



Evidence of a Strong 19.5 Hz Flux Oscillation in Swift BAT and Fermi GBM Gamma-Ray Data from GRB 211211A

Cecilia Chirenti^{1,2,3,4}, Simone Dichiara⁵, Amy Lien⁶, and M. Coleman Miller⁷¹Department of Astronomy, University of Maryland, College Park, MD 20742, USA; chirenti@umd.edu²Astroparticle Physics Laboratory, NASA/GSFC, Greenbelt, MD 20771, USA³Center for Research and Exploration in Space Science and Technology, NASA/GSFC, Greenbelt, MD 20771, USA⁴Center for Mathematics, Computation, and Cognition, UFABC, Santo Andre, 09210-170, SP, Brazil⁵Department of Astronomy and Astrophysics, The Pennsylvania State University, 525 Davey Lab, University Park, PA 16802, USA⁶Department of Chemistry, Biochemistry, and Physics, University of Tampa, 401 W. Kennedy Boulevard, Tampa, FL 33606, USA⁷Department of Astronomy and Joint Space-Science Institute, University of Maryland, College Park, MD 20742-2421 USA

Received 2023 October 19; revised 2024 April 1; accepted 2024 April 3; published 2024 May 15

Abstract

The gamma-ray burst (GRB) GRB 211211A is believed to have occurred due to the merger of two neutron stars or a neutron star and a black hole, despite its duration of more than a minute. Subsequent analysis has revealed numerous interesting properties including the possible presence of a ~ 22 Hz quasiperiodic oscillation (QPO) during precursor emission. Here we perform timing analysis of Fermi and Swift gamma-ray data on GRB 211211A and, although we do not find a strong QPO during the precursor, we do find an extremely significant 19.5 Hz flux oscillation, which has higher fractional amplitude at higher energies, in a ~ 0.2 s segment beginning ~ 1.6 s after the start of the burst. After presenting our analysis we discuss possible mechanisms for the oscillation.

Unified Astronomy Thesaurus concepts: [Black holes \(162\)](#); [Gamma-ray bursts \(629\)](#); [Gamma-rays \(637\)](#); [Neutron stars \(1108\)](#); [Relativistic binary stars \(1386\)](#)

1. Introduction

There are believed to be two basic categories of gamma-ray bursts (GRBs; Kouveliotou et al. 1993): those powered by a particular type of core-collapse supernova, which typically produce long bursts (durations of tens of seconds or larger), and those produced by the merger of two neutron stars or possibly a neutron star and a black hole, which typically produce short bursts (durations of a few seconds or shorter). A growing number of GRBs blur the lines between these categories. For example, GRB 211211A lasted for more than a minute, yet the spectrum and especially the apparent presence of a kilonova after the burst suggest a merger rather than a core-collapse supernova (Rastinejad et al. 2022; Troja et al. 2022; Yang et al. 2022).

Further elucidation of the nature of GRBs in general, and anomalous bursts such as GRB 211211A in particular, could be obtained with the detection of quasiperiodic oscillations (QPOs) in the gamma-ray light curve. For example, QPOs with frequencies $\nu \gtrsim 1000$ Hz could be related to oscillations of a hypermassive neutron star or an accretion disk shortly after merger (Chirenti et al. 2019); indeed, evidence for such oscillations was found in GRB 910711 and GRB 931101B by Chirenti et al. (2023), and similar frequencies were seen in a magnetar giant flare by Castro-Tirado et al. (2021).

Lower-frequency oscillations, on the order of ~ 10 – 100 Hz, are also predicted in several models. For example, if a neutron star merges with a rapidly spinning, low-mass black hole, then coherent Lense–Thirring precession of the resulting accretion disk could lead to QPOs in this frequency range (Stone et al. 2013; Li et al. 2023). Searches for QPOs in this range

were performed by, e.g., Dichiara et al. (2013) with negative results.

Here we report the detection of a strong oscillation, at a frequency $\nu \approx 19.5$ Hz, shortly after the beginning of the main part of GRB 211211A. The oscillation is independently evident in both Swift and Fermi gamma-ray data. Interestingly, a ~ 22 Hz QPO from the precursor of this same burst was reported by Xiao et al. (2022) and modeled as a seismic aftershock from a resonant crust shattering by Suvorov et al. (2022). We see a power excess at this frequency during the precursor, but not with a high enough significance to claim detection. The 19.5 Hz oscillation (which we will sometimes call a QPO although we do not formally resolve the frequency width of the oscillation) lasts for ~ 0.2 s, starts and ends abruptly, and has a higher fractional amplitude at higher photon energies. Compared with a red-noise-only model, the Bayes factor in favor of a QPO is $\sim 6 \times 10^{10}$ in the Swift Burst Alert Telescope (BAT) data alone, and $\sim 4 \times 10^4$ in the Fermi Gamma-ray Burst Monitor (GBM) data alone. We discuss different interpretations of this oscillation in Section 4, after presenting the description of our data in Section 2 and our analysis in Section 3.

2. Description of Data and Burst

GRB 211211A triggered the Swift BAT (see Barthelmy et al. 2005) at 13:09:59.634 UT on 2021 December 11 (D’Ai et al. 2021), and triggered the Fermi GBM (see Meegan et al. 2009) just 0.017 s later, at 13:09:59.651 UT on 2021 December 11 (Fermi GBM Team 2021). The burst was in the direction R.A. = 14h09m10^s.12, decl. = +27:53:18.1 (J2000), with an estimated redshift of $z = 0.076$ ($D \approx 350$ Mpc) based on a galactic association (Malesani et al. 2021). The burst lasted more than a minute, but later association with a kilonova (Rastinejad et al. 2022; Troja et al. 2022; Yang et al. 2022) suggests that this was a long-duration “short” GRB likely



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

associated with the merger of two neutron stars or of a neutron star and a black hole. Because we are interested in the timing properties of this burst, in this section we give details about how we extracted the Swift BAT and Fermi GBM data and how we time-aligned the data from the two satellites, as well as displaying the light curves and power spectra.

2.1. Extraction of Swift BAT Data

For the QPO analysis, we created a 100 μ s non-mask-weighted light curve in 15–350 keV using the BAT event data from the Swift BAT GRB catalog.⁸ The event data was created using the standard BAT GRB analysis tool, “batgrbproduct⁹,” version 2.48, which is part of the HEASoft analysis package. The 100 μ s non-mask-weighted light curve was created using the BAT analysis tool, “batbinevt”,¹⁰ version 1.48. The non-mask-weighted (i.e., non-background-subtracted) light curve was used because the QPO analysis in this work requires that the data obey Poisson statistics and do not have negative or fractional photon counts (which can be the case for a mask-weighted light curve).

2.2. Extraction of Fermi GBM Data

To study the prompt emission of GRB 211211A we used the time-tagged event (TTE) data obtained from the two most illuminated sodium iodide (NaI) detectors (N2 and Na). We processed the data using the HEASoft (version 6.30.1; NASA High Energy Astrophysics Science Archive Research Center Heasarc 2014) and the FermiTools software packages (version 2.0.8; Fermi Science Support Development Team 2019) following standard procedures.¹¹ Light curves were reconstructed using a 0.1 ms time bin, and considering different energy ranges (8–1000 keV, 4–37 keV, 37–88 keV, 88–166 keV, and >166 keV) using the FSELECT and GTBIN tools.

2.3. Time Alignment of the Swift BAT and Fermi GBM Data

Because the segment of interest is short and the 19.5 Hz signal appears to emerge and disappear suddenly, we align the starting times of the Swift BAT and Fermi GBM data prior to performing our analysis.

We start with the trigger times: as indicated above, the Fermi GBM trigger time is 0.017 s later than the Swift BAT trigger time. We then note that at the time of trigger, the Fermi satellite was at longitude 197°17 and latitude 24°57, at an altitude of 524 km (derived using the Fermi GBM Data Tools; Goldreich et al. 2022), whereas the Swift satellite was at longitude 7°12 and latitude of 20°36, at an altitude of 538 km. Given the time, R.A., and decl. of the burst, at the time of the burst, it was above latitude 27°89 and longitude 294°.

To figure out the projection of the direction to the GRB onto the Fermi–Swift vector, we first compute the Fermi–Swift vector in Cartesian coordinates. The average radius of the Earth is 6371 km, so at the time of the burst the distance from the center of the Earth to Swift was $r_S = 6371 + 538 = 6909$ km and the distance from the center of the Earth to Fermi was

$r_F = 6371 + 524 = 6895$ km. At the time of the GRB, the colatitude for Swift was $\theta_S = 90^\circ - 20^\circ 36' = 69^\circ 24' = 1.2154$ radians and the azimuth for Swift was $\phi_S = 7^\circ 12' = 0.1243$ radians. Similarly, during the time of the GRB, the colatitude for Fermi was $\theta_F = 90^\circ - 24^\circ 57' = 65^\circ 3' = 1.1420$ radians and the azimuth for Fermi was $\phi_F = 197^\circ 17' = 3.4413$ radians. The three-dimensional locations of Swift and Fermi at the time of the burst were then

$$\begin{aligned} \text{Swift} &= r_S (\sin \theta_S \cos \phi_S, \sin \theta_S \sin \phi_S, \cos \theta_S) \\ &= (6427, 803, 2404) \text{ km} \\ \text{Fermi} &= r_F (\sin \theta_F \cos \phi_F, \sin \theta_F \sin \phi_F, \cos \theta_F) \\ &= (-5991, -1851, 2866) \text{ km}. \end{aligned} \quad (1)$$

Therefore the Fermi–Swift vector at the time of the burst was

$$\text{Swift} - \text{Fermi} = (12418, 2654, -462) \text{ km}. \quad (2)$$

The direction to the burst at the time of the burst had $\theta_{\text{GRB}} = 90^\circ - 27^\circ 89' = 62^\circ 11' = 1.0840$ radians and $\phi_{\text{GRB}} = 294^\circ = 5.1313$ radians. The projected Fermi–Swift distance along the direction to the GRB equals the dot product of the unit vector toward the GRB, with the Fermi–Swift vector calculated above; positive means that the signal reached Swift first, whereas negative means that the signal reached Fermi first. The unit vector is

$$\begin{aligned} \hat{\Omega}_{\text{GRB}} &= (\sin \theta_{\text{GRB}} \cos \phi_{\text{GRB}}, \sin \theta_{\text{GRB}} \sin \phi_{\text{GRB}}, \cos \theta_{\text{GRB}}) \\ &= (0.3595, -0.8074, 0.4678), \end{aligned} \quad (3)$$

and the dot product of this with the Fermi–Swift vector is 2105 km, which is a light travel time of 0.007 s. Subtracted from the 0.017 s difference in trigger times, this means that to align the Swift BAT and Fermi GBM light curves we shift the Fermi count arrival times by 0.01 s compared with their nominal values.

2.4. Light Curves, Power Spectra, and Energy Dependence

The light curve of GRB 211211A is highly complex, as is evident from the Swift BAT and Fermi GBM light curves in Figure 1. Here, to show the structure more clearly, we show data binned to 12.8 ms intervals, although we use 0.1 ms intervals in our analysis. The Fermi GBM curve has been shifted following the procedure described in Section 2.3, and after this shift the correlations between the curves are evident. The right-hand panels show the segment that displays the strong ≈ 19.5 Hz signal in both the Swift BAT and the Fermi GBM data, and the middle panel also shows the precursor to the burst, from which Xiao et al. (2022) reported a ≈ 22.5 Hz QPO.

In Figure 2 we see the power spectra, using Fermi GBM and Swift BAT data independently, for our 0.2048 s segment (left panel) and for the first 0.2048 s of the precursor (right panel), which is the segment from which Xiao et al. (2022) reported a ≈ 22.5 Hz QPO. In our featured segment the excess power is clear in both Swift BAT and Fermi GBM data at ≈ 19.5 Hz, compared with the power at the next lower (≈ 14.6 Hz) and next higher (≈ 20.4 Hz) frequencies. The power spectra in Figure 2 show excess power in a single frequency (19.5 Hz), which therefore does not have a resolved width. The analysis in Section 3 demonstrates that the excess power is significant for both sets of data even when compared with a flexible red-noise

⁸ <https://swift.gsfc.nasa.gov/results/batgrbcat/>

⁹ <https://heasarc.gsfc.nasa.gov/ftools/caldb/help/batgrbproduct.html>

¹⁰ <https://heasarc.gsfc.nasa.gov/ftools/caldb/help/batbinevt.html>

¹¹ https://fermi.gsfc.nasa.gov/ssc/data/p7rep/analysis/scitools/gbm_grb_analysis.html

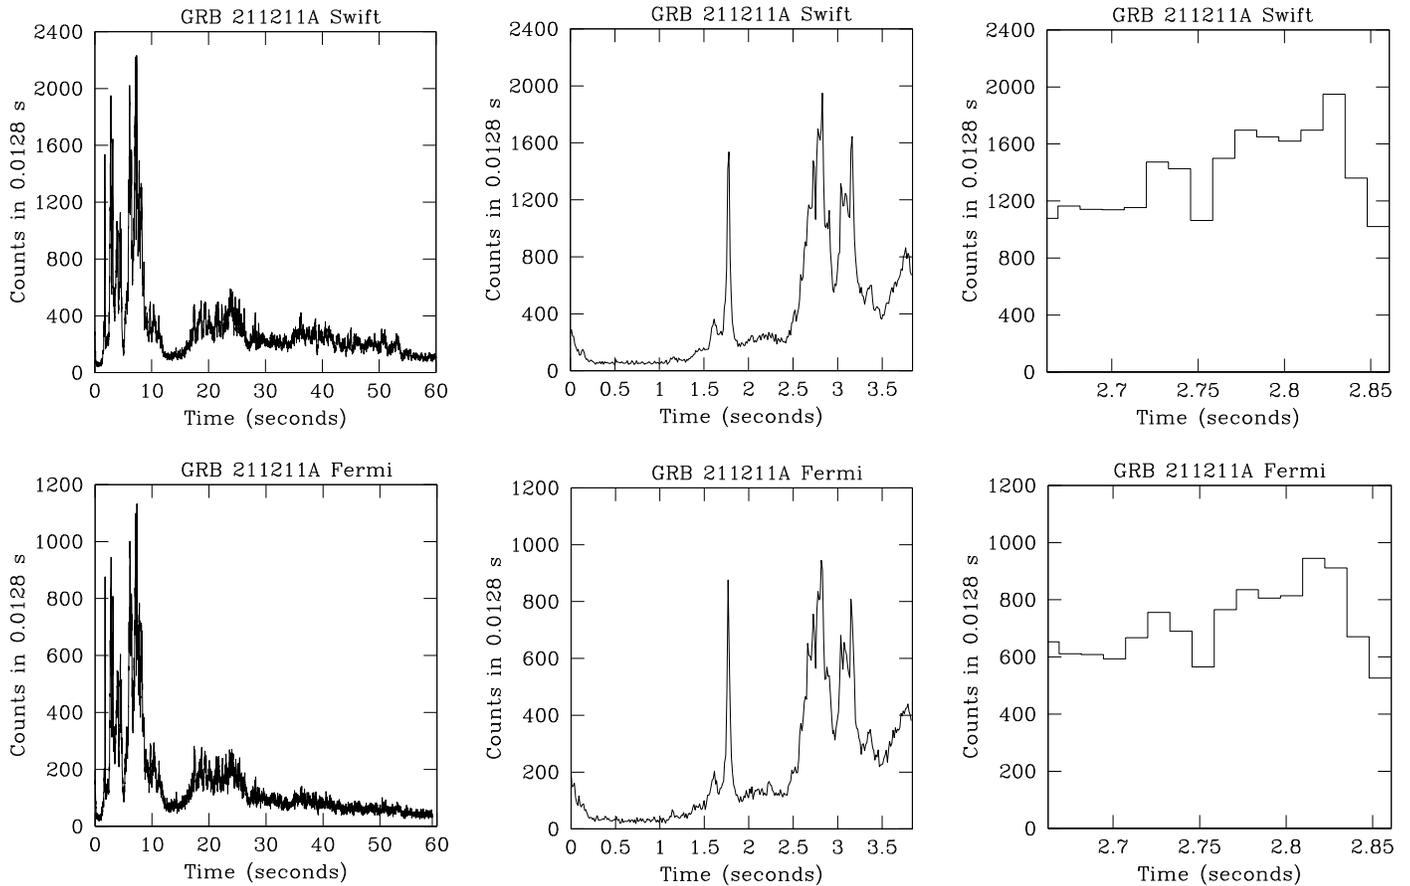


Figure 1. Light curves from Swift BAT (top panels) and Fermi GBM after the alignment procedure described in Section 2.3 (bottom panels) over different timescales. In each case, we bin the data in intervals of 12.8 ms to make the development of the light curve more evident, although all of our analysis is performed using a time resolution of 0.1 ms. The left, middle, and right panels show, respectively, the first minute of the burst, approximately the first 4 s (including the precursor at the beginning), and the 0.2048 s segment on which we focus our analysis. Note that the instruments triggered on the precursor, which is time zero in these plots, but that the main burst started ~ 1 s later. Thus our segment starts 2.66 s after the trigger but ≈ 1.6 s after the beginning of the main burst. The burst evidently has a long and complex light curve, and there is a strong correlation of substructure between the Fermi GBM and Swift BAT light curves, although BAT registers $\sim 2 \times$ the number of counts as GBM.

model that can accommodate multiple slopes. In contrast, although there is some excess power in the vicinity of the ≈ 22.5 Hz signal noted by Xiao et al. (2022), the significance is not high.

In Figure 3 we see the fractional rms amplitude versus energy for the Swift BAT data (left panel) and for the Fermi GBM data (right panel). For each data set, we break the data into four energy ranges with approximately equal numbers of counts. For the Swift BAT data, the energy ranges were roughly <37 keV, 37–70 keV, 70–126 keV, and >126 keV, up to a maximum energy of about 500 keV. For the Fermi GBM data, the energy ranges were roughly 4–37 keV, 37–88 keV, 88–166 keV, and >166 keV, up to a maximum energy of about 2000 keV. The vertical location of the solid red square in each energy range is the median amplitude (which, using our power normalization, equals $\sqrt{2P/N_{\text{counts}}}$ for a power P with N_{counts} counts) estimated using Equation (16) of Groth (1975), and the upper and lower error bars show the $\pm 1\sigma$ amplitude using the same equation. For both data sets, the amplitude rises with energy.

The dynamical power spectra shown in Figure 4 allow us to estimate the approximate duration of the 19.5 Hz signal. We can see excess power in the signal frequency from approximately 2500 ms to 2700 ms after T_0 . The strongest

part of the signal, highlighted in the right panel, lasts about 100 ms.

3. Analysis Methods and Results

It is notoriously difficult to establish the presence of a periodic signal in data dominated by red noise. This lesson has recently been reinforced in ongoing searches for binary supermassive black holes, where promising evidence for periodicity has often weakened with additional data (e.g., Liu et al. 2018; Dotti et al. 2023; see Vaughan et al. 2016 for a general discussion of false periodicities). One of the reasons for the difficulty, which applies equally well to GRB data, is that the red noise itself can have a structure that can be mistaken for periodicity. In this section, we discuss our approach, which allows the red noise to have a wide variety of shapes, and show that even with this flexibility the 19.5 Hz QPO stands out.

Our analysis follows the method of QPO detection described first in Miller et al. (2019), in the context of a search for QPOs in the tail of the giant flare from the soft gamma-ray repeater SGR 1806–20, where there is also significant red noise. The method was then used in Chirenti et al. (2023) to discover kilohertz QPOs in Burst and Transient Source Experiment data on GRB 910711 and GRB 931101B. In brief, the method performs a Bayesian model comparison between a model

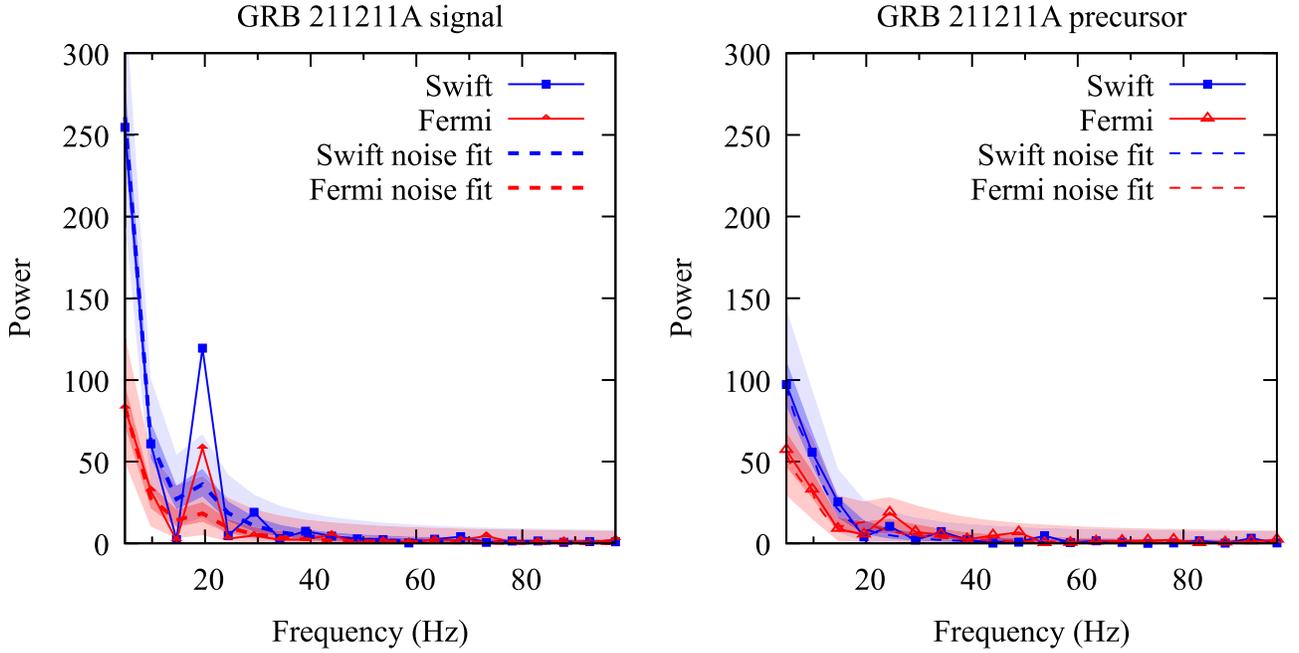


Figure 2. (Left panel) Power spectrum of our featured 0.2048 s segment. Here we use the normalization from Groth (1975), in which the average power is 1 from an intrinsically constant signal with purely Poisson noise. The frequency steps in the fast Fourier transform used to produce this power spectrum are $1/0.2048 \text{ s} \approx 4.88 \text{ Hz}$. The dashed lines show the fits to the data sets using the multislope-noise model described in Section 3, without a QPO; the noise-only model clearly underpredicts the power at 19.5 Hz. We also show with colored bands the $\pm 1\sigma$ (darker) and $\pm 3\sigma$ (lighter) single-trial power ranges for the best noise-only fits in each case (red for Fermi and blue for Swift). The key feature that makes the signal stand out in our analysis is the high power at $\approx 19.5 \text{ Hz}$, flanked by low powers at $\approx 14.6 \text{ Hz}$ and $\approx 20.4 \text{ Hz}$. This feature is seen independently in the Fermi GBM data (red lines and open red triangles) and in the Swift BAT data (blue lines and solid blue squares), which argues against an instrumental origin for this signal. No other interval in this burst has such a strong feature. (Right panel) Here, as a contrast to the left panel, we show the power spectrum from the 0.2048 s precursor to the burst. The normalization and line/point types are the same as in the left panel. This is the segment for which Xiao et al. (2022) reported a moderately significant $\sim 22.5 \text{ Hz}$ QPO. There is indeed an excess of power near that frequency in both the Swift BAT and the Fermi GBM data, but it is much weaker than the signal we feature.

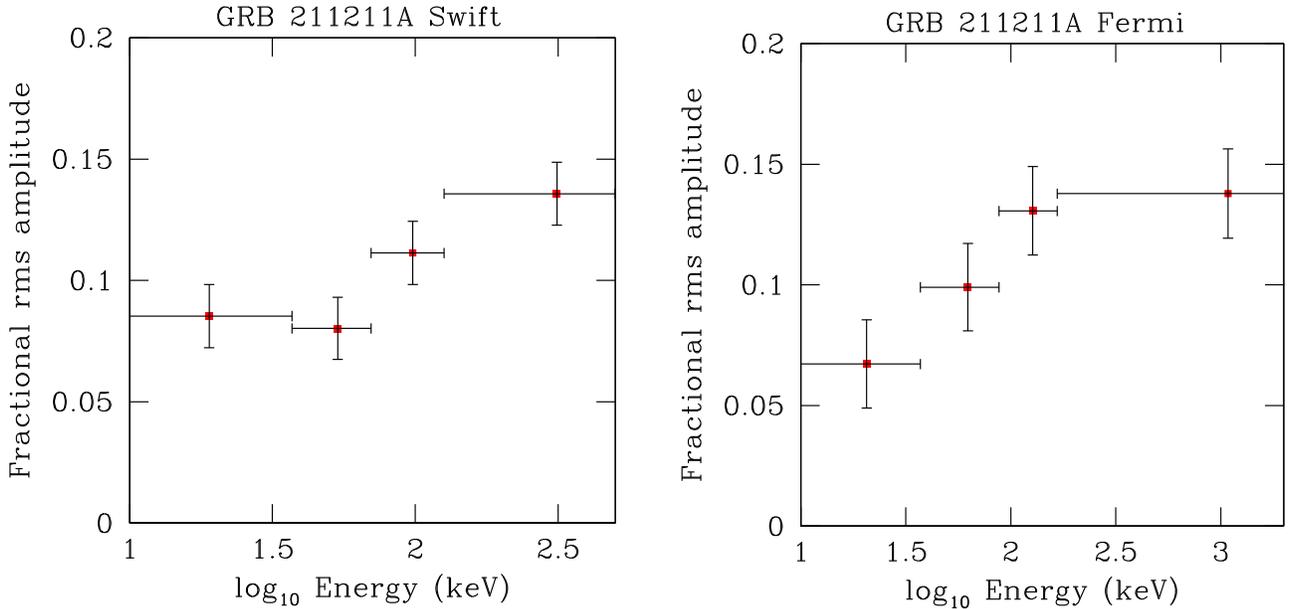


Figure 3. (left panel) Fractional rms amplitude of the $\approx 19.5 \text{ Hz}$ signal in four different energy ranges, selected to have approximately equal numbers of counts, from the Swift BAT data. The energy ranges were roughly $<37 \text{ keV}$, $37\text{--}70 \text{ keV}$, $70\text{--}126 \text{ keV}$, and $>126 \text{ keV}$, up to a maximum energy of about 500 keV, and are indicated by the horizontal bars. For each energy range, the solid red squares indicate the median of the estimated amplitude and the vertical error bars indicate the $\pm 1\sigma$ ranges of the amplitude, as inferred using the power distributions discussed in Groth (1975); see text for details. (Right panel) The same, for the Fermi GBM data. The energy ranges were roughly $4\text{--}37 \text{ keV}$, $37\text{--}88 \text{ keV}$, $88\text{--}166 \text{ keV}$, and $>166 \text{ keV}$, up to a maximum energy of about 2000 keV. We see that in both data sets there is a clear increase in fractional rms amplitude with increasing energy.

without a QPO (which could have excess red, white, or blue noise) and one with one or more Lorentzian QPOs (which can also have excess noise), using power spectral data. Here, within

our QPO model, we also encompass the possibility of a periodic oscillation, which in practice means a QPO with an unresolvably small frequency width. As indicated in

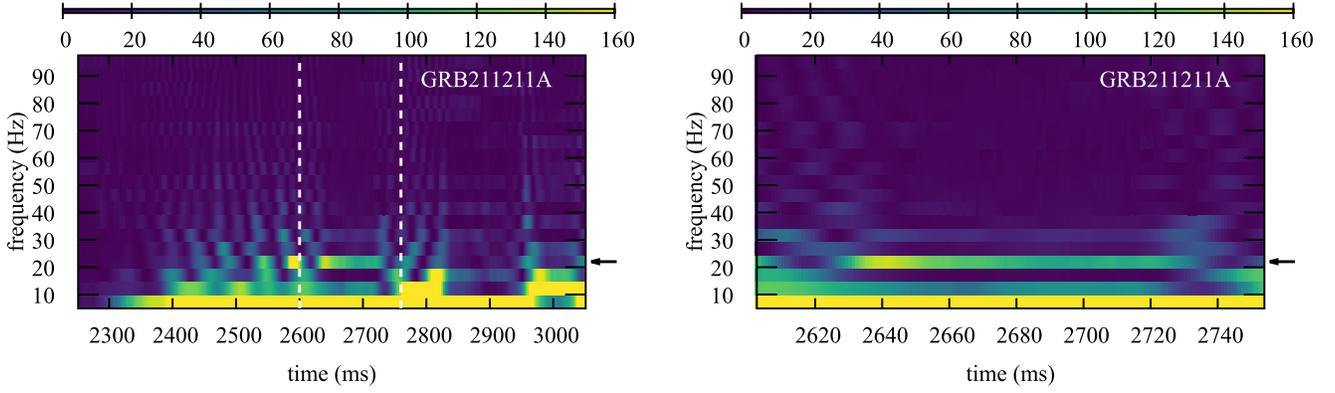


Figure 4. (Left panel) Spectrogram of the time range surrounding the segment in the Swift BAT data of GRB 211211A that shows a strong 19.5 Hz signal. We use a 0.2048 s window and shift it by 1 ms to cover the full segment. The power scale is saturated to match the highest power shown in the right panel. (Right panel) Same as the left panel but highlighting the strong signal.

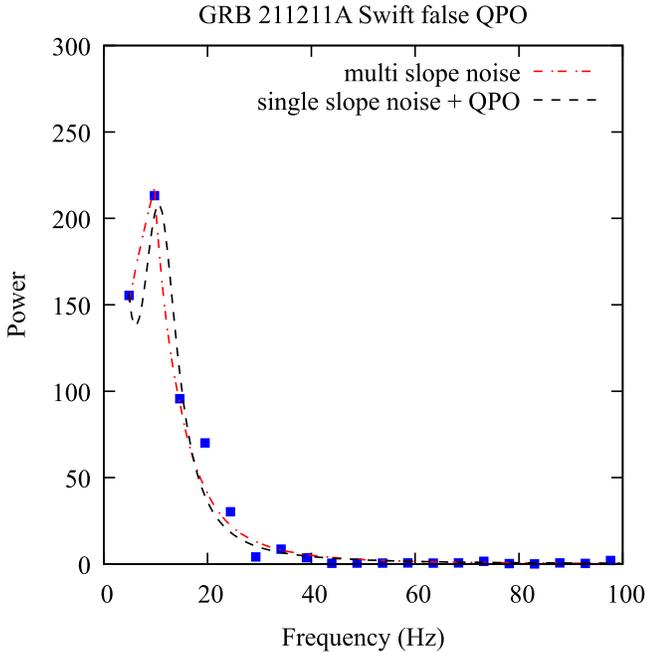


Figure 5. Example of a power spectrum from a segment of Swift BAT data (solid blue squares) from GRB 211211A which registers as an extremely strong QPO when compared with a single-slope red-noise model. The large magnitudes of the powers at low frequencies, combined with the local maximum in power at ≈ 9.7 Hz, gives a Bayes factor $>10^{31}$ in favor of a QPO (dashed black line) vs. the single-slope red-noise model. For this reason, we chose to employ a more flexible red-noise model with multiple slopes at low frequencies (dashed red line), which provides an adequate fit without requiring a QPO.

Section 2.4, we use the Groth (1975) power spectrum normalization, in which the mean power is 1 from an intrinsically constant signal with only Poisson noise.

As is evident from Figure 2, our segments, and indeed most segments of most GRBs, have substantial red noise at the low frequencies of interest in our analysis. We emphasize that this is a real, physical, variation; using the Groth (1975) normalization the chance probability of a power $P > P_0$ from purely Poisson noise with no intrinsic variability is e^{-P_0} . Thus in practice powers larger than a few tens are not produced by Poisson fluctuations. Note that in the formalism of Groth (1975), if there is a nonzero signal then the probability distribution of observed power is a series expansion (Equation (15) from Groth 1975) rather than a simple exponential.

Table 1
Priors on Power Spectral Models

Quantity	Prior (Flat in Indicated Range)
$A_{\text{noise}}(4.88 \text{ Hz})$	0–2000
Slope (4.88–9.77 Hz)	–1 to 3
Slope (9.77–14.65 Hz)	–1 to 3
Slope (14.65–19.53 Hz)	–1 to 3
Slope (19.53–24.41 Hz)	–1 to 3
A_{QPO}	0–200
$\log_{10} \nu_{\text{QPO}}(\text{Hz})$	1.0–3.7
$\log_{10} \Delta \nu_{\text{QPO}}(\text{Hz})$	–1 to 3

Note. Priors on our noise and noise+QPO models for the power spectra that we analyze. All quantities have flat priors in the indicated range, and the quantities with QPO subscripts are only used in the noise+QPO model.

However, we are focused not on the general continuum of red noise but on the possibly special implications of a QPO, which could point to a characteristic frequency in the system. With that in mind, the strong excess at ≈ 19.5 Hz in both Swift BAT and Fermi GBM data, flanked by much lower powers on either side, is worthy of investigation.

To pursue our analysis we need to decide on a red-noise-only model to compare with a model that has a QPO. In our initial analysis, we used red noise described by a single power law: $P(f) \propto f^{-\alpha}$, where α could range between $\alpha = -1$ (which is thus actually blue noise) and a fairly steep red noise slope of $\alpha = +3$. But in the long and complex light curve of GRB 211211A, there are segments with power spectra such as that featured in Figure 5. In this segment, the power at the second-lowest frequency is higher than the power at the lowest frequency, and the powers are large enough that a QPO model is favored overwhelmingly compared with a single-slope red noise-only model (Bayes factor $>10^{31}$).

Although the increase in power to the second-lowest frequency is formally significant, we elect to employ a more flexible red-noise model to ensure that local maxima in the power need to stand out substantially from background red noise. Note that for our segment length of 0.2048 s, which was inspired by the report from Xiao et al. (2022) of a ≈ 22.5 Hz QPO in a ~ 0.2 s portion of the precursor, the frequency resolution is $1/0.2048 \text{ s} \approx 4.88$ Hz and thus the initial several frequencies are 4.88 Hz, 9.77 Hz, etc. The priors on our models are displayed in Table 1.

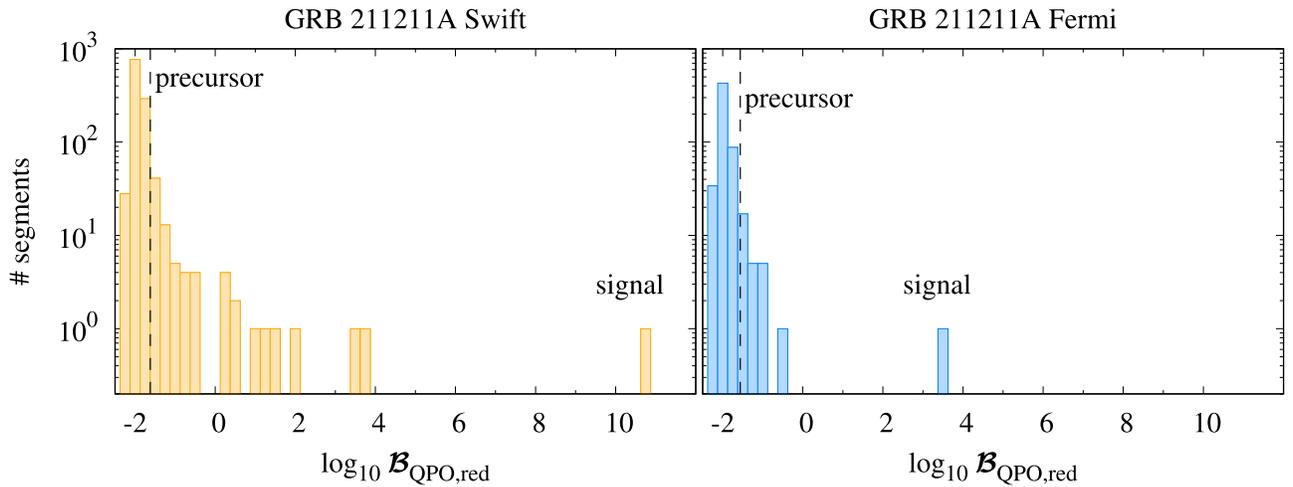


Figure 6. Differential distribution of \log_{10} Bayes factors between models with red noise plus a QPO and models with just red noise (see text for details) for half-overlapping 0.2048 s segments of Swift BAT data (left panel) and Fermi GBM data (right panel) from GRB 211211A. In each panel, our featured segment in the main burst with a 19.5 Hz QPO is highlighted, and the vertical dashed line indicates the Bayes factor for the precursor segment used by Xiao et al. (2022) to suggest the presence of a ≈ 22.5 Hz QPO. The evidence for a signal in our featured segment stands out overwhelmingly and independently in the Swift BAT and the Fermi GBM data. In contrast, although the Bayes factor for the precursor segment is above average in both data sets, the evidence for a QPO is not especially strong.

Table 2
Summary of Best Fits and Bayes Factors

Detector	ν (Hz)	$\Delta\nu$ (Hz)	Bayes Factor
Swift BAT	19.5	0.15	6.9×10^{10}
Fermi GBM	19.4	0.12	4.5×10^4

Note. Best fits and Bayes factors for our featured segment in the Swift BAT and Fermi GBM data. In practice, frequency widths of $\lesssim 1$ Hz cannot be distinguished from each other in the data. The centroid frequency ν and frequency width $\Delta\nu$ of the fitted Lorentzian QPOs are consistent between the two data sets, and the large Bayes factors compared with a noise-only model indicate that the signal is strong for both data sets independently.

We then compute the Bayes factor \mathcal{B} between the two models using the standard Bayesian prescription:

$$\mathcal{B}_{\text{QPO,red}} = \frac{\int \mathcal{L}(d|\theta_{\text{QPO}})q(\theta_{\text{QPO}})d\theta_{\text{QPO}}}{\int \mathcal{L}(d|\theta_{\text{red}})q(\theta_{\text{red}})d\theta_{\text{red}}}. \quad (4)$$

Here θ_{QPO} represents the vector of parameters for the QPO model, θ_{red} represents the vector of parameters for the red noise-only model, q is the (normalized) prior, and \mathcal{L} is the likelihood of the data d given the model. We assume that prior to our analyzing the data the models are equally probable, which means that the odds ratio $\mathcal{O}_{\text{QPO,red}}$ between the models equals $\mathcal{B}_{\text{QPO,red}}$.

We divide the Swift BAT data, and independently the Fermi GBM data, into segments of duration 0.2048 s ($=2^{11}$ times our time resolution of 0.0001 s), with consecutive segments overlapping by half their duration, i.e., by 0.1024 s (to reduce the probability that a short-lived signal will be missed). This results in 1171 segments of Swift BAT data and 579 segments of Fermi GBM data.

Figure 6 shows the resulting distribution of Bayes factors in Swift BAT and Fermi GBM data, and highlights those for our featured segment ($\mathcal{B} = 6.9 \times 10^{10}$ for Swift BAT and $\mathcal{B} = 4.5 \times 10^4$ for Fermi GBM) and the precursor (vertical dashed line).

The evidence for a signal in our segment stands out overwhelmingly, in both the Swift BAT and the Fermi GBM data, compared with any other segment. The evidence is stronger from the Swift BAT data than from the Fermi GBM data, due to the larger number of counts, but in both data sets independently the signal is strong (see Table 2 for a summary of the best fits and Bayes factors).

We also checked that our model provides an acceptable description of the data. Using the formulae of Groth (1975) we generated numerous synthetic data sets from our best-fit noise +QPO models and computed the log-likelihood of the synthetic data (up to 100 Hz, i.e., the first 20 frequencies). The log-likelihood of the Swift BAT data is at the 8th percentile and that of the Fermi GBM data is at the 30th percentile of the corresponding sets of synthetic log-likelihoods. Thus our model has captured the essential features of the low-frequency portions of the power spectra.

It is therefore clear that for these data sets the red noise plus QPO model that we employ fits the data far better than the red noise-only model. It is, however, difficult to judge whether this is the correct red-noise model and whether it would be reasonably common for the natural high-amplitude variability of GRBs to counterfeit a signal similar to what we see in GRB 211211A.

To provide an independent measure of the significance of the signal we would like to use a model without QPOs, generate synthetic light curves with that model, and then compare the results with the data. We lack a physical picture with which to select such models. We therefore follow the guidance of Hübner et al. (2022) and use Gaussian processes. More specifically, we use as a smooth light curve model a triangular shape with a slow rise and a faster decline (which is roughly similar to the average light curve in the right-hand panels of Figure 1). A least-squares fit of that functional form gives the following for counts per 0.0064 s interval (where t is in units of seconds and $t = 0$ is the start of the segment). Swift BAT:

$$\begin{aligned} t \leq 0.1674 \text{ s: counts} &= 547.4 + 2426.674t \\ t > 0.1674 \text{ s: counts} &= 953.625 \\ &- 14794.124(t - 0.1674). \end{aligned} \quad (5)$$

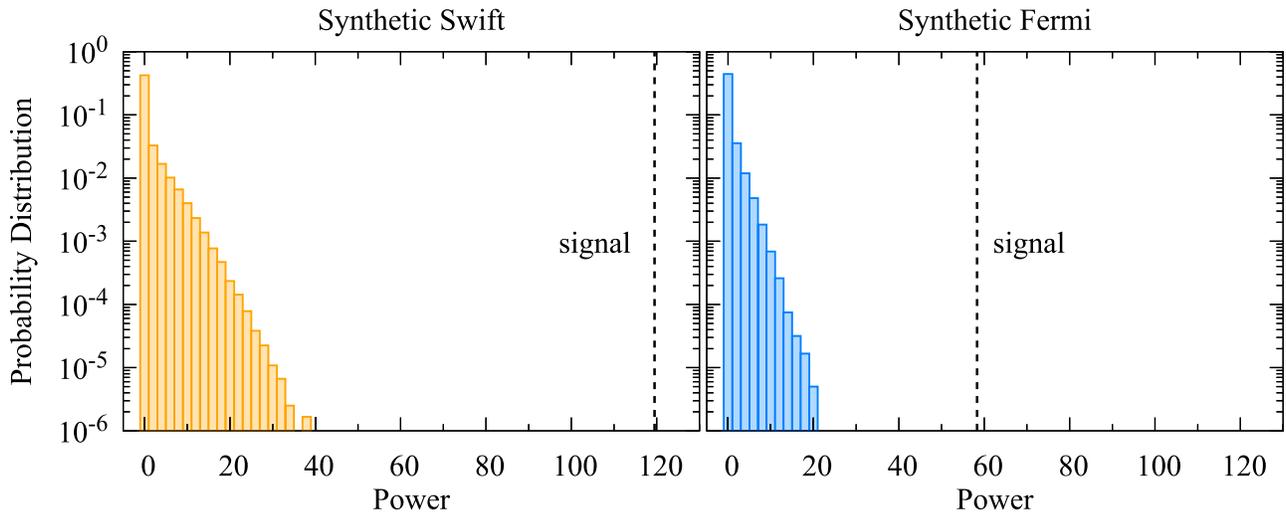


Figure 7. Differential distribution of powers at $f \geq 19.5$ Hz for light curves generated using Gaussian processes (see text for details) for synthetic light curves similar to the Swift BAT data (left panel) and the Fermi GBM data (right panel), for 50,000 realizations each. The vertical dashed line in each panel shows the observed power in the 19.5 Hz QPO. For both the Swift BAT and the Fermi GBM data sets on our segment, the power is much larger than what emerges in our synthetic data sets.

Fermi GBM:

$$\begin{aligned} t \leq 0.1668 \text{ s: counts} &= 284.622 + 943.776t \\ t > 0.1668 \text{ s: counts} &= 442.044 \\ &- 6409.395(t - 0.1668). \end{aligned} \quad (6)$$

We then sampled from a Gaussian process and added the result to this overall shape (see Hübner et al. 2022 for a discussion of Gaussian processes in this context). We used a squared exponential kernel such that the covariance for two samples separated by time τ is

$$\kappa(\tau) = \sigma^2 \exp[-\tau^2/2\ell^2], \quad (7)$$

with parameters σ (the overall scale of the variance) and ℓ (the duration over which correlations decline). Based on an approximate fit to the covariances of the Swift BAT data, we chose $\sigma = 80$ and $\ell = 0.05$ s, and from a fit to the Fermi GBM data we chose $\sigma = 40$ and $\ell = 0.1$ s.

We then generated 50,000 synthetic light curves for Swift BAT, and 50,000 synthetic light curves for Fermi GBM, using these Gaussian processes. As a check that our approach gives red noise levels comparable to what we see in the data, we note that the observed Swift BAT powers at 4.9 Hz and 9.8 Hz are, respectively, at the 44th and 53rd percentiles of the powers at those frequencies in the synthetic Swift BAT data, and the observed Fermi GBM powers at those frequencies are, respectively, at the 45th and 95th percentiles at those frequencies in the synthetic Fermi GBM data. The power distributions are thus roughly consistent with what we see in the observations.

Figure 7 shows the cumulative distribution of the synthetic powers seen at a frequency of 19.5 Hz or higher, and compares that distribution with what is observed. For both the Swift BAT and the Fermi GBM data sets, the observed power is much larger than any power seen in the synthetic data. The distribution of powers in the synthetic data sets is not clear, but at the highest powers, an exponential distribution appears roughly consistent with the data. Using this extrapolation suggests that perhaps a billion times as many samples would be needed for there to be a good chance of obtaining powers as large as are observed in either data set.

Because we do not have a well-understood physical model for the details of the light curve of GRB 211211A or other GRBs, our quantitative results cannot be considered definitive. However, it does suggest that signals as strong as we see are not easily produced even given large variations in the count rate.

Our final clue regarding this signal is that it starts and ends abruptly. We see this in the dynamical power spectrum Figure 4, where the 19.5 Hz power is large only for a brief time. We can also see this from the Bayes factors: $\mathcal{B}_{\text{QPO,red}} \approx 6.9 \times 10^{10}$ for our featured segment in the Swift BAT data, but for the previous segment (which we recall overlaps half of our featured segment) $\mathcal{B}_{\text{QPO,red}} = 0.32$, and for the following segment (also half-overlapping), $\mathcal{B}_{\text{QPO,red}} = 0.043$.

We draw the following conclusions from the results of this section:

1. There is strong red noise in our featured segment of GRB 211211A and in many other segments from this burst.
2. However, the power at ≈ 19.5 Hz, lasting for 0.2048 s and starting roughly 2.66 s after the burst trigger, plus the much lower powers at the next lowest and at the next highest frequencies in our power spectra, makes this segment stand out from any other in the burst. The exact significance of the feature depends on the model of red noise, but the Bayes factor of a QPO model relative to a red-noise model is orders of magnitude greater than it is for any other segment.
3. The 19.5 Hz signal is seen strongly, and independently, in the Swift BAT and the Fermi GBM data, and has similar characteristics (e.g., frequency and frequency width). Thus the signal is very unlikely to be an instrumental artifact.
4. The 19.5 Hz signal is narrow: there is negligible excess power at ± 5 Hz compared with the main signal.
5. The 19.5 Hz signal has a higher fractional amplitude at higher photon energies in both data sets.
6. The 19.5 Hz feature starts and ends abruptly; the full duration of the signal is not much longer than the ≈ 0.2 s of our segment.

7. The 22.5 Hz QPO suggested by Xiao et al. (2022) to exist in the burst precursor does not, in our analysis, appear to be especially significant. We note that Xiao et al. (2022) base their conclusions on the analysis of the detrended light curve and consider only Poisson noise in their estimate of the QPO significance. This approach can cause false positives of periodic pulsations in Fourier analysis (see, e.g., Auchère et al. 2016).

It is clear, based on the magnitudes of the powers, that the apparent 19.5 Hz signal has an astrophysical origin rather than being caused by instrumental effects or statistical fluctuations. It is less certain that the feature we discovered indicates the presence of a narrow, coherent frequency. However, in the next section we will proceed under the assumption that during the short duration of the signal, 19.5 Hz is characteristic, and seek possible causes.

4. Implications and Explanations for the QPO

From the previous section, we found that (1) there is a strong ≈ 19.5 Hz signal in (2) a short (≈ 0.2 s) interval of GRB 211211A, which (3) is much narrower than the ≈ 5 Hz resolution of our power spectra and (4) has a higher fractional amplitude at higher photon energies. Assuming that this is a characteristic frequency of the system that is evident for only a short time, what are some possible physical causes?

Of the QPO features listed above, the one that is likely to be the easiest to explain in the widest variety of models is the increase of fractional amplitude with increasing photon energy. Any model with a periodically changing spectrum that has a steeply decreasing flux at higher energies will show this behavior. For example, if the temperature T of a blackbody changes periodically then the fractional amplitude of the modulation at energies many times kT will be much larger than the fractional amplitude at energies $\sim kT$. Thus although this observed feature of our QPO might be considered a rough confirmation of the physical reality of the feature, it does not discriminate between models.

We thus instead begin by considering what sources can produce frequencies of order 19.5 Hz. The characteristic frequency of an object of average density $\bar{\rho}$ is $\sim (G\bar{\rho})^{1/2}$. Because 19.5 Hz is well above the < 1 Hz maximum for white dwarfs and less dense objects, these are disfavored (see the similar discussion in Gold 1968 for why pulsars cannot be white dwarfs). Thus, even independently of GRB 211211A being a GRB, the 19.5 Hz QPO points to a neutron star or black hole origin.¹²

The fundamental ringdown frequency of a black hole is $\sim 10^4 \text{ Hz} (M_{\odot}/M)$, multiplied by a factor of order unity that depends on the black hole spin parameter and the harmonic of the mode. Thus a $\sim 500 M_{\odot}$ black hole would have a frequency in the vicinity of our signal (see, e.g., Table 1 of Kokkotas & Schmidt 1999). The observed quality factor of $Q = \pi f / \Delta f \gtrsim \pi 20 \text{ Hz} / 1 \text{ Hz} \approx 60$ is quite high for a black hole ringdown but would be possible if the spin parameter is $\gtrsim 0.99$ (e.g., Echeverria 1989). However, the observation of a kilonova from this GRB, which suggests that a neutron star was

disrupted, is not consistent with such a high-mass black hole because a neutron star would enter the horizon without being torn apart. It therefore appears that the signal originated from a mode or rotation of a neutron star, or from some aspect of an accretion disk around either a neutron star or a black hole.

Neutron star p-modes, including the fundamental f-mode, are much too high-frequency (> 1000 Hz) to explain our signal. Neutron star g-modes are lower in frequency but are expected to be at least hundreds of Hz and are thus also too high in frequency, although the exact value is dependent on the equation of state and could be lower (see, e.g., Figure 2 of Kuan et al. 2022, which shows g-mode frequencies as low as 60 Hz). The frequencies could be lower for a proto neutron star because of its lower density, but this state is expected to evolve rapidly in density and thus it is difficult to understand how it would produce a frequency as sharp as our signal.

QPOs with frequencies comparable to the 19.5 Hz signal have been seen in giant flares from the soft gamma-ray repeaters SGR 1900 + 14 and SGR 1806–20 (Israel et al. 2005; Strohmayer & Watts 2005, 2006; Huppenkothen et al. 2014; Pompe et al. 2018; Miller et al. 2019), and the frequency width is often comparable to the ~ 1 Hz upper limit that we infer for the 19.5 Hz signal (Miller et al. 2019). There is not a clear consensus about the origin of these SGR QPOs, but candidates include torsional modes of the crust and magneto-hydrodynamic (MHD) modes in the core. A challenge to crustal models of our signal is that because the oscillation is evident less than 2 s after the start of the main burst, it is implausible that a hard crust would have formed. MHD modes are not as easy to disprove, although for this and other frameworks, there remains the question of why the signal starts and stops abruptly.

A neutron star could rotate at a frequency compatible with our signal. The initial rotation rate after the merger would be high, in the vicinity of ~ 1500 Hz, which means that it would need to slow down within ~ 1.6 s–19.5 Hz. Candidate mechanisms for the slowdown include pulsar-like vacuum magnetic dipole radiation, the interaction of a stellar magnetic field with matter falling back onto the remnant, and gravitational radiation from an asymmetric star (see, e.g., Shapiro & Teukolsky 1983 for the relevant formulae).¹³ We find that even for a star of ellipticity unity, gravitational radiation would take tens of thousands of seconds to spin a star down to 19.5 Hz, so this is insufficient. The mechanisms involving magnetic fields both need field strengths $\sim 10^{18}$ G to work in 1.6 s, which is two orders of magnitude larger than has been inferred from any other star but might not be impossible.

However, the strongest argument against this scenario is that the energy released due to spindown is larger by a factor of several than even the isotropic equivalent energy release for GRB 211211A. Minaev et al. (2021) estimate a total isotropic equivalent energy release from 1 keV to 10 MeV of 1.16×10^{52} erg. High-density equations of state which sustain masses $> 2 M_{\odot}$ have maximum-mass moments of inertia

¹² The characteristic frequency limit may be slightly exceeded for oscillations with low amplitude. E.g., QPOs with frequencies of a few Hz have been predicted for the X-ray spectra of magnetic cataclysmic variables. Observations of such systems have reported QPOs with amplitudes of 1%–5% rms and periods of 1–2.5 s in their optical light curves (Busschaert et al. 2015).

¹³ It has been suggested that a plateau phase in the X-ray afterglow could be a signature of a magnetar central engine in short GRBs, and Rowlinson et al. (2013) reported that a plateau may be present in about half of all short GRBs detected with Swift. Within this scenario, Lasky & Glampedakis (2016) constrained the ellipticity of a nascent magnetar in eight short GRBs from X-ray observations and concluded that the gravitational waves emitted during the spindown are unlikely to be detectable with LIGO, requiring next-generation ground-based gravitational wave detectors. Upper limits for pulsations in the X-ray data were obtained by Rowlinson et al. (2017).

$I \sim 2 \times 10^{45} \text{ g cm}^2$ (e.g., Cook et al. 1994). At an angular frequency $\Omega = 2\pi \times 1500 \text{ rad s}^{-1}$, the rotational energy is $E_{\text{rot}} = \frac{1}{2}I\Omega^2 \approx 9 \times 10^{52} \text{ erg}$. Thus if the star spun down to 19.5 Hz, the fluence we would see would be much larger than what was observed from GRB 211211A.

Another possibility is free precession of the merger remnant if it is not rotating around one of its principal axes (we thank Zorawar Wadiasingh for suggesting this possibility). For an oblate star, the precession frequency is roughly the rotation frequency multiplied by the fractional difference in the moments of inertia (see Goldreich 1970; Ruderman 1970 for early discussion of neutron star precession). This would imply an oblateness on the order of $\sim 1\%$ – 2% , which seems plausible. Physically, if the direction of the jet is modified by precession then the observed flux could be modulated at this frequency.

The last possibility involves an accretion disk. It has been suggested (e.g., Stone et al. 2013; Li et al. 2023) that if a rapidly rotating black hole tidally disrupts a neutron star, and if the resulting accretion disk has an axis that is not aligned with the black hole rotation axis, then at high accretion rates Lense–Thirring precession could drive the disk to solid-body precession (Fragile & Anninos 2005) which would have a frequency in the ~ 10 – 100 Hz range. The modulation we see could be due to the precession of a jet aligned with the disk axis. If this is the explanation, then it suggests that the black hole had low mass (because otherwise the neutron star would not be disrupted outside the horizon) and high enough spin to produce strong Lense–Thirring precession.

One of the most significant challenges to any model of the 19.5 Hz signal is to explain how it starts and then ends abruptly. We could speculate that, for example in the precessing disk idea, it takes a certain amount of time for the disk to lock into solid-body rotation; prior to that, there would not be a clear direction to the disk axis and thus no definite frequency. Once the disk is in solid-body rotation, it could be that alignment with the black hole rotation axis and/or rapid draining of the disk into the black hole drops the amplitude quickly. Another consideration could be optical depth: perhaps the system needed to clear out some amount of matter before the QPO could be observed. A full explanation almost certainly will require convincing numerical simulations, which are beyond the scope of this paper.

5. Conclusions

We have presented evidence for a strong 19.5 Hz signal in the Swift BAT data, and independently in the Fermi GBM data, for GRB 211211A. Although this burst lasted for more than a minute, other characteristics (most notably the evidence for an associated kilonova: Rastinejad et al. 2022; Troja et al. 2022; Yang et al. 2022) suggest that it was a prolonged burst after the merger of two compact objects, rather than resulting from the core collapse of a massive star.

The oscillation is evident only in a $\sim 0.2 \text{ s}$ interval beginning $\approx 1.6 \text{ s}$ after the start of the main burst. The signal is also very narrow in frequency, with a width that is $\lesssim 1 \text{ Hz}$, and its fractional amplitude increases with increasing energy in both the Swift BAT and the Fermi GBM data sets.

Of the models we considered, precession seems most consistent with the observed features. One possibility is Lense–Thirring precession of a remnant accretion disk after the disruption of a neutron star by a low-mass and rapidly spinning black hole (e.g., Stone et al. 2013; Li et al. 2023). This

would involve a black hole with a mass and spin that might not be represented in the current gravitational wave samples (Abbott et al. 2019, 2021; The LIGO Scientific Collaboration et al. 2023). Another possibility, which does not seem to have been explored in this context, is the free-body precession of the merger remnant, which in that case would not have collapsed to a black hole by the time the QPO is evident. In either case, targeted numerical simulations will be needed to determine whether a compact object coalescence could produce the behavior that we see, and in particular to produce a coherent signal that lasts for only a short time.

Acknowledgments

We thank Chuck Horowitz, Konstantinos Kalapotharakos, Stephen Lasage, David Radice, Jeff Scargle, and Zorawar Wadiasingh for discussions. C.C. acknowledges support by NASA under award No. 80GSFC21M0002. M. C.M. was supported in part by NASA ADAP grants 80NSSC20K0288 and 80NSSC21K0649. S.D. was supported by NASA under award No. 80NSSC22K1516. This work was partially conducted at the Aspen Center for Physics, which is supported by National Science Foundation grant PHY-1607611.

ORCID iDs

Cecilia Chirenti  <https://orcid.org/0000-0003-2759-1368>
 Simone Dichiarà  <https://orcid.org/0000-0001-6849-1270>
 Amy Lien  <https://orcid.org/0000-0002-7851-9756>
 M. Coleman Miller  <https://orcid.org/0000-0002-2666-728X>

References

- Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2019, *PhRvX*, 9, 031040
 Abbott, R., Abbott, T. D., Abraham, S., et al. 2021, *PhRvX*, 11, 021053
 Auchère, F., Froment, C., Bocchialini, K., Buchlin, E., & Solomon, J. 2016, *ApJ*, 825, 110
 Barthelmy, S. D., Barbier, L. M., Cummings, J. R., et al. 2005, *SSRv*, 120, 143
 Busschaert, C., Falize, É., Michaut, C., Bonnet-Bidaud, J. M., & Mouchet, M. 2015, *A&A*, 579, A25
 Castro-Tirado, A. J., Østgaard, N., Göçüç, E., et al. 2021, *Natur*, 600, 621
 Chirenti, C., Dichiarà, S., Lien, A., Miller, M. C., & Preece, R. 2023, *Natur*, 613, 253
 Chirenti, C., Miller, M. C., Strohmayer, T., & Camp, J. 2019, *ApJL*, 884, L16
 Cook, G. B., Shapiro, S. L., & Teukolsky, S. A. 1994, *ApJ*, 424, 823
 D’Ai, A., Ambrosi, E., D’Elia, V., et al. 2021, *GCN*, 31202, 1
 Dichiarà, S., Guidorzi, C., Frontera, F., & Amati, L. 2013, *ApJ*, 777, 132
 Dotti, M., Bonetti, M., Rigamonti, F., et al. 2023, *MNRAS*, 518, 4172
 Echeverria, F. 1989, *PhRvD*, 40, 3194
 Fermi GBM Team 2021, *GCN*, 31201, 1
 Fermi Science Support Development Team 2019, *Fermitools: Fermi Science Tools*, Astrophysics Source Code Library, ascl:1905.011
 Fragile, P. C., & Anninos, P. 2005, *ApJ*, 623, 347
 Gold, T. 1968, *Natur*, 218, 731
 Goldreich, A., Cleveland, W. H., & Kocevski, D. 2022, *Fermi GBM Data Tools: v1.1.1*, <https://fermi.gsfc.nasa.gov/ssc/data/analysis/gbm>
 Goldreich, P. 1970, *ApJL*, 160, L11
 Groth, E. J. 1975, *ApJS*, 29, 285
 Hübner, M., Huppenkothen, D., Lasky, P. D., et al. 2022, *ApJ*, 936, 17
 Huppenkothen, D., Watts, A. L., & Levin, Y. 2014, *ApJ*, 793, 129
 Israel, G. L., Belloni, T., Stella, L., et al. 2005, *ApJL*, 628, L53
 Kokkotas, K. D., & Schmidt, B. G. 1999, *LRR*, 2, 2
 Kouveliotou, C., Meegan, C. A., Fishman, G. J., et al. 1993, *ApJL*, 413, L101
 Kuan, H.-J., Krüger, C. J., Suvorov, A. G., & Kokkotas, K. D. 2022, *MNRAS*, 513, 4045
 Lasky, P. D., & Glampedakis, K. 2016, *MNRAS*, 458, 1660
 Li, Y., Shen, R.-F., & Zhang, B.-B. 2023, *ApJ*, 955, 98
 Liu, T., Gezari, S., & Miller, M. C. 2018, *ApJL*, 859, L12
 Malesani, D. B., Fynbo, J. P. U., de Ugarte Postigo, A., et al. 2021, *GCN*, 31221, 1

- Meehan, C., Lichti, G., Bhat, P. N., et al. 2009, *ApJ*, 702, 791
- Miller, M. C., Chirenti, C., & Strohmayer, T. E. 2019, *ApJ*, 871, 95
- Minaev, P., Pozanenko, A. & GRB IKI FuN 2021, GCN, 31230, 1
- NASA High Energy Astrophysics Science Archive Research Center (Heasarc) 2014, HEASoft: Unified Release of FTOOLS and XANADU, Astrophysics Source Code Library, ascl:1408.004
- Pumpe, D., Gabler, M., Steininger, T., & Enßlin, T. A. 2018, *A&A*, 610, A61
- Rastinejad, J. C., Gompertz, B. P., Levan, A. J., et al. 2022, *Natur*, 612, 223
- Rowlinson, A., O'Brien, P. T., Metzger, B. D., Tanvir, N. R., & Levan, A. J. 2013, *MNRAS*, 430, 1061
- Rowlinson, A., Patruno, A., & O'Brien, P. T. 2017, *MNRAS*, 472, 1152
- Ruderman, M. 1970, *Natur*, 225, 838
- Shapiro, S. L., & Teukolsky, S. A. 1983, Black holes, white dwarfs and neutron stars. The physics of compact objects (New York: Wiley)
- Stone, N., Loeb, A., & Berger, E. 2013, *PhRvD*, 87, 084053
- Strohmayer, T. E., & Watts, A. L. 2005, *ApJL*, 632, L111
- Strohmayer, T. E., & Watts, A. L. 2006, *ApJ*, 653, 593
- Suvorov, A. G., Kuan, H. J., & Kokkotas, K. D. 2022, *A&A*, 664, A177
- The LIGO Scientific Collaboration, the Virgo Collaboration, the KAGRA Collaboration, et al. 2023, *PhRvX*, 13, 041039
- Troja, E., Fryer, C. L., O'Connor, B., et al. 2022, *Natur*, 612, 228
- Vaughan, S., Uttley, P., Markowitz, A. G., et al. 2016, *MNRAS*, 461, 3145
- Xiao, S., Zhang, Y.-Q., Zhu, Z.-P., et al. 2022, arXiv:2205.02186
- Yang, J., Ai, S., Zhang, B.-B., et al. 2022, *Natur*, 612, 232