism outlined in Section 2.2. In Figure 13 we plot a histogram of the posterior distribution for the likelihood ratio from the simulated periodograms. The black vertical line indicates the value of the likelihood ratio of the observed data. For bn080823478, $p(\text{LRT}) = 0.003 \pm 0.002$, hence we consider observing these data unlikely under the null hypothesis (a simple power law), and we adopt model H_1 for the rest of our analysis of this burst.

5.2.2. Searching for Periodicities

We use the broken power-law fit to the periodogram to draw another sample of MCMC parameter sets, in the same fashion as outlined above, and simulate 1000 fake periodograms all following a broken power law. From these, we computed the summed-square residuals T_{SSE} and search for the highest data/model outlier, T_R in the unbinned and binned periodogram. The latter should tell us about any features narrower than the frequency resolution $n\Delta\nu$ (where $n = 1$ for the unbinned periodogram and $n > 1$ for binned periodograms), while the former will give information about the overall fit of the model to the data. For this burst, we computed the posterior predictive distribution for the square-summed residuals and compared this distribution to the observed value, finding $p_{\text{SSE}} = 0.49 \pm 0.01$. As this statistic is an indicator for how well the model fits the data, we expect a low p_{SSE} to indicate that the model fit could

be improved, either by implementation of a different model or addition of model components. For this burst, we conclude the model fits the data rather well. The highest outlier in the residuals is at $\nu_{\text{max}} = 2317\text{Hz}$ with a power $2I/S = 15.71$ and a posterior predictive p -value $p(T_R) = 0.42 \pm 0.02$. The observed maximum power seen in the residuals is well within the distribution of outliers produced by the Monte Carlo simulations of the broadband model without any periodicity (see Figure 14), and is hence unlikely to represent a real periodic process. Similarly, we find maximum outliers as well as posterior probabilities of these outliers for the smoothed and binned periodograms, neither of which show no significant model parameters as well as the other bursts in Tables 1 and 2. None of the outliers were significant, thus we conclude that under our assumption of red noise, there are no narrow (quasi-)periodic signals in this data set. The posterior distributions and the observed values for the unsmoothed and smoothed periodograms are presented in Figure 14.

We searched the burst for broader quasi-periodic signals using an additional Lorentzian component and comparing the mixture model of broadband noise process and Lorentzian to the pure broadband model. With a posterior probability of the pure broadband model of $p(\text{LRT}) = 0.51 \pm 0.02$ (i.e., the probability that this model is sufficient in explaining the observed data), we conclude that there is no QPO in the burst.

5.3. Whole Sample

For all bursts, we followed the same procedure as for 080823478. All of the preferred models had a fairly high $p(\text{SSE})$, which indicates that the overall fit of the preferred model to the data is good. A summary of the results is presented in Table 1. Periodicity searches on the data are summarized in Table 2. While we compute posterior p -values for all binned and smoothed spectra, we only report the results for the unbinned spectra here for reasons of brevity, and only point out significant results in the binned spectra where appropriate.

None of the 27 bursts shows periodicities of any noteworthy significance in any of the unbinned (see Table 2, column $p(T_R)$) periodograms. The highest data/model outlier significance is seen in burst bn080823847a (see Figure 15 for a light curve and periodogram), $p(T_R) = 0.11 \pm 0.01$, at frequency $\nu_{\text{max}} = 4057\text{Hz}$ with a power $P(2I/S) = 18.88$, well below the power required to reach the detection threshold corresponding to a posterior p -value of 5%. However, the same burst shows significant signals in the binned periodograms, as summarized in Table 3. Note that p -values quoted there are corrected for the number of frequencies searched, but neither for the number of bursts searched nor the number of binned spectra searched for each burst. While the former is straightforward (a simple multiplication factor of 27 for the number of bursts searched), the latter is more complicated, owing to the fact that searching different binnings for a single periodogram does not result in independent trials. The most conservative assumption is to consider them independent, including another multiplication factor equal to the number of binnings searched (here: 9). This would rule out all but the two signals with frequency bins of 95 and 158 Hz, which remain significant even after a correction for the number of trials.

Comparing the results in Table 3 to the periodogram of the same burst in Figure 15 allows for several interesting observations, whose implications will be discussed in detail in Section 6. The frequencies at which significant excess powers




Figure 12. The MCMC sample of the parameters for the preferred broadband model for burst 080823478 (here: a broken power law). The posterior distributions for individual parameters are presented on the diagonal. If a posterior distribution is very broad, then the parameter is not very well constrained (indicating a high standard deviation on that parameter), and a simpler model might be adequate. The off-diagonal panels show correlations between parameters (panels opposite of each other, mirrored on the diagonal, are equivalent): scatter plots for 1000 randomly picked parameter pairs from the entire sample of 250,000 parameter sets, and contours of number density. One can observe for example a very tight correlation between low-frequency power-law index and normalization, and very little correlation between the normalization and the noise. Other parameters may correlate in more complex ways with each other. The trails and “clumpiness” in some of the scatter plots indicate that the distributions are not perfectly unimodal, and that even for highly peaked distributions, there are parameters far off the mean that are nevertheless not entirely unlikely.

(A color version of this figure is available in the online journal.)

are detected in the binned spectra are all at integer multiples of a suspicious feature at around 30 Hz, which in itself is not significant in any of the searches. The periodogram itself shows fairly prominent features at the frequencies at which signals are detected (see arrows in Figure 15, right panel), which become more prominent in the binned spectra, lending confidence that these might be real signals, and not false positive detections. If indeed there is a feature at 30 Hz, whose significance is diminished by the presence of red noise, then the higher-frequency detections may correspond to harmonics of this signal. The implications of these findings are discussed in more

detail in Section 6. Since we only search for the highest peak in each periodogram, there is a chance that several frequencies may be significant in each binned periodogram. This would require a more extensive search, including, e.g., the second- and third-highest peaks in the analysis. Additionally, a potential signal may have an energy dependence, thus an energy-resolved timing analysis may yield more conclusive results.

In none of the 27 bursts do we find any significant QPOs (see last column in Table 2). Posterior probabilities for the broadband model alone are largely in the range 0.2–0.8, indicating that the broadband model alone is an adequate fit to the data, and an

Table 2
Results for the Search for Periodicities and Quasi-periodicities in the Entire Sample of Bursts

Burst ID	Maximum Measured Power			Sensitivities (% rms)				$p(\text{LRT})$
	T_R	ν_{max}	$p(T_R)$	40 Hz	70 Hz	100 Hz	500 Hz	
080822529	1.43	1973	0.77 ± 0.01	19	19	19	19	0.47 ± 0.02
080822647	14.36	3283	0.56 ± 0.02	78	40	28	16	0.48 ± 0.02
080822981	7.79	2118	0.98 ± 0.004	72	16	0.76 ± 0.02
080823020	13.77	3145	0.31 ± 0.01	80	32	23	13	0.50 ± 0.02
080823091	0.44	2367	0.80 ± 0.01	9	9	9	9	0.51 ± 0.02
080823174	18.14	1711	0.24 ± 0.01	13	12	12	12	0.48 ± 0.02
080823248	12.49	95	0.91 ± 0.01	18	17	17	17	0.53 ± 0.02
080823293a	15.51	2069	0.32 ± 0.02	23	14	12	10	0.49 ± 0.02
080823293b	11.03	3593	0.54 ± 0.02	35	35	35	35	0.50 ± 0.02
080823319	14.39	1542	0.43 ± 0.02	28	18	16	14	0.54 ± 0.02
080823330	15.07	3695	0.46 ± 0.02	40	19	12	6	0.88 ± 0.01
080823354	12.13	2407	0.72 ± 0.01	46	26	21	18	0.81 ± 0.02
080823429	12.97	3689	0.55 ± 0.02	81	24	17	13	0.49 ± 0.02
080823478	15.71	2317	0.42 ± 0.02	28	12	8	4	0.51 ± 0.02
080823623	18.60	902	0.09 ± 0.009	24	17	16	14	0.51 ± 0.02
080823714	15.03	1301	0.69 ± 0.01	19	14	13	11	0.85 ± 0.01
080823847a	18.88	4057	0.11 ± 0.01	43	23	15	8	0.49 ± 0.02
080823847b	11.03	2515	0.94 ± 0.007	...	57	40	15	0.50 ± 0.02
080823986	10.07	2791	0.74 ± 0.01	19	19	19	19	0.47 ± 0.02
080824346	9.39	2968	0.83 ± 0.01	...	70	55	26	0.52 ± 0.02
080824828	2.54	74	0.66 ± 0.01	23	23	23	23	0.24 ± 0.02
080825401	12.36	496	0.80 ± 0.01	73	36	26	7	0.94 ± 0.007
080826136	12.80	2868	0.77 ± 0.01	43	20	11	5	0.54 ± 0.02
080826236	13.16	3536	0.49 ± 0.02	70	38	28	15	0.56 ± 0.02
080828875	12.24	667	0.56 ± 0.02	54	25	17	17	0.53 ± 0.02
080903421	9.97	3781	0.66 ± 0.02	...	33	24	22	0.56 ± 0.02
080903787	14.04	2817	0.40 ± 0.02	73	42	28	16	0.50 ± 0.02

Notes. We show the $T_R = \max_j(\hat{R}_j)$ statistics for each burst, along with the associated frequency and the posterior probability to find this outlier given a pure noise process. For each burst, we also quote sensitivities, i.e., the fractional rms amplitude a periodic process would have needed to have in order to be detectable for our method, given the noise process and parameters determined for that burst. Note that due to the excess power in the low-frequency part spectrum being modeled as red noise, the sensitivity will depend on frequency, and be generally less constrained in the low-frequency part of the spectrum than in the white-noise dominated high-frequency spectrum. Where no sensitivity is given, the derived value exceeded 100%. A signal with more than 100% fractional rms amplitude would have negative counts, and is therefore not physical. A sensitivity limit on the amplitude $>100\%$ merely indicates that we cannot constrain the signal amplitude at the given frequency at all. Finally, we also present the posterior probability on the likelihood ratio for a model containing a QPO versus a model without QPO, which is an indicator for the presence of a QPO in the spectrum.




Figure 13. Distribution of likelihood ratios for 1000 simulations of the null hypothesis (power-law model). The observed value of T_{LRT} for burst 080823478 is indicated as a black vertical line. The further to the right (i.e., the further in the tail of the distribution) this observed value is located, the more unlikely the null hypothesis becomes, indicating that a more complex model (in this case the broken power law) may be more appropriate in modeling the broadband variability.

(A color version of this figure is available in the online journal.)

additional Lorentzian does not result in a better fit. However, we also note that the posterior probability of the likelihood ratio is clustered around 0.5 for 19 out of 27 bursts. Since for a well-behaved probability statistic applied to data consistent with the null hypothesis, p -values should be uniformly distributed between 0 and 1, this clustering indicates that the test is conservative in the sense that it does not overstate the rejection of either null hypothesis or alternative hypothesis. In practice, results on the LRT should be combined with those on the binned spectra to yield reliable detections.

5.4. Broadband Variability

The broadband variability observed in the bursts is not just a nuisance when searching for (quasi-) periodicities, but is of interest in its own right: it shows that something is varying in the source, although not periodically. Here, we have chosen a purely phenomenological approach, selecting empirical models that are both simple and widely observed in many astrophysical contexts, without physical justification. Hence, the question of whether we can derive any physical knowledge from these empirical models is an interesting and important one.

Almost all bursts in the sample are well-modeled by a simple power law, although the lower and upper bounds on the 90% credible interval show a large variation in some




Figure 14. Histograms of posterior distributions and observed values (black vertical lines, burst 080823478) of the T_R statistic defined in Equation (11). (A color version of this figure is available in the online journal.)

indices, indicating that they are not very well constrained. Some caution should be used when interpreting these values, since they were derived from a sample of bursts with very diverse properties overall, such as burst length and fluence. Additionally, an unmodeled burst envelope may significantly change the slope of the power law-like part of the periodogram. A reliable characterization of the broadband properties hinges on our knowledge of the processes (both noise and non-stochastic) involved, and will be deferred to a future paper involving a larger sample of bursts. Figure 16 shows the distribution of power-law indices for all bursts where at least a simple power law was required to adequately represent the data. The distribution of indices ranges from 1.7 to 4.3 and peaks around 2.5, which is higher than commonly seen for example in GRBs (see, e.g., Beloborodov et al. 2000 and Guidorzi et al. 2012).

Four bursts could be modeled without invoking the presence of red noise at all; contributions by burst variability were confined to very low frequencies and standard Fourier techniques apply for all but the first three or four frequency bins. Only two bursts required the more complex model (bursts 080823478 and 080826236). While these were not the longest bursts, they had

Table 3
Posterior Summary for Various Binned Periodograms Derived from Burst 080823847a

$d\nu_{\text{bin}}$ (Hz)	ν_{max} (Hz)	T_R	$p(T_R)$
16	310	7.24	0.017 ± 0.004
22	2094	7.01	0.004 ± 0.002
32	301	6.01	0.003 ± 0.002
47	307	4.90	0.002 ± 0.001
63	4067	4.27	0.003 ± 0.002
95	2096	4.10	$(< 2.0 \times 10^{-5})$
158	2129	3.26	$(< 2.0 \times 10^{-5})$
316	2050	2.59	0.027 ± 0.005
633	2050	2.45	0.003 ± 0.002
949	2050	2.34	0.007 ± 0.003
1583	2050	2.16	0.029 ± 0.005

Notes. p -values were derived using an increased number of 50,000 simulations in order to increase the resolution on small probabilities. The first column holds the binned frequency resolution, $d\nu_{\text{bin}}$, the second column the frequency at which the highest outlier T_R was found, ν_{max} , column three the corresponding value of the T_R statistic and finally column four the associated posterior p -value to find that value in a pure noise spectrum, binned to the same frequency resolution. Note the probabilities without errors in brackets at binning frequencies of 95 and 158 Hz. The p -value there turned out to be zero at these binning frequencies. Of course, the p -value is not actually zero; however, since we approximate the posterior distribution of T_R with a finite number of simulations, there is a possibility that the true probability to achieve the observed value with only noise is small enough that none of the simulations will exceed the observed T_R , giving rise to a zero p -value. We computed the p -values with up to 50,000 simulations, and hence state an upper limit on the p -value of 2.0×10^{-5} .

the highest fluence (except for the excluded saturated bursts), indicating a potential correlation between power spectral shape and burst fluence. This is expected: the normalization (i.e., the relative strength to the noise) of the broadband noise model depends directly on the number of counts detected, thus bright bursts may enable us to see the cut-off frequency of the power law as set by the burst duration, whereas many of the other bursts have too low a fluence to observe the same behavior. Alternatively, it is possible that the difference in power spectral shape is intrinsic to the source, that is, the physical processes creating this kind of variability vary in some way with burst fluence. Without a model for the emission processes producing the burst




Figure 15. *Fermi* GBM observation of burst bn080823847a from SGR J0501+4516. Left: light curve with a time resolution of 0.002 s. Structure in the burst profile is clearly visible. Right: unbinned (blue) and binned (magenta: 16 Hz binning; orange: 65 Hz binning) periodogram for this burst. There is a feature in the periodogram around 30 Hz (leftmost arrow), which is by itself not significant. However, significant features reported in Table 3 are all at integer multiples of this frequency (within the uncertainty imposed by the frequency resolution), indicating the presence of harmonics at 150 Hz, 300 Hz, 900 Hz, and 2100 Hz (arrows 2–5).

(A color version of this figure is available in the online journal.)




Figure 16. Distribution of power-law indices (α for single power law, α_2 for broken power law) for all bursts where the power law or broken power law was preferred.

(A color version of this figure is available in the online journal.)

in the first place, however, it is difficult to assess the validity of the latter hypothesis. Additionally, with only two bursts preferring the broken power-law model, the numbers are too low to draw conclusions, and we defer the discussion on the physical implications of the broadband noise modeling to a later paper utilizing a larger sample of magnetar bursts.

6. DISCUSSION AND CONCLUSION

Magnetar bursts are a potential window into the interior of neutron stars, via the oscillations measured in magnetar giant flares. Finding analogous signals in the wealth of short SGR bursts, however, poses something of a challenge. We have shown that timing analysis of astrophysical transients is a non-trivial problem. Standard Fourier techniques are defined for infinitely long time series, an assumption that is clearly broken by the non-stationary nature of transient events in general, and magnetar bursts in particular.

Monte Carlo simulations of light curves fail to be predictive when there is no precise knowledge of the underlying burst light curve: there is a degeneracy between the overall, aperiodic burst shape, a potential red noise component, and the very thing we would like to measure: a QPO. When the light curve is not adequately modeled, then the periodograms produced from the Monte Carlo simulations will not reproduce the low-frequency part of the periodogram, where it clearly diverges from the statistical distributions expected for pure white noise.

In the absence of better knowledge about the emission processes in magnetar bursts, we advocate a conservative Bayesian method that models the burst light curve as a pure red noise process. It is purely empirical in the sense that it does not require additional assumptions on the underlying physical processes, except for fairly broad assumptions on what the power spectrum shape might be. Assuming pure red noise is, in a way, the complementary extreme to Monte Carlo simulations of the light curve: in one, we assume only a deterministic burst profile without the presence of red noise, with the price that inadequate modeling of the periodogram shape will lead to spurious detections. Here, we assume only red noise, at the

cost that weak signals are likely missed. This is the greatest weakness of our approach. We have shown in Section 4 that even strong signals may be undetectable at low frequencies, where burst envelope and red noise dominate. These, however, are exactly the frequencies at which many of the QPOs in giant flares have been seen (e.g., 18 Hz, 30 Hz for the 2004 giant flare; see Israel et al. 2005). This limitation is in part not only due to restrictions of our method, but also to the short lengths of the SGR bursts, where at these frequencies only one or two cycles may be seen in the light curve. Upper limits below 100 Hz are often fairly unconstraining, and range from 10% to more than 100% fractional rms amplitude for a signal to be detectable. At frequencies close to and above 100 Hz, sensitivities approach the white noise limit, which is strongly dependent on the number of photons from a particular burst. Thus, for a bright burst with good count statistics, sensitivities are quite constraining, down to less than 10% (e.g., burst 080823478, see Table 2). This is comparable to what was observed in giant flares: for example, a QPO at 93 Hz, as seen in the 2004 flare, at roughly 10% rms amplitude (Israel et al. 2005; Watts & Strohmayer 2006), should be detectable in at least the brightest bursts of our sample. Similarly, a high-frequency QPO like the one at 625 Hz seen in the 2004 flare with a fractional rms amplitude of up to 20% should be clearly detected with our method as well.

However, QPOs in SGR bursts may be less strong than in the giant flares, owing to the lower energy injected in SGR bursts, and hence more likely to be misclassified as non-detections, if their fractional rms amplitudes fall below 5%. Additionally, we restrict ourselves when searching for periodicities by considering only the highest peak in the spectrum, which is clearly not adequate when there are multiple signals in the periodogram. On the other hand, if even the highest peak is not significant, any other peak in the periodogram will be even less significant.

The burst 080823847a presents an interesting case that illustrates the limits of a pure signal-processing approach to the timing analysis shown here. Two of several significant signals detected in the binned spectra remain significant even after the most conservative correction for the number of trials is applied, indicating that there is indeed a real signal present. However, the nature of this signal is at present unclear. The detected signals are possibly harmonics of a lower-frequency signal around 30 Hz, corresponding to a timescale of $\tau = 1/\nu = 33$ ms. This timescale roughly corresponds to the two sharp peaks seen in the burst light curve in Figure 15 (left side). Whether we consider this to be a QPO atop a burst envelope or not cannot be answered from Fourier analysis alone; it becomes a matter of interpretation and prior knowledge. The presence of harmonics indicates a signal repeating on the timescale of the fundamental, but with variability on shorter timescales in the signal. One could well interpret the two peaks in the light curve as a strongly damped (quasi-)periodic signal that is amplified together with some underlying burst profile and dies away after two, possibly three, cycles. The frequency of this signal is similar to that observed from the 2004 giant flare (Israel et al. 2005; Strohmayer & Watts 2005), and thus not unlikely. On the other hand, this kind of signal can equally well be derived from a red noise process. The fact that red noise is a stochastic process means that at some point, two or even three peaks will follow each other, as in the present burst. While red noise itself would not introduce harmonics, the signal could be boosted by an underlying burst envelope, introducing the observed harmonics. At present, without any knowledge about emission processes and the kind of light curve they produce, it is impossible to

distinguish the two, and we choose the conservative approach and interpret the observed feature as part of a noise process.

Another problem as yet unsolved is that of false non-detections we expect, i.e., weak signals missed due to the fact that we assume a pure red noise process. There are several ways to break this dilemma, but all require more detailed knowledge of the variability-producing processes in the neutron star, and this is where both theoretical efforts and development of novel statistical techniques are required. What produces the burst emission? What produces the aperiodic variability seen in the red noise part of the periodograms? Until we can answer these questions, finding QPOs in magnetar bursts will always suffer from the essential degeneracy between burst envelope and red noise. If we knew the overall burst shape, one could, for example, simulate light curves, but as a combination of a burst profile and a red noise process, as done in Section 4, and compare this sample to the observed periodogram. Other approaches involve leaving behind the Fourier domain and its incorrect assumption of stationarity behind altogether.

Knowledge about the burst envelope, on the other hand, would also offer us an additional source of observational information to exploit: if we can use the existing information on the hundreds of bursts available to learn something about the burst envelope shape, we may be able to put tight constraints on potential QPO detections and provide additional observational constraints for burst emission models in general. Clearly, with the right statistical techniques, there is a wealth of information yet to be extracted from the SGR bursts observed with *Fermi* GBM. Additionally, for bursts with high count rates, it is possible to study variability properties of the bursts with energy, thanks to *Fermi*/GBM's excellent energy resolution. These studies may provide additional information on QPOs that depend on energy. The methods developed here, however, while developed with SGR bursts in mind, are by no means limited to magnetars. They are applicable in fairly general circumstances, for any light curve that is phenomenologically similar to what we observe from magnetars: highly variable, transient events with complex light curves. This includes, for example, other known transients such as GRBs, tidal disruption events, and supernova light curves.

D.H. and A.L.W. acknowledge support from a Netherlands Organization for Scientific Research (NWO) Vidi Fellowship (PI: A. Watts), and would like to thank Jason Farquhar for useful discussions. C.K. was partially supported by NASA grant NNH07ZDA001-GLAST. E.G. acknowledges support from the Scientific and Technological Research Council of Turkey (TÜBİTAK) through grant 109T755.

APPENDIX A

COMPARISON WITH RECENT RESULTS

El-Mezeini & Ibrahim (2010) searched a data set of SGR bursts from SGR J1806–20, found in data taken with the *RXTE*. They report the significant detection of QPOs in 5 different SGR bursts out of a sample of 30, with frequencies of 84, 103, and 648 Hz with significances $\geq 4.3\sigma$, estimated using Monte Carlo simulations of light curves equivalent to those described in Section 2.1, by smoothing the burst light curve and subsequently adding Poisson detector noise to form a null hypothesis against which to test the data. For the reasons stated in Section 2.1, we do not believe these estimates to be conservative, and consequently reanalyze this data set from 1996 November 5–18 after barycentering the data and filtering out photons outside

the range 2–60 keV, where the response curve of the instrument indicates that noise will dominate at these energies. We use data from all Proportional Counter Units of the Proportional Counter Array detector. We use the Bayesian analysis presented in Section 2.2 to choose a broadband noise model and search for (quasi-)periodic signals in the same data set. The results are presented in Table 4 and Figure 17.

For the most part, we cannot confirm the detections shown in El-Mezeini & Ibrahim (2010) using our Bayesian methods. Only one burst shows a marginally significant signal: burst 3 ($p = 0.03 \pm 0.01$) at a frequency of around 3706 Hz, and only in one binned spectrum. Given the posterior probability is very close to our (rather high) detection threshold, we are inclined to disregard this detection as insignificant as well, as it would indeed become insignificant as soon as we take the number of bursts searched into account. This result is in stark contrast with the probabilities quoted in El-Mezeini & Ibrahim (2010). There are, however, errors in their analysis: taking into account the varying nature of the background light curve should make the significance of any claimed detection drop compared with the significance computed from the ideal χ^2 distribution. Although El-Mezeini & Ibrahim (2010) do carry out simulations, the significances that they quote rise substantially, indicating a problem in their Monte Carlo simulation method. Indeed their simulated power spectra show far fewer high noise powers than one would expect given the number of simulations carried out and the number of independent frequency bins (enhancing the significance of any tentative detection).

APPENDIX B

DATA ANALYSIS RECIPES

B.1. How to Fit a Noise Model

Fitting a noise model is essentially a model selection task. In the following, we will lay out the individual steps in a recipe-like style.

1. Compute a periodogram of the burst light curve.
2. Fit the periodogram with both a simple model (the null hypothesis) and a more complex model (the alternative hypothesis we wish to test against), to get MAP estimates for the parameters in each model
3. Using the MAP estimates, compute the likelihoods of the data given each model and MAP parameters, then compute the likelihood ratio of the complex model versus the simple (null) model.
4. Produce a large MCMC sample approximating the posterior distributions of the parameters of the null model.
5. Pick a (large) number n of parameter vectors from this sample (e.g., $n = 1000$), and create power spectra from these parameters and simulate a periodogram from each by drawing a realization from the random process the power spectrum represents. This will yield n fake periodograms.
6. Fit each simulated periodogram with both simple and complex model, exactly in the same way as done for the burst periodogram, and compute the likelihood ratio of this simulation. The sample of likelihood ratios from fitting the simulated periodograms will be representative of the likelihood ratios one expects when fitting both the simple and complex model to a sample of periodograms derived entirely from the null hypothesis.
7. Compare the distribution of simulated likelihood ratios with that of the observed burst, and compute the tail area

Table 4
Summary of the Bayesian Analysis for Five Bursts Presented in El-Mezeini & Ibrahim (2010)

Quantity	Burst 1	Burst 2	Burst 3	Burst 4	Burst 5
Start time (MET s)	90915519.65	90909708.72	90925017.31	90915519.65	90937076.56
length (s)	0.13	0.31	0.12	0.11	0.36
Model	PL	PL	PL	PL	PL
$p(\text{LRT})$	0.84 ± 0.01	0.59 ± 0.01	0.13 ± 0.01	0.41 ± 0.01	0.38 ± 0.02
$p(\text{SSE})$	0.91 ± 0.008	0.85 ± 0.01	0.39 ± 0.02	0.26 ± 0.01	0.75 ± 0.01
α_1	2.63	2.68	2.52	2.67	2.29
α_1 5%	2.07	2.31	2.06	2.10	2.11
α_1 95%	3.31	3.12	3.04	3.37	2.47
Unsmoothed					
$\max(2I/S)$	15.11	13.99	13.67	13.18	14.46
v_{\max}	2464	2690	3696	3694	3246
$p(T_R)$	0.13 ± 0.01	0.41 ± 0.02	0.08 ± 0.01	0.32 ± 0.01	0.95 ± 0.02
Three bins					
$\max(2I/S)$	8.34	9.21	8.65	9.4	9.24
v_{\max}	2466	2701	3706	3698	1394
$p(T_R)$	0.20 ± 0.01	0.05 ± 0.01	0.03 ± 0.01	0.11 ± 0.01	0.18 ± 0.01
Five bins					
$\max(2I/S)$	7.38	8.51	8.73	6.33	7.15
v_{\max}	2352	2703	3695	3687	1395
$p(T_R)$	0.11 ± 0.01	0.05 ± 0.01	0.11 ± 0.01	0.24 ± 0.01	0.20 ± 0.01
Eleven bins					
$\max(2I/S)$	4.27	6.91	5.05	4.98	5.85
v_{\max}	2416	2708	3712	3648	2218
$p(T_R)$	0.07 ± 0.01	0.41 ± 0.01	0.23 ± 0.01	0.29 ± 0.02	0.53 ± 0.01
QPO $p(\text{LRT})$	0.49 ± 0.02	0.20 ± 0.02	0.47 ± 0.02	0.54 ± 0.02	0.25 ± 0.01

Notes. The choice of broadband model is recorded, as well as the results of the periodicity and QPO searches. Note that the posterior probability for the summed-squared residuals ($p(\text{SSE})$) is shown for the model with the better fit, whereas the posterior probability for the likelihood ratio $p(\text{LRT})$ is shown for the simpler model always.

probability of seeing the observed likelihood ratio, if the data were entirely drawn from the null hypothesis. If this tail area probability is large, then the data are consistent with the null hypothesis. The converse, however, is not necessarily true. A small p -value indicates that the data are unlikely to be drawn from the null hypothesis. This is not a direct proof that the complex model is the underlying process that produced the observed burst, however, it is an indication that the more complex model is likely a better representation of the data than the null hypothesis.

B.2. Searching for Periodicities

1. Fit a broadband noise model (e.g., the preferred model as defined as above in Appendix B.1) to the burst, and compute the residuals $R_j = 2I_j/S_j$, where I_j are the periodogram powers and S_j is the broadband noise model at j th frequency ν_j .
2. From the residuals, pick the highest outlier $\max R_j$; this is the candidate single-bin periodicity.
3. Simulate a large number of periodograms from an MCMC sample in the same way as done for the choice of noise model above.
4. Fit each simulated periodogram with the preferred noise model, compute the data/model residuals R_j and find the maximum outlier $\max R_j$ in the residuals of each simulated periodogram.
5. Compare the distribution of maximum outliers from the set of simulations derived from the broadband noise model with no periodicity present, to the outlier in the real burst periodogram. One may compute the posterior predictive p -value for the observed outlier in much the same way as for the LRT. If the p -value is large, then the outlier is consistent with a pure noise distribution.

B.3. Searching for QPOs

We search for QPOs as a model selection problem, where we compare the broadband noise model to a more complex model combining both the broadband noise model and a Lorentzian to account for the QPO. Because we do not know the centroid frequency of the potential QPO inherently, the task is slightly more complex than that which we use for the choice of noise model.

1. Fit the observed periodogram with the broadband noise model and compute residuals R_j . Smooth the residuals with a Wiener filter with a width of five frequency bins in order to reduce the probability of the minimization algorithm terminating in local minima due to sharp noise features.
2. At each frequency, fit a flat line and a Lorentzian of variable width and intensity, but fixed centroid frequency to the smoothed residuals and compute the likelihood of that fit. We leave out the first five and last five bins in order to avoid fitting the edge of the periodogram. The result of this process will be the maximum likelihood as a function of frequency.
3. Pick the frequency bin with the largest maximum likelihood as given from modeling each frequency.
4. Fit the full-resolution periodogram with the broadband noise model alone as well as a combined model of broadband red noise and Lorentzian, using the model parameters for the largest maximum likelihood fit in the previous step as starting parameters for this fit.
5. Compute the likelihood ratio for these two models.
6. Simulate a large number of periodograms (in our case, 500 to reduce computational load) from an MCMC sample of the broadband noise model.

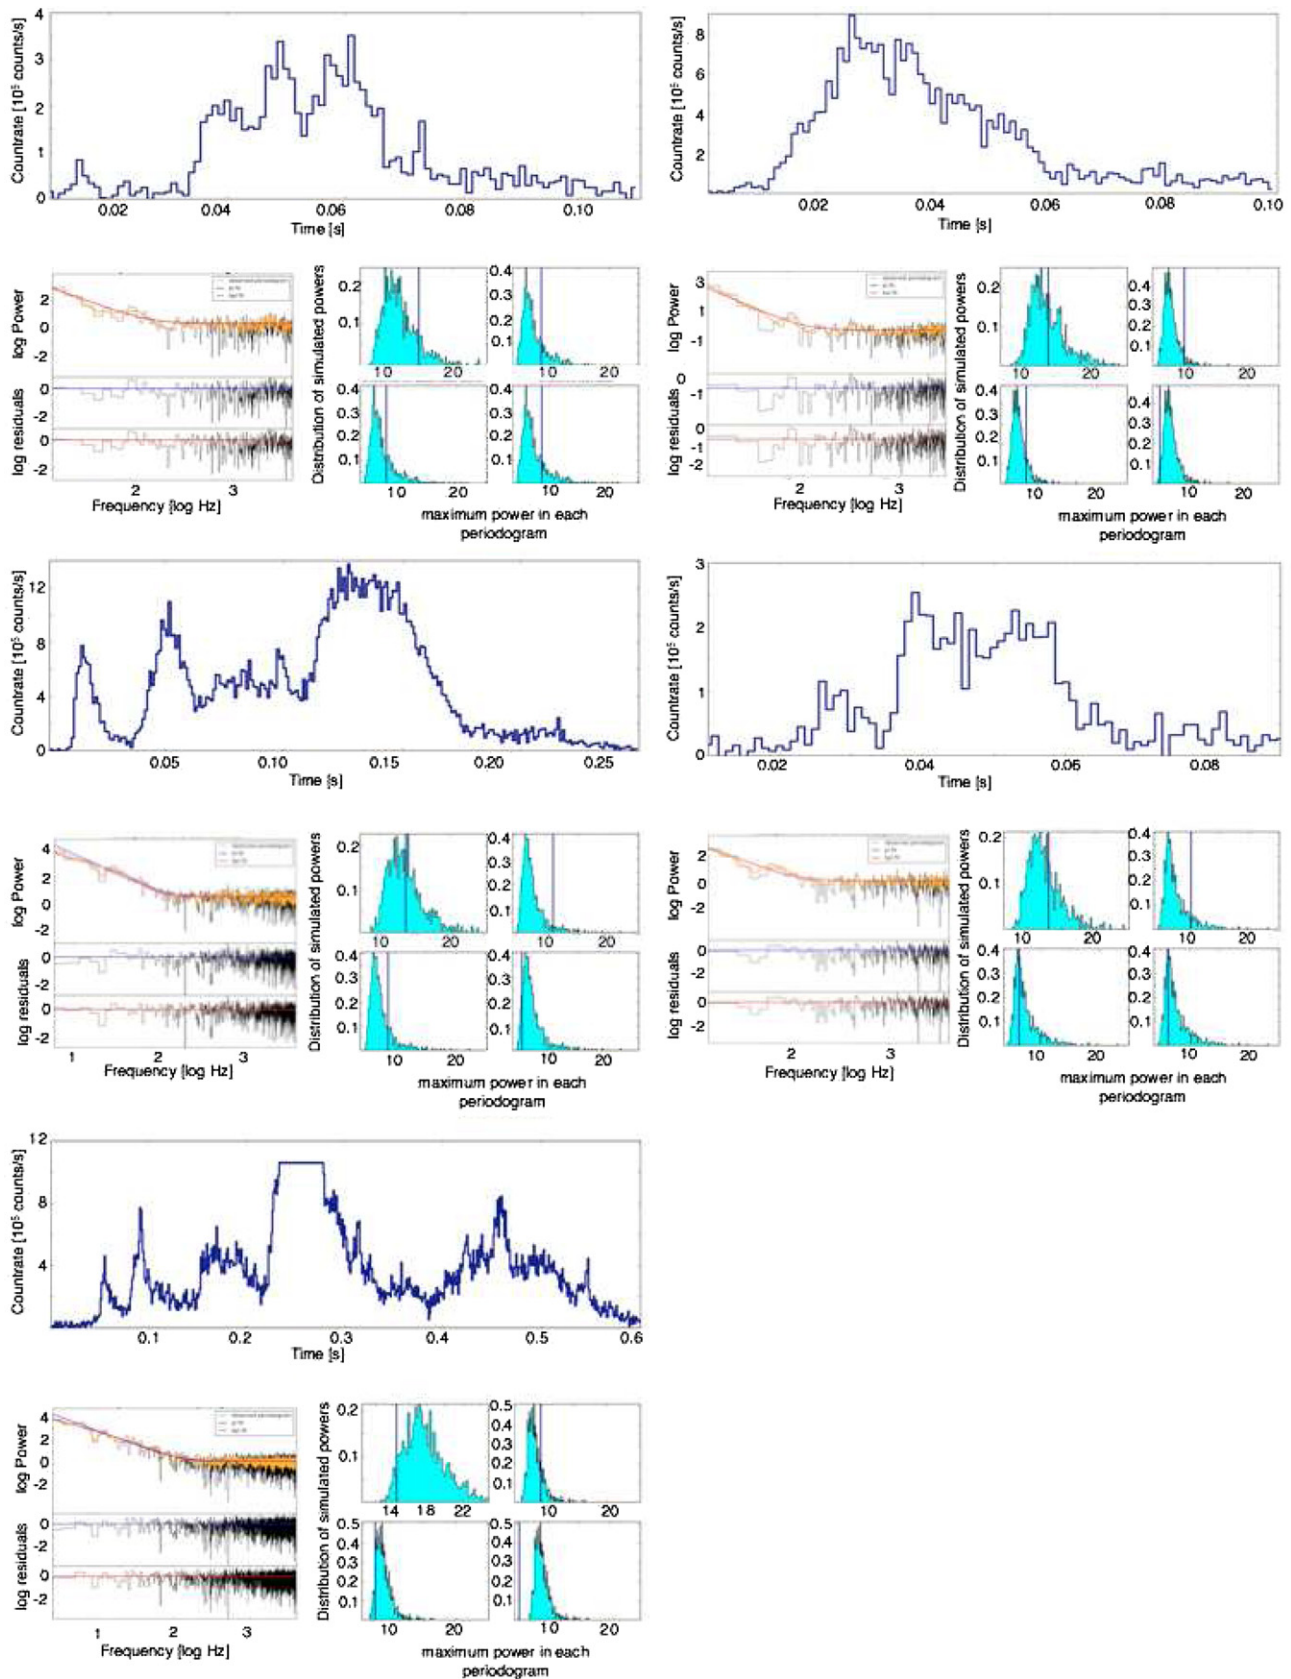


Figure 17. Light curves, periodograms, and posterior distributions for the five bursts from SGR J1806 – 20 observed with the *Rossi X-Ray Timing Explorer* between 1996 November 5–18 presented in El-Mezzeini & Ibrahim (2010). Long, upper plots for each burst are the light curves binned to 0.005 s. On the lower left, the unsmoothed periodograms (black) the five-bin smoothed periodogram (orange) as well as the power law and broken power-law fits. Underneath each periodogram is a plot of the residuals of dividing the periodogram by the broadband model (power law in the middle, and bent-power law on bottom). On the right, the posterior distributions of the highest data/model outlier for the unsmoothed (upper left) and smoothed periodograms (rest; upper right: 3-bin smoothing, lower left: 5-bin smoothing; lower right: 11-bin smoothing). The observed value is overplotted as a black vertical line.

(A color version of this figure is available in the online journal.)

7. For each simulated periodogram, follow the exact same procedure in steps 1–4 to produce an approximation to the distribution of likelihood ratios from the null hypothesis.
8. Compare the distribution of likelihood ratios derived from the simulated periodograms to the likelihood ratio for the observed burst, and compute a posterior predictive p -value for the probability of obtaining the observed likelihood ratio, if the data were consistent with pure red noise.

REFERENCES

- Anderson, E. R., Duvall, T. L., Jr., & Jefferies, S. M. 1990, *ApJ*, **364**, 699
- Andersson, N., Glampedakis, K., & Samuelsson, L. 2009, *MNRAS*, **396**, 894
- Barret, D., & Vaughan, S. 2012, *ApJ*, **746**, 131
- Barthelmy, S. D., Baumgartner, W. H., Beardmore, A. P., et al. 2008, *ATel*, **1676**, 1
- Beloborodov, A. M., Stern, B. E., & Svensson, R. 2000, *ApJ*, **535**, 158
- Colaiuda, A., & Kokkotas, K. D. 2012, *MNRAS*, **423**, 811
- D’Angelo, C. R., & Watts, A. L. 2012, *ApJL*, **751**, L41
- Duncan, R. C. 1998, *ApJL*, **498**, L45
- Duncan, R. C. 2004, in *Cosmic Explosions in Three Dimensions*, ed. P. Höflich, P. Kumar, & J. C. Wheeler (Cambridge: Cambridge Univ. Press), 285
- El-Mezzeini, A. M., & Ibrahim, A. I. 2010, *ApJL*, **721**, L121
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, in press (arXiv:1202.3665)
- Fox, D. W., Lewin, W. H. G., Rutledge, R. E., et al. 2001, *MNRAS*, **321**, 776
- Gabler, M., Cerdá Durán, P., Font, J. A., Müller, E., & Stergioulas, N. 2011, *MNRAS*, **410**, L37
- Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. S. 2004, *Bayesian Data Analysis* (London: Chapman & Hall)
- Gill, R., & Heyl, J. S. 2010, *MNRAS*, **407**, 1926
- Gögüş, E., Woods, P., & Kouveliotou, C. 2008, *ATel*, **1677**, 1
- Gögüş, E., Woods, P. M., Kouveliotou, C., et al. 1999, *ApJL*, **526**, L93
- Gögüş, E., Woods, P. M., Kouveliotou, C., et al. 2000, *ApJL*, **532**, L121
- Gögüş, E., Woods, P. M., Kouveliotou, C., et al. 2010, *ApJ*, **722**, 899
- Groth, E. J. 1975, *ApJS*, **29**, 285
- Guidorzi, C. 2011, *MNRAS*, **415**, 3561
- Guidorzi, C., Frontera, F., Montanari, E., et al. 2004, *A&A*, **416**, 297
- Guidorzi, C., Margutti, R., Amati, L., et al. 2012, *MNRAS*, **422**, 1785
- Hastings, W. K. 1970, *Biometrika*, **57**, 97
- Holland, S. T., Barthelmy, S. D., Baumgartner, W. H., et al. 2008, *GCN*, **8112**, 1
- Hurley, K., Cline, T., Mazets, E., et al. 1999, *Natur*, **397**, 41
- Ibrahim, A. I., Strohmayer, T. E., Woods, P. M., et al. 2001, *ApJ*, **558**, 237
- Israel, G. L., Belloni, T., Stella, L., et al. 2005, *ApJL*, **628**, L53
- Israel, G. L., Romano, P., Mangano, V., et al. 2008, *ApJ*, **685**, 1114
- Kaneko, Y., Gögüş, E., Kouveliotou, C., et al. 2010, *ApJ*, **710**, 1335
- Kashyap, V. L., van Dyk, D. A., Connors, A., et al. 2010, *ApJ*, **719**, 900
- Kouveliotou, C., Dieters, S., Strohmayer, T., et al. 1998, *Natur*, **393**, 235
- Kouveliotou, C., Strohmayer, T., Hurley, K., et al. 1999, *ApJL*, **510**, L115
- Lattimer, J. M., & Prakash, M. 2007, *PhR*, **442**, 109
- Leahy, D. A., Darbro, W., Elsner, R. F., et al. 1983, *ApJ*, **266**, 160
- Lenters, G. T., Woods, P. M., Goupell, J. E., et al. 2003, *ApJ*, **587**, 761
- Levin, Y. 2007, *MNRAS*, **377**, 159
- Levin, Y., & Lyutikov, M. 2012, *MNRAS*, **427**, 1574
- Lin, L., Kouveliotou, C., Baring, M. G., et al. 2011, *ApJ*, **739**, 87
- Lyutikov, M. 2003, *MNRAS*, **346**, 540
- McDermott, P. N., van Horn, H. M., & Hansen, C. J. 1988, *ApJ*, **325**, 725
- Meegan, C., Lichti, G., Bhat, P. N., et al. 2009, *ApJ*, **702**, 791
- Mereghetti, S. 2008, *A&ARv*, **15**, 225
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. 1953, *JChPh*, **21**, 1087
- Palmer, D. M., Barthelmy, S., Gehrels, N., et al. 2005, *Natur*, **434**, 1107
- Park, T., van Dyk, D. A., & Siemiginowska, A. 2008, *ApJ*, **688**, 807
- Peng, Z. Y., Yin, Y., Bi, X. W., et al. 2010, *ApJ*, **718**, 894
- Perna, R., & Pons, J. A. 2011, *ApJL*, **727**, L51
- Protassov, R., van Dyk, D. A., Connors, A., Kashyap, V. L., & Siemiginowska, A. 2002, *ApJ*, **571**, 545
- Rea, N., Israel, G. L., Turolla, R., et al. 2009, *MNRAS*, **396**, 2419
- Steiner, A. W., & Watts, A. L. 2009, *PhRvL*, **103**, 181101
- Strohmayer, T. E., & Watts, A. L. 2005, *ApJL*, **632**, L111
- Strohmayer, T. E., & Watts, A. L. 2006, *ApJ*, **653**, 593
- Thompson, C., & Duncan, R. C. 1995, *MNRAS*, **275**, 255
- Timmer, J., & Koenig, M. 1995, *A&A*, **300**, 707
- Timokhin, A. N., Eichler, D., & Lyubarsky, Y. 2008, *ApJ*, **680**, 1398
- Uttley, P., McHardy, I. M., & Vaughan, S. 2005, *MNRAS*, **359**, 345
- van der Klis, M. 1989, in *NATO ASI Ser. C 262, Timing Neutron Stars*, ed. H. Ögelman & E. P. J. van den Heuvel (Dordrecht: Kluwer), 27
- van der Klis, M. 2006, in *Compact Stellar X-Ray Sources*, ed. W. H. G. Lewin & M. van der Klis (Cambridge: Cambridge Univ. Press), 39
- van Hoven, M., & Levin, Y. 2008, *MNRAS*, **391**, 283
- van Hoven, M., & Levin, Y. 2011, *MNRAS*, **410**, 1036
- Vaughan, S. 2010, *MNRAS*, **402**, 307
- Watts, A. L. 2012, in *Neutron Star Crust*, ed. C. A. Bertulani & J. Piekarewicz (Hauppauge, NY: Nova Science Pub. Inc.), 265
- Watts, A. L., & Reddy, S. 2007, *MNRAS*, **379**, L63
- Watts, A. L., & Strohmayer, T. E. 2006, *ApJL*, **637**, L117
- Watts, A. L., Strohmayer, T. E., & Markwardt, C. B. 2005, *ApJ*, **634**, 547
- Woods, P. M., Gögüş, E., & Kouveliotou, C. 2008, *ATel*, 1691, 1
- Woods, P. M., Kouveliotou, C., Gavriil, F. P., et al. 2005, *ApJ*, **629**, 985
- Woods, P. M., Kouveliotou, C., Gögüş, E., et al. 2001, *ApJ*, **552**, 748
- Woods, P. M., & Thompson, C. 2006, in *Compact Stellar X-Ray Sources*, ed. W. H. G. Lewin & M. van der Klis (Cambridge: Cambridge Univ. Press), 547