



Fermi-GBM Observations of GRB 210812A: Signatures of a Million Solar Mass Gravitational Lens

P. Veres¹ , N. Bhat¹ , N. Fraija² , and S. Lesage¹

¹ Center for Space Plasma and Aeronomic Research, University of Alabama in Huntsville, 320 Sparkman Drive, Huntsville, AL 35899, USA; peter.veres@uah.edu

² Instituto de Astronomía, Universidad Nacional Autónoma de México, Apartado Postal 70-264, 04510 México, CDMX, Mexico

Received 2021 September 27; revised 2021 October 11; accepted 2021 October 11; published 2021 November 4

Abstract

Observing gravitationally lensed objects in the time domain is difficult, and well-observed time-varying sources are rare. Lensed gamma-ray bursts (GRBs) offer improved timing precision for this class of objects, complementing observations of quasars and supernovae. The rate of lensed GRBs is highly uncertain, approximately one in 1000. The Gamma-ray Burst Monitor on board the Fermi Gamma-ray Space Telescope has observed more than 3000 GRBs, making it an ideal instrument to uncover lensed bursts. Here we present observations of GRB 210812A showing two emission episodes, separated by 33.3 s and with a flux ratio of about 4.5. An exhaustive temporal and spectral analysis shows that the two emission episodes have the same pulse and spectral shape, which poses challenges to GRB models. We report multiple lines of evidence for a gravitational lens origin. In particular, modeling the lightcurve using nested sampling, we uncover strong evidence in favor of the lensing scenario. Assuming a point-mass lens, the mass of the lensing object is about 1 million solar masses. High-resolution radio imaging is needed for future lens candidates to derive tighter constraints.

Unified Astronomy Thesaurus concepts: [Gamma-ray bursts \(629\)](#); [Gravitational lensing \(670\)](#)

1. Introduction

Strong gravitational lensing is a tool that serendipitously enhances our observing capabilities and offers new opportunities to study the universe (see, e.g., Oguri 2019). Gamma-ray bursts (GRBs) are energetic transient sources at cosmological distances, involving relativistic jets from stellar-mass black hole (BH) central engines. The GRBs last from a fraction of a second to about 1000 s and typically show nonthermal spectra (see, e.g., Kumar & Zhang 2015; Beloborodov & Mészáros 2017, for reviews). Given that the distance scale of GRBs spans a wide range (up to redshift $z \lesssim 9$; Cucchiara et al. 2011), a fraction of GRBs will show the imprints of strong gravitational lensing.

Strong gravitational lensing produces multiple images of the same source. The images differ in their intensity, but importantly, their spectral shapes will be the same. Similarly, in time-varying sources, the temporal profile will be the same but shifted in time for different images (Schneider et al. 1992). The temporal and spectral invariance is the main defining feature of gravitational lensing. Lensed images are separated by up to arcseconds, which is clearly below the resolution of current gamma-ray detectors. Gamma-ray instruments, however, have excellent time resolution, and temporal structures can be recorded with unparalleled accuracy (Meegan et al. 2009).

In one incarnation of the lensing scenario (Refsdal 1964; Rodney et al. 2021) applied to GRBs, also called macrolensing, a single event triggers the same instrument twice and can be separated anywhere from a few hours to decades. The two triggers will have similar light-curve shapes and spectra. The delay between the events scales linearly with the lens mass. Lens candidates for this scenario have masses in the 10^8 – $10^{12} M_{\odot}$ range. Objects in this mass range include supermassive BHs (up

to a few $\times 10^{10} M_{\odot}$), galaxies, and clusters of galaxies. In practice, however, all traditional strong-lensing time-delay measurements are from galaxies or clusters of galaxies. The GRB lensing rates are highly uncertain, but somewhere on the order of one in 1000 GRBs should be affected (Mao 1992). After roughly 10,000 observed GRBs during three decades, no convincing macrolensing candidate GRB pair has been found (Nemiroff et al. 1994; Veres et al. 2009; Davidson et al. 2011; Hurley et al. 2019; Ahlgren & Larsson 2020). The negative result is likely a combination of two effects. First, GRB detectors, typically in low Earth orbit, will miss a sizable fraction of GRB lens echoes (see, however, Hurley et al. 2019; Hui & MoonBEAM Team 2021, for existing and future all-sky instruments). Second, for weaker GRBs, with pulses close to the noise level, it is difficult to distinguish between the lensing scenario and just two unrelated but similar looking GRBs (Ahlgren & Larsson 2020).

In a different scenario, also called millilensing (because the expected separation between the images is on the order of milliarcsseconds; Nemiroff et al. 2001), the gravitational lens signature is imprinted upon the lightcurve of a single trigger. In this case, we have, e.g., two emission episodes with similar light-curve patterns that can be separated by timescales spanning from a fraction of a second to a few minutes.

Recently, there has been an increase in claims of millilensing events. Paynter et al. (2021) presented convincing evidence for lensing in the short-duration BATSE GRB 950830. Mukherjee & Nemiroff (2021a) raised some concerns based on the inconsistent flux ratio between the two pulses. Wang et al. (2021) and Yang et al. (2021) independently argued that the likely short Gamma-ray Burst Monitor (GBM) GRB 200716C shows millilensing signatures. Kalantari et al. (2021) reported on a different GBM lensing candidate, the long-duration GRB 090717, selected based on the analysis of the autocorrelation function. The claim of Kalantari et al. (2021) was challenged by Mukherjee & Nemiroff (2021b), arguing that the two pulse shapes differ significantly. For the above-reported lensing

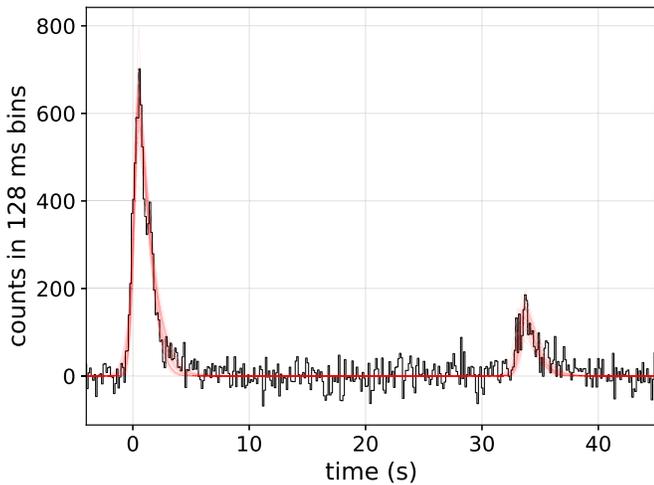


Figure 1. Background-subtracted lightcurve of GRB 210812A. The data are summed over all GBM detectors with good coverage. The red lines show a subset of MCMC fits using the N1 pulse (see Equation (3) and Section 4.1.1).

candidates, the flux of the first and second emission episodes is either at the same level or a ratio of $\lesssim 1.5$.

We present observations of the long-duration GRB 210812A and show that it is consistent with a gravitational-lensing scenario. It is the first lensing claim with a flux ratio $\gtrsim 3$. We list multiple lines of evidence to support the lensing interpretation. We perform spectral and temporal analysis using Fermi-GBM data and complement it with additional data from INTEGRAL-SPI/ACS and Swift/BAT.

In Section 2, we present the observations, followed by tests for gravitational lensing in Section 3. We discuss the power of the tests in Section 5 and conclude in Section 6.

2. Observations

Fermi-GBM triggered on GRB 210812A (Figure 1; trigger number 650479626/210812699) on 2021 August 12 at 16:47:01.014 UT (T0; Fermi-GBM Team 2021; Veres & Fermi-GBM Team 2021). Fermi-GBM consists of 12 NaI (referred to as n0, n1, ..., n9, na, and nb) and two BGO (referred to as b0 and b1) detectors covering the entire unocculted sky in the 8–1000 keV and ~ 0.1 –40 MeV energy range, respectively.

As reported by the automatic pipeline, the GRB location is R.A. = $39^\circ.7$, decl. = $69^\circ.7$, with an error radius of $1^\circ.1$ (statistical only). At the trigger time, this position corresponds to an LAT boresight angle of 149° . The GRB location was behind the spacecraft, meaning most GBM detector normals have a large angle to the source. Based on previous experience (e.g., Connaughton et al. 2016), this type of geometry results in a large number of detectors showing approximately equal count rates compared to \sim three detectors with a dominating signal for typical GRBs with a small LAT boresight angle. Upon visual inspection of the NaI detectors’ lightcurves in the 50–300 keV range, where the GRBs are brightest, we find that all 12 NaI detectors detected the brighter first peak. Moreover, detectors n1 and n6–nb also detected the fainter second pulse. Both of the pulses were detected by the two BGO detectors.

Detectors n8, na, and nb showed the strongest signals. We used data from these detectors, along with b1, for spectral analysis. For the temporal analysis, we used detectors n1, n6–nb, b0, and b1. For both spectral and temporal analysis, we

used the 128 energy channel time-tagged event (TTE) data. We chose the prebinned, eight energy channel (ctime) data to carry out the flux ratio test.

The targeted search (Blackburn et al. 2015; Goldstein et al. 2019) was designed to search for coherent subthreshold signals (weak signals that did not trigger the instrument). As expected, we recovered both pulses of GRB 210812A with high significance. The search also provides a location for the burst (R.A. = $40^\circ.5$, decl. = $69^\circ.4$), consistent with the location of the automatic pipeline used for standard GRB analysis. We also find that the locations of the two pulses are consistent, meaning they do indeed belong to the same source.

2.1. Other Observations

The location of GRB 210812A was outside the coded mask of Swift-BAT (Gehrels et al. 2004; Barthelmy et al. 2005) at the time of the trigger. The first peak of GRB 210812A is clearly present in the continuous four-channel data,^{3,4} but the second peak is only discernible in the summed lightcurve.

INTEGRAL/SPI-ACS (von Kienlin et al. 2003; Winkler et al. 2003) detected GRB 210812A and clearly shows the two-peak structure.⁵ We calculate the time delay between Fermi-GBM and ACS due to the higher altitude of the INTEGRAL spacecraft to be $dt_{\text{ACS}} = -0.396$ ms. We applied this correction to the ACS lightcurve in Figure 2.

2.2. Temporal Properties

Object GRB 210812A consists of two pulses separated by a quiescent period of about 30 s (Figure 1). Using the GBM targeted search, we found no emission between the two pulses. For background estimates, we fit a third-degree polynomial based on quiescent segments of the lightcurve before, after, and between the two pulses. We also report no discernible emission around 33 s after the second pulse, indicating that this is not a periodic source. We further note that low-level emission discernible in the summed lightcurve just before the start of the second pulse (T0+26.5 to T0+29.6 s) is inconsistent with coming from the location of GRB 210812A; thus, it is unrelated.

The duration of GRB 210812A is $T_{90} = 39.9 \pm 3.6$ s (10–1000 keV; T_{90} marks the time interval between 5% and 95% of the cumulative flux).⁶ Analyzed separately, the two pulses have consistent duration; the duration of the first pulse is $T_{90,1} = 5.31 \pm 0.68$ s, while the second pulse is $T_{90,2} = 3.84 \pm 1.64$ s. Taking the first pulse as a separate GRB, we classify it based on the GBM T_{90} distribution (Bhat et al. 2016; von Kienlin et al. 2020) as a likely long GRB originating from the collapse of a massive star, with a probability of 87%. Conversely, the likelihood of GRB 210812A being a short GRB from a compact binary merger, based on the $T_{90,1}$ information and observed distribution of Fermi-GBM GRBs and calculated consistently with Goldstein et al. (2017) and Rouco Escorial et al. (2021), is 13%.

Autocorrelation function—We measure the time delay between the two pulses using the autocorrelation function of

³ <https://www.swift.ac.uk/archive/reproc/00096363010/bat/rate/sw00096363010brtms.lc.gz>

⁴ <https://www.swift.ac.uk/archive/reproc/00036688035/bat/rate/sw00036688035brtms.lc.gz>

⁵ <http://isdc.unige.ch/~savchenk/spiacs-online/spiacs.pl>

⁶ <https://gcn.gsfc.nasa.gov/gcn3/30633.gcn3>

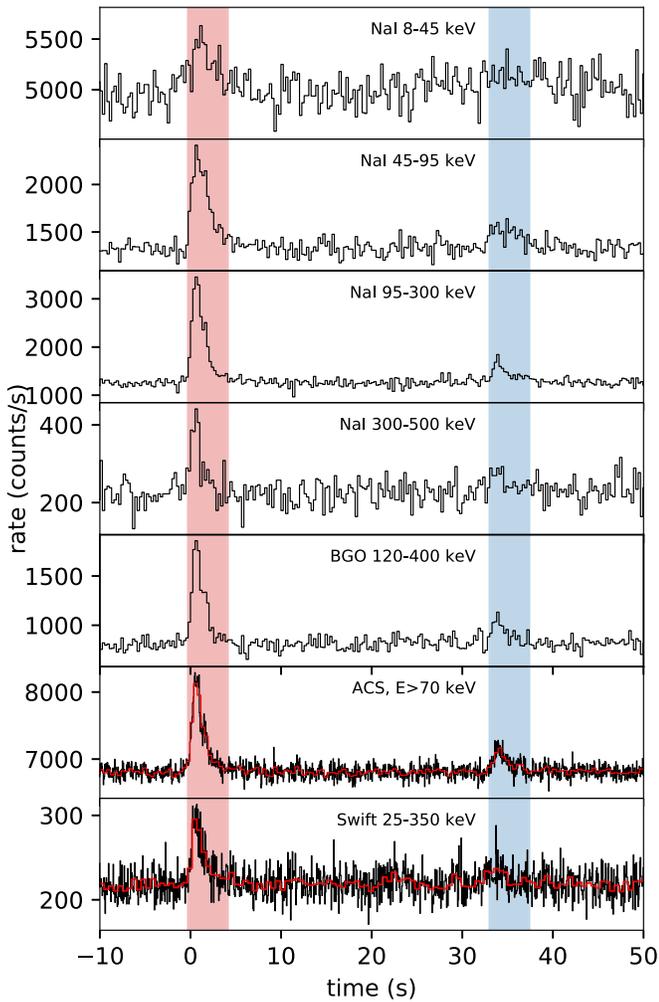


Figure 2. The GBM lightcurves in different energy ranges, summed over all detectors with good coverage. The ACS and Swift lightcurves are also shown in the bottom two panels with the native resolution (black; 50 ms for ACS and 64 ms for BAT) and binned (red; 0.4 s for ACS and 0.512 s for BAT).

the lightcurve between 10 and 1000 keV, with 64 ms resolution. To accurately determine the autocorrelation peak, we fit a ninth-degree polynomial to the peak region. We estimate the uncertainty of the delay by adding Poisson noise to the original lightcurve (see, e.g., Ukwatta et al. 2010; Hakkila et al. 2018; Ji et al. 2018 for a similar approach). After repeating the peak finding procedure, we take the uncertainty as the 1σ confidence region of the delays of the modified lightcurves and get

$$\Delta t_{\text{ACF}} = 33.30^{+0.12}_{-0.11} \text{ s}. \quad (1)$$

This is the first of many time-delay measurements between the pulses (Figure 3). We note here that this accuracy is typical of what one would expect for a lensed GRB separated by a longer timescale (macrolensing).

The peak of the autocorrelation curve shows a 3.15σ excess over the smoothed curve calculated using the Savitzky–Golay filter. This excess satisfies the criteria of Paynter et al. (2021) for lensing, which classifies GRB 210812A as a lensing candidate for further scrutiny.

Spectral lag—The lag is a measure of the delay between high- and low-energy photons (e.g., Norris et al. 2000). Typically, a GRB is initially harder, and the higher-energy

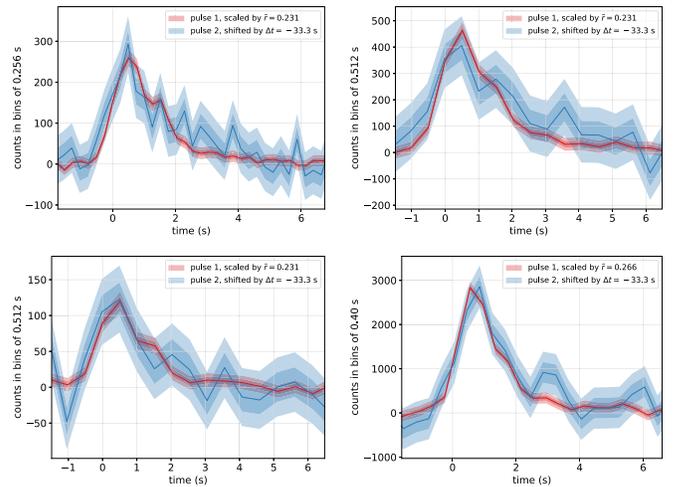


Figure 3. Comparison of the lightcurves of the two pulses. The stronger first pulse is scaled down to match the second pulse. The fainter second pulse is shifted in time to match the first pulse. Dark and light regions are 1σ and 2σ regions, respectively. Top left: NaI (45–300 keV) + BGO (120–400 keV); top right: NaI (22–800 keV); bottom left: BGO (120–400 keV); bottom right: ACS (>70 keV).

photons (100–300 keV) arrive earlier than the lower-energy (25–50 keV) photons. This relation is also used to estimate the redshift based on the empirical correlation between lag and luminosity (Norris et al. 2000). For the first pulse, we find that the lag is $\tau_{\text{lag},1} = 221.6^{+123.4}_{-139.6}$ ms, while the second pulse has $\tau_{\text{lag},2} = -89.2^{+255.9}_{-206.5}$ ms. The error on the second pulse is much larger because of its weakness, especially in the 25–50 keV range.

2.3. Spectral Analysis

The spectrum of the first pulse is best fit by a power law with an exponential cutoff, also known as the Comptonized model (see Table 1 for the parameters and Figure 4 for the spectrum). The fluence in the 10–1000 keV range is $F_1 = (80.1 \pm 0.3) \times 10^{-7} \text{ erg cm}^{-2}$.

The second pulse is weaker, and it is fit equally well by a simple power law and the Comptonized model. We chose the Comptonized model because it is more physical (the power law with photon index > -2 integrates to infinite energy). The difference in the goodness-of-fit measure (ΔC -stat) when going from the simpler power law (two parameters) to the Comptonized model (three parameters) is ~ 6 . This is just below the ΔC -stat ≈ 8 used in, e.g., Poolakkil et al. (2021) to select the more complex model. However, all parameters of the Comptonized model are well constrained (Table 1), which justifies the use of this model. The fluence of pulse 2 is $F_2 = (21.6 \pm 2.6) \times 10^{-7} \text{ erg cm}^{-2}$.

We can also compare the time-resolved spectra of the two pulses. Because of the relative weakness of the second pulse, we chose to fit the simple power-law model in bins of 0.256 s. We show the temporal evolution of the power-law indices for the two pulses and a linear fit to both as a function of time (Figure 5). We shifted the second pulse by 33.3 s to highlight their similar behavior.

3. Indicators of Lensing Origin

Proving the gravitational-lensing origin involves showing that the two pulses have identical pulse shapes, and their

Table 1
Spectrum of the Two Pulses Fit by a Power Law with Exponential Cutoff

Interval (s)	E_{peak} (keV)	Photon Index	C-stat/dof	Photon Flux (ph cm $^{-2}$ s $^{-1}$)	Energy Flux (10 $^{-7}$ erg cm $^{-2}$ s $^{-1}$)
−0.256 to 4.096	324 ± 28	−0.88 ± 0.08	373.3/310	10.35 ± 0.46	18.4 ± 0.7
33.024 to 37.376	283 ± 90	−1.10 ± 0.24	295.3/304	3.48 ± 0.49	4.96 ± 0.59

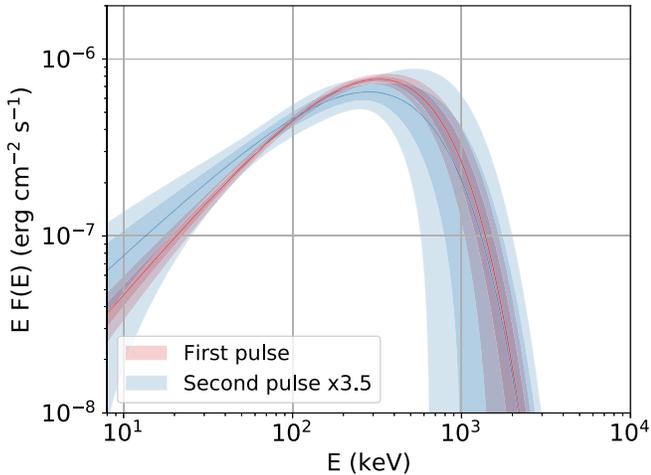


Figure 4. Spectra of the two pulses. The second pulse is shifted to match the first. Shaded regions mark 1σ and 2σ confidence regions.

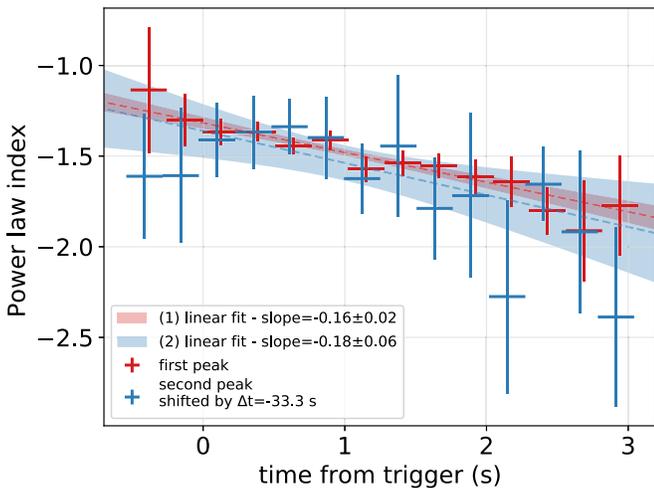


Figure 5. Temporal evolution of the spectral power-law index compared for the two pulses. The slopes of the fitted linear functions are indicated in the legend, and the shaded regions mark 1σ confidence intervals.

spectra are similar as well. It is conceivable that some physical mechanism produces two pulses with identical pulse shapes and spectra with no lensing involved. Without imaging the sources based solely on the gamma-ray observations, a lensing scenario is difficult to prove beyond the shadow of a doubt.

In this section, we will show, however, that the gamma-ray properties of GRB 210812A pass all the tests in the literature for a lensing origin. Furthermore, we apply the Bayesian evidence criterion that can select among models, and this method indicates strong evidence in favor of the lensing scenario.

The basic idea for establishing the statistical likelihood of lensing in the temporal domain is comparing two scenarios. First, we assume no lensing. We fit the pulse model and allow every parameter to vary freely. In the alternative scenario, where lensing is assumed, the second pulse is forced to have the same shape as the first one and differ only by a normalization factor and a time delay.

We use two types of pulses (Norris et al. 1996, 2005). The first pulse with five parameters, referred to as “N1,” is defined as

$$I_{N1}(t|A, t_{\text{max}}, \sigma_r, \sigma_d, \nu) = A \begin{cases} e^{-\left(\frac{t_{\text{max}}-t}{\sigma_r}\right)^\nu} & \text{if } t < t_{\text{max}} \\ e^{-\left(\frac{t-t_{\text{max}}}{\sigma_d}\right)^\nu} & \text{if } t > t_{\text{max}} \end{cases}$$

Here A is the amplitude at peak time, t_{max} ; σ_r and σ_d mark the rise and decay timescales of the pulse; and ν is a shape parameter.

The second pulse with four parameters, “N2,” has the following temporal dependence:

$$I_{N2}(t|A, \xi, \Delta, \tau) = A \begin{cases} e^{-\xi\left(\frac{t-\Delta}{\tau}\right)} & \text{if } t > \Delta \\ 0 & \text{if } t < \Delta, \end{cases} \quad (2)$$

where A is the amplitude, ξ is the asymmetry parameter, Δ is the start time of the pulse, and τ is a duration parameter.

3.1. Indirect Evidence

We explore a few properties of GRB 210812A that are not direct proofs of lensing, but they are necessary to any such claim.

Hard-to-soft evolution—The GRBs typically become softer with time (e.g., Kaneko et al. 2006). This trend can be observed in the individual pulses of GRB 210812A as well (see Figure 5). For GRBs in general, the hard-to-soft evolution can be observed even across pulses. Specifically, for GRBs that show two pulses separated by a quiescent period, the second pulse is, in general, softer (e.g., Zhang et al. 2012; Lan et al. 2018). We find that the second pulse of GRB 210812A has the same spectral shape within the errors or, subsequently, the same hardness, which is uncommon for typical GRBs.

Leading pulse is brighter—In the case of simple lens models, the light ray traveling closer to the lens arrives later. It has a lower magnification than the first arriving light ray with the larger impact parameter (Krauss & Small 1991). We find that the first pulse is indeed visibly brighter and thus consistent with the expectation from a lensed source. We note that this criterion may not hold for complex lens models (Keeton & Moustakas 2009).

3.2. Spectra of the Pulses

For the spectral analysis, we first selected the interval containing the second pulse by visually identifying contiguous temporal bins with significant signal. For the first pulse, we selected a source interval that is the same length as the second

pulse (see Table 1 and Figure 2). The spectral fits of the two pulses yield consistent spectral shapes within the errors; the peaks of the energy-per-decade or νF_ν spectra are $E_{\text{peak},1} = 324 \pm 28$ and $E_{\text{peak},2} = 283 \pm 90$ keV. To compare the two spectra, we plot the 1σ and 2σ confidence regions of the spectral shapes (Figure 4), accounting for the correlations between the parameters. We multiply the second pulse by the fiducial 3.5 number to show that the two spectral shapes are consistent.

3.3. Count Ratio Test

If the pulses are gravitationally lensed, the ratio between pulses should not depend on energy. Mukherjee & Nemiroff (2021a) investigated the count ratio (CR) of the two pulses as a function of energy for GRB 950830 and found a $\gtrsim 2\sigma$ inconsistency. This test has the advantage that it is independent of the particular spectral shape.

For GBM, we used the eight-channel `ctime` data to carry out this test on GRB 210812A. We considered all detectors where the second pulse was visible in any channel (n1, n6–nb, energy channels 1–6, b0 and b1, and energy channels 0 and 1) and data from ACS and BAT (25–350 keV). We find that the ratio of the pulses in all channels and all three instruments is consistent with the mean value within 1.6 standard deviations (see Figure 5). We thus conclude that GRB 210812A passes the CR test for lensing.

3.4. χ^2 Test

A simple and robust test that does not assume any pulse shape was introduced by Nemiroff et al. (2001). This test was recently applied to the claim of Kalantari et al. (2021) on GRB 090717 by Mukherjee & Nemiroff (2021b). Mukherjee & Nemiroff (2021b) conclude that based on the χ^2 test, the claim of gravitational lensing can be excluded at the 5σ level.

This test considers the binned lightcurves of the two pulses as representing two distributions and asks if they are consistent with coming from the same parent distribution. After appropriately rescaling the first pulse and taking the background into account, we perform a χ^2 test for the hypothesis that the two lightcurves are drawn from the same distribution.

The test statistic is defined the following way:

$$\chi^2 = \sum_i \frac{(\tilde{r}P_1(t) - P_2(t + \Delta t))^2}{\tilde{r}^2P_1(t) + P_2(t + \Delta t) + B_1(t) + B_2(t + \Delta t)}, \quad (3)$$

where the t is the time, $\tilde{r} < 1$ is the scaling factor, P_i marks the background-subtracted counts in the two pulses, and B_i is the background counts ($i = \{1, 2\}$).

As an example (see Figure 3, top left), we consider data from the $T_0 - 2$ to $T_0 + 7$ s interval (first pulse) and compare it with the interval shifted by $\Delta t = 33.3$ s (second pulse). We use 0.256 s resolution, summed over the detectors with good signal in the 45–300 keV range, and add the signal of the BGO detectors (120–400 keV). We use the `tte` instead of `ctime` data because the beginning of the first pulse has uneven temporal binning in the `ctime` data.

First, we find the minimum of the χ^2 expression in Equation (3) as a function of \tilde{r} , and we get $\tilde{r} = 0.231$ (see Figure 3, top left). Next, we calculate the minimum $\chi^2 = 24.3$ value for 34 degrees of freedom (dof), which corresponds to a p -value of 0.89. Thus, there is no statistically significant difference between the two distributions.

We compare the lightcurves of the two pulses for different temporal resolutions, energy ranges, and instruments in Figure 3. We consistently find that there is no statistically significant difference between the two pulses. For example, Mukherjee & Nemiroff (2021c) performed a preliminary χ^2 analysis of GRB 210812A and concluded that there is an $\sim 2.8\sigma$ discrepancy between the two pulses. Using the same energy range and detector selection as Mukherjee & Nemiroff (2021c; assuming that in their notation detectors are numbered 1–12 and `ctime` energy channels 1–8), we performed the χ^2 test on 512 ms resolution lightcurves in the energy range 22–800 keV (NaI detectors only; Figure 3, top right). We find that $\tilde{r} = 0.231$ minimizes χ^2 at a value of $\chi^2 = 11.5$ for 16 dof (p -value of 0.78), indicating that the two lightcurves are consistent with being drawn from the same distribution. The BGO lightcurve (512 ms resolution) yields $\chi^2 = 14.027$ (dof = 16) and a p -value of 0.597, and for ACS (0.4 s resolution), $\chi^2 = 24.662$ (dof = 22) and a p -value of 0.314 (bottom two panels of Figure 3).

3.5. Bayesian Model Comparison

Applying an idea from gravitational-wave model selection, Paynter et al. (2021) introduced the Bayesian evidence to compare the lensing scenario to the case where there is no lensing.

In the *no-lens* scenario, we fit the lightcurve with two pulses, with all parameters left to vary. We derive the Bayesian evidence \mathcal{Z}_{NL} by integrating the likelihood over the multidimensional parameter space. For example, for the N1 pulse, this involves 10 parameters: $I(t) = I_{\text{N1}}(t|A_1, \sigma_{r,1}, \sigma_{d,1}, t_{\text{max},1}, \nu_1) + I_{\text{N1}}(t|A_2, \sigma_{r,2}, \sigma_{d,2}, t_{\text{max},2}, \nu_2)$. In the *lensing* scenario, the second pulse is constrained. It has the same shape parameters as the first pulse, only differing in the normalization (r) and the shift in the peak time (Δt), resulting in $5 + 2$ parameters: $I(t) = I_{\text{N1}}(t|A, \sigma_r, \sigma_d, t_{\text{max}}, \nu) + I_{\text{N1}}(t|A/r, \sigma_r, \sigma_d, t_{\text{max}} + \Delta t, \nu)$. In this case, \mathcal{Z}_{L} is the evidence.

Formally, the Bayesian evidence (or simply evidence) has the following meaning (e.g., Speagle 2020): from Bayes' rule, the probability of the model parameters (in our case, the peak times, pulse widths, amplitudes, etc., denoted by Θ) given the observations (D) and a model (M ; e.g., the lensing scenario using the N1 pulse shape) is $P(\Theta|D, M) = P(D|\Theta, M)P(\Theta|M)/P(D|M)$. Here $P(D|\Theta, M)$ is the likelihood of the data given the model and its parameters, and $P(\Theta|M)$ represents our prior knowledge of the parameters. The evidence is the denominator in the expression of Bayes' rule: $\mathcal{Z} = P(D|M) = \int_{V_\Theta} P(D|\Theta, M)P(\Theta|M)d\Theta$. The integral is performed over the hypervolume V_Θ constructed from all of the parameters.

Calculating the evidence \mathcal{Z} is computationally intensive. It requires integrating the likelihood over a multidimensional parameter space. We use the `bilby` python package (Ashton et al. 2019) and `dynesty` nested sampler (Speagle 2020) to carry out the integration over the parameters to find \mathcal{Z} . As in Paynter et al. (2021), the difference in $\ln \mathcal{Z}$ is the natural logarithm of the Bayes factor ($\ln \text{BF} = \ln \mathcal{Z}_{\text{L}} - \ln \mathcal{Z}_{\text{NL}}$), and it can be used to decide between the models. The Bayes factor is additive; values from different independent (e.g., different energy ranges) measurements can be added to perform model comparison.

We present the results of this analysis in Table 2. Conveniently, the nested sampling also yields the best-fitting flux ratio and time delay. Data from every instrument and

Table 2Bayes Factor Compilation for Different Instruments, Energy Ranges, and Pulse Models Using the `bilby` Nested Sampling Method

Detector	Energy Range (keV)	Pulse Model	ln(BF)	r	Δt (s)
NaI	8–45	N1	0.96	$4.14_{-1.79}^{+3.29}$	$33.85_{-2.66}^{+2.42}$
NaI	45–95	N1	0.50	$5.53_{-1.93}^{+2.67}$	$33.37_{-0.76}^{+1.05}$
NaI	95–300	N1	1.79	$5.59_{-1.27}^{+2.03}$	33.26 ± 0.24
NaI	300–500	N1	0.28	$5.96_{-2.81}^{+2.73}$	$33.68_{-2.39}^{+3.15}$
BGO	120–400	N1	2.29	$4.69_{-1.14}^{+1.91}$	$33.16_{-0.23}^{+0.25}$
ACS	>70	N1	2.89	$4.26_{-0.97}^{+1.71}$	$33.37_{-0.24}^{+0.23}$
Swift	25–350	N1	0.50	$6.02_{-2.67}^{+2.71}$	$33.12_{-2.65}^{+3.55}$
NaI	8–45	N2	1.94	$3.61_{-1.60}^{+3.23}$	$33.25_{-1.30}^{+1.07}$
NaI	45–95	N2	2.83	$5.09_{-1.67}^{+2.48}$	$33.11_{-0.56}^{+0.79}$
NaI	95–300	N2	5.13	$5.35_{-1.24}^{+1.91}$	$33.30_{-0.25}^{+0.23}$
NaI	300–500	N2	0.81	$5.51_{-2.82}^{+2.97}$	$33.11_{-1.30}^{+1.11}$
BGO	120–400	N2	5.25	$4.53_{-1.08}^{+1.80}$	33.13 ± 0.23
ACS	>70	N2	5.18	$4.20_{-0.93}^{+1.69}$	33.34 ± 0.24
Swift	25–350	N2	0.55	$5.72_{-2.71}^{+2.90}$	$32.81_{-1.08}^{+1.26}$

Note. The flux ratio and time delay resulting from the fits are shown in the last two columns.

energy range where the second pulse was detectable provided positive Bayesian evidence in favor of the lensing scenario.

4. Gravitational Lens Modeling

4.1. Point-mass Lens

The simplest mass model is the point-mass lens, when the mass of the lens is concentrated in a projected region smaller than the Einstein radius of the source. In this scenario, we can derive the lens mass from the flux ratio and time delay (e.g., Mao 1992),

$$(1 + z_l)M_l = \frac{c^3 \Delta t}{2G} \left(\frac{r-1}{\sqrt{r}} + \ln r \right)^{-1}, \quad (4)$$

where z_l is the lens redshift, M_l is the lens mass, c is the speed of light, and G is the gravitational constant.

4.1.1. MCMC Light-curve Fitting

We fit the summed NaI and BGO lightcurves with the Markov Chain Monte Carlo (MCMC) method using the `emcee` (Foreman-Mackey et al. 2013) python package with both the N1 and N2 pulses (see Figure 1). This method cannot select between the lensing and no-lensing scenarios; however, it is fast and robust compared to the more computation-intensive nested sampling (see Section 3.5). We can take the result of the light-curve fits in the lensing scenario and derive the lens mass, $M_l(1 + z_l)$, in the point-mass approximation.

The N1 pulse model leads to a flux ratio of $r = 4.47_{-0.73}^{+1.06}$ and delay time of $\Delta t_1 = 33.16_{-0.21}^{+0.19}$ s. The corresponding point-mass lens value is

$$(1 + z_l)M_l = 1.07_{-0.15}^{+0.16} \times 10^6 M_\odot. \quad (5)$$

For the N2 pulse, the flux ratio is $r = 4.19_{-0.25}^{+0.30}$, and the time delay $\Delta t_1 = 33.11 \pm 0.06$ s. These lead to a point-mass lens of

mass

$$(1 + z_l)M_l = 1.13 \pm 0.06 \times 10^6 M_\odot. \quad (6)$$

4.2. Singular Isothermal Sphere Lens Model

This lens model is characterized by the line-of-sight velocity dispersion, σ_v , of its mass distribution. While in the point-mass lens case, we could constrain the lens mass, here it is only possible to restrict the velocity dispersion up to a distance scale D ,

$$\begin{aligned} \sigma_v &= c \left(\frac{1}{32\pi^2} \frac{c\Delta t}{D(1+z_l)} \frac{r+1}{r-1} \right)^{1/4} \\ &\approx 15 \left(\frac{D}{0.6 \text{ Gpc}} \right)^{-1/4} \text{ km s}^{-1}, \end{aligned} \quad (7)$$

where $D = D_{OL}D_{LS}/D_{OS}$ and D_{ij} mark the angular diameter distance combinations between the observer (O), lens (L), and source (S). We assumed a GRB redshift $z_s = 1$ and a lens redshift $z_l = 0.4$.

A simple mass estimate based on the virial theorem yields a mass $M \approx 8 \times 10^5 (\sigma_v/15 \text{ km s}^{-1})^2 (R/10 \text{ pc}) M_\odot$, where $R = 10 \text{ pc}$ is an assumed size considered typical for, e.g., globular clusters for which the singular isothermal sphere (SIS) model is a good approximation. The mass is broadly consistent with the point-mass lens approximation.

5. Discussion

In the previous section, we performed tests to confirm that GRB 210812A is affected by strong gravitational lensing. Here we discuss the strengths of each test, analyze the unlensed GRB properties, and present future detection prospects.

5.1. Spectrum

The most basic spectral test for a lensing scenario is the flux ratio test. Because gravitational lensing is achromatic, the flux ratio of the two pulses has to remain constant across energy ranges and different instruments. This is a robust measure because it does not depend on the assumed spectrum. The only caveat to consider is if the GBM detectors' pointing has changed between the two pulses. In that case, the spectral responses change significantly between the pulses, and the recorded counts cannot be compared across the emission episodes of GRB 210812A. The pointing of the detectors, however, has not changed by more than 5° between the pulses, which means the detectors' response in the direction of GRB 210812A is essentially the same throughout the duration of the burst. The flux ratios across the energy ranges and instruments are clearly consistent with being equal (Figure 6). The weighted mean is 3.45, and we measure the most significant deviation for the ACS data point (green), which is 1.6 standard deviations away.

Mukherjee & Nemiroff (2021c) showed a preliminary analysis of the hardness ratio (HR) for `ctime` channels 3 and 4 (4 and 5 in their notation) and claimed a 2.2σ discrepancy between the HRs of the two pulses. Formally, the CR test is equivalent to the HR test; the HR of pulse 1 is $\text{HR}(P1) = C(P1, \text{Ch2})/C(P1, \text{Ch1})$, where $C(P1, \text{Ch2})$ denotes the count rate in pulse 1 for channel 2 (higher-energy channel for the hardness), and $C(P1, \text{Ch1})$, $\text{HR}(P2)$ can be calculated analogously. The

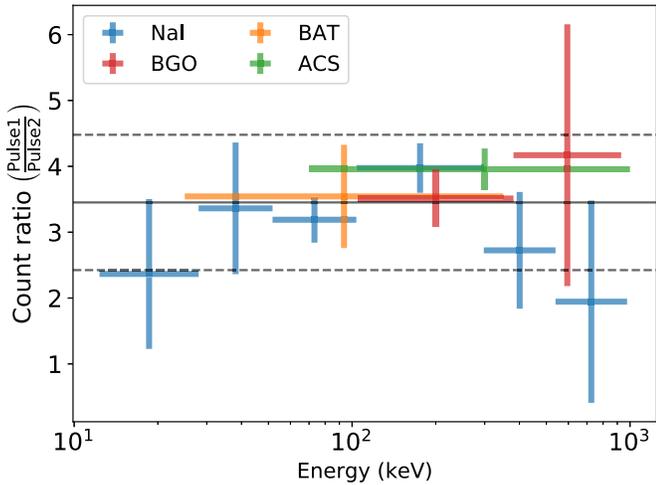


Figure 6. Ratio of counts in the first and second pulses as a function of energy and different instruments. The average CR is 3.45 ± 1.03 (horizontal lines.)

CR in energy channel 1 is $CR(\text{Ch1}) = C(\text{P1}, \text{Ch1})/C(\text{P2}, \text{Ch1})$. Thus, from $CR(\text{Ch1}) = CR(\text{Ch2})$, it follows that $HR(\text{P1}) = HR(\text{P2})$. Because there are no strong outliers in the CR test, we expect the data to reflect this in the HR test. For the first pulse, we find $HR = 1.405 \pm 0.063$, and for the second pulse, $HR = 1.127 \pm 0.155$. We confirm the finding of Mukherjee & Nemiroff (2021c) that the second pulse indeed shows a lower HR. However, taking their difference and adding the errors in quadrature, we find that the discrepancy is only 1.66σ , which does not invalidate the lensing scenario.

Next, we carried out a spectral analysis of the two pulses using the GBM data. In the Swift-BAT data, the second pulse was only visible in the summed lightcurve, and INTEGRAL-SPI/ACS only had data in one energy channel. Therefore, we only used Fermi-GBM data for the detailed spectral analysis. The spectral parameters are consistent within the errors (Table 1), and a plot of the spectral shapes (Figure 4) also shows that the two spectra overlap when considering the confidence regions. We thus conclude that the spectra of the pulses are consistent with the lensing interpretation. We note the advantage of the continuous 128 energy channel data of GBM over the four-channel $\text{t}t\text{e}$ data available for BATSE, for which precise spectral fits were not feasible (Paynter et al. 2021).

5.2. Time History

Gamma-ray instruments have better temporal than spectral resolution. For example, the number of spectral resolution elements in 10–1000 keV is $\lesssim 10$, while the 5 s duration pulse with ~ 0.1 s resolution has 50 temporal resolution elements, where 0.1 s is the approximate timescale of variations for GRB 210812A (Bhat et al. 2012). For this reason, the temporal study of the lensing scenario can provide more constraints.

The weaker second pulse is closer to the noise level of the detectors. Thus, the shorter duration of the second pulse is in line with expectations for a weaker burst with a similar pulse shape. Nonetheless, the duration of the two pulses is still consistent within the errors, as expected from a lensing scenario.

Independent of any pulse models, we first performed a χ^2 test to compare the two lightcurves. The χ^2 test determines if the two lightcurves are consistent with being drawn from the

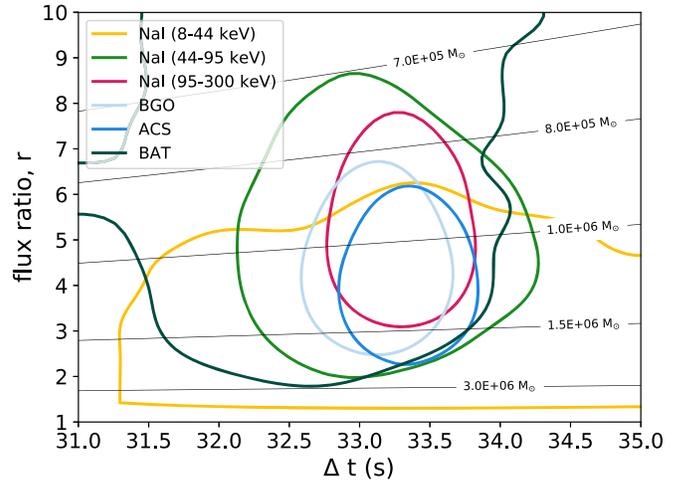


Figure 7. The 1σ confidence region of the time delays and flux ratios using the N2 pulse model and *bilby* nested sampling across energy ranges and instruments. Gray contour lines mark the associated $(1+z)M_l$ lens mass values. Note that the region encompassing $\Delta t \approx 33.3$ s and $r = 4-5$ is consistently part of the solutions.

same distribution. The χ^2 test showed that there is no significant difference between the lightcurves in different energy ranges and across instruments (Figure 3). We note, however, that the weakness of the second pulse results in relatively large Poisson errors, and fine temporal structures in the second pulse, if there are any, are washed out. This somewhat reduces the power of this test, showing only that the general shapes of the pulses are consistent.

Next, we introduced pulse models from the literature and fit the lightcurve in different energy bands and instruments. Using nested sampling, we evaluated the evidence in favor of the lensing scenario using the *bilby* code. Independent of the detector, energy range, or pulse model (Table 2), the evidence is consistently for the lensing scenario, as opposed to the nonlensing scenario ($\text{BF} > 0$ in all cases). We note that the logarithm of the Bayes factor is not necessarily positive for the model with fewer parameters (the lensing model, in our case). Indeed, e.g., Wang et al. (2021; their Table 1) showed some cases with negative $\ln(\text{BF})$.

The Bayes factor differs depending on which pulse model we apply. We find that in energy ranges where the second pulse is relatively weak, the evidence for lensing is not as strong (but it still favors the lensing). The evidence using the N2 pulse model is more compelling than in the case of the N1 shape. This can be due to the larger number of parameters in the case of the N1 pulse shape.

We consider the NaI data alone and the N2 pulse shape fits. The sum of $\ln(\text{Bayes factors})$ for the NaI detectors yields $\ln(\text{BF}) \approx 10.7$. Following Kass & Raftery (1995) and Thrane & Talbot (2019), we can assign colloquial meaning to this number. A difference of more than 8 is considered strong evidence in favor of the lensing model. We conclude that the Bayesian evidence thus supports the lensing interpretation. We note that the BGO and ACS lightcurves also provide additional evidence and, to a lesser extent, the Swift-BAT data as well.

The nested sampling provides a selection criterion between the models through the Bayesian evidence and, at the same time, the parameters for time delay and flux ratio. We show the values in Table 2 and Figure 7. Different instruments, detectors, and energy ranges all yield a consistent solution,

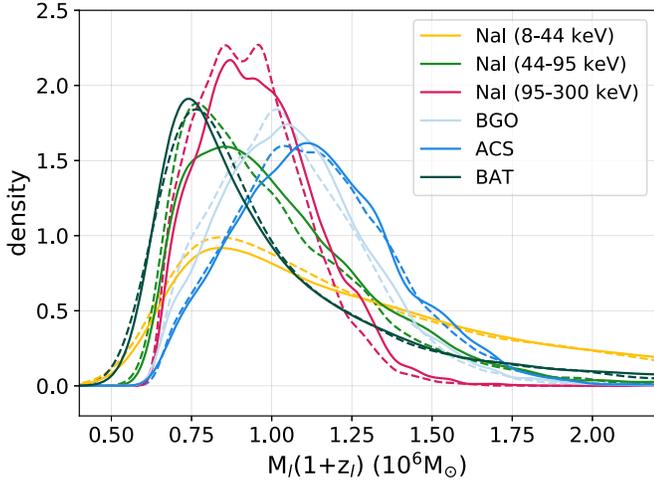


Figure 8. Mass posterior density distribution from the `bilby` modeling assuming a point-mass lens. Different colors show different energy ranges or instruments, as indicated in the legend. Solid lines represent the N2 pulse model, and dashed lines represent the N1 model.

pointing to the lensing origin. Most importantly, we can derive the lens mass for both pulse shapes and arrive at a consistent picture indicating an $\sim 10^6 M_\odot$ lens (Figure 8). The MCMC method yields smaller errors than the nested sampling. This is due to the different fitting approaches of the two methods (see, e.g., Speagle 2020, for details).

5.3. GRB Properties

If GRB 210812A had not been lensed, its fluence would have been $F_0 = F_1 \left(1 - \frac{1}{r}\right) \approx 6.2 \times 10^{-6} \text{ erg cm}^{-2}$, where F_1 is the fluence of the first pulse (see Section 3.2), and we took $r = 4.5$ for the numerical value. Using this flux value, we can get a broad range for the possible redshift of this GRB using empirical correlations between gamma-ray properties.

We scan the 0.1–5 redshift range and find that $z \gtrsim 0.5$ is consistent to within 1σ with the Amati relation (Amati et al. 2002) between the isotropic equivalent energy, E_{iso} , and the redshift-corrected peak energy, $(1+z)E_{\text{peak}}$. While this is a very crude estimate, it is in line with the average measured redshift of GRBs (Bagoly et al. 2006; Jakobsson et al. 2006) and consistent with our fiducial value of $z_s = 1$.

The lag–luminosity relationship (Norris et al. 2000; Gehrels et al. 2006) provides another estimate of the redshift ($L \propto (\tau_{\text{lag}}/(1+z))^{-0.74}$). Taking the median value of the lag (221.6 ms) and scanning the 0.1–5 redshift range, we find that the redshift of GRB 210812A within the range $0.9 < z < 1.5$ is consistent at the 1σ level with the lag–luminosity relation. While this is similarly an empirical relation, it further reinforces that a redshift of $z_s = 1$ is a reasonable approximation.

5.4. Future Events

Even though GRB 210812A had no multiwavelength follow-up, as more lensed GRB candidates are observed, we expect to eventually have a well-localized counterpart. Identifying the afterglow of a lensed GRB showing two consistently fading images would be the smoking-gun evidence for the lensing scenario. The angular separation of the two lensed images on the sky are on the order of the Einstein radius, which, for a $10^6 M_\odot$ point mass, is $\theta_E = ((4GM/c^2)(D_{LS}/D_{OL} D_{OS}))^{1/2} \approx 3 \text{ mas}$

$(M/10^6 M_\odot)^{1/2} (D/0.6 \text{ Gpc})^{-1/2}$ (assuming $z_l = 0.4$ and $z_s = 1$). This falls just short of the resolution of 10 m class optical telescopes ($\approx 40 \text{ mas}$). The only conceivable way of resolving the two images is through very large baseline interferometry (VLBI) radio imaging (Casadio et al. 2021). The sensitivity of VLBI can be as low as $\gtrsim 10 \mu\text{Jy}$ (Venturi et al. 2020), and radio afterglows are detected in significant numbers at or above this flux (Chandra & Frail 2012). The VLBI imaging of the two sources will provide additional information on the source and lens redshift and help constrain the lens model. Capturing a lensed GRB with a well-understood lens will allow one to fully exploit the accurate, subsecond time-delay measurement achievable for GRBs and will allow for time-delay cosmography.

5.5. Lensing Object

It is difficult to identify the type of lens that produced the two pulses of GRB 210812A. Possible objects include BHs or globular clusters. Populations of objects can be ruled out based on their number density and contribution to the total lensing probability (Nemiroff 1989). The total number of millilensed GRBs can be estimated based on a search for lensing candidates in the entire Fermi-GBM GRB catalog. Our goal in this paper was to report only on GRB 210812A, and we leave population-level studies for future work. Nonetheless, we can make a few general observations based on, e.g., the previously mentioned claims by Kalantari et al. (2021) and Wang et al. (2021), Yang et al. (2021). For one to three lens events and the total number of Fermi-GBM GRBs, $N > 3100$, the lensing rate is $(3\text{--}9) \times 10^{-4}$. This is on par with the rate based on BATSE observations by Paynter et al. (2021).

A BH mass of $M \approx 10^6 M_\odot$ lies at the lower end of the supermassive BH population with measured masses (e.g., Woo et al. 2010) and at the upper end of the intermediate-mass BH population (Greene et al. 2020). Without detailed counterpart observations, it is unclear to which group, if any, the lens of GRB 210812A belongs. Further lensed events, however, can provide essential constraints on the rates and origin of BHs in this mass range.

Globular clusters are similarly good lens candidates, and their masses can indeed reach $10^6 M_\odot$ (Baumgardt & Hilker 2018). An SIS model is a good approximation to the velocity dispersion in a globular cluster. Paynter et al. (2021) found, however, that even globular clusters with 1 order of magnitude smaller mass, $\sim 10^5 M_\odot$, do not exist in sufficient numbers to produce the 1/2700 rate of lensed GRBs for BATSE. In our case, we require similar lensing probabilities but with an $\sim 10^6 M_\odot$ globular cluster population. The $10^6 M_\odot$ globular cluster lies above the approximately $2 \times 10^5 M_\odot$ turnover mass in the mass function (Jordán et al. 2007). This means that the larger $10^6 M_\odot$ cannot compensate the drop in number density. We thus conclude that the globular cluster lens can be tentatively ruled out, and a point-mass lens, e.g., a BH, is more likely. For more definitive statements on the nature of the lens, precise localization and high-resolution observations will be necessary.

6. Conclusion

In this paper, we presented multiple lines of evidence for GRB 210812A being gravitationally lensed. The two peaks in GRB 210812A have consistent spectra, time profiles, and spectral evolution. We determined the flux ratio and time delay

with multiple methods and arrived at a consistent picture. The first pulse is approximately 4.5 times brighter, and the delay between the pulses is 33.3 s. Assuming a point-mass lens, this flux ratio and delay correspond to a lens mass of $(1+z_l)M_l = 10^6 M_\odot$. There are only a few unchallenged claims in the literature for lensed GRB lightcurves. GRB 210812A presents the first strong evidence for lensing a long GRB with a flux ratio larger than 2. Future events will benefit from high-resolution radio observations for definitive proof of lensing origin and detailed lens modeling.

The authors thank the anonymous referee for a prompt report. P.V. thanks Michael S. Briggs, Zsolt Paragi, Sándor Frey, and Rosa Leticia Becerra for input and discussions and support from NASA grants 80NSSC19K0595 and NNM11AA01A. N.F. acknowledges financial support from UNAM-DGAPA-PAPIIT through grant IN106521. This work used INTEGRAL public data provided by Dr. V. Savchenko through the INTEGRAL Science Data Centre infrastructure hosted at the University of Geneva.

ORCID iDs

P. Veres  <https://orcid.org/0000-0002-2149-9846>
 N. Bhat  <https://orcid.org/0000-0001-7916-2923>
 N. Fraija  <https://orcid.org/0000-0002-0173-6453>
 S. Lesage  <https://orcid.org/0000-0001-8058-9684>

References

- Ahlgren, B., & Larsson, J. 2020, *ApJ*, **897**, 178
 Amati, L., Frontera, F., Tavani, M., et al. 2002, *A&A*, **390**, 81
 Ashton, G., Hübner, M., Lasky, P. D., et al. 2019, *ApJS*, **241**, 27
 Bagoly, Z., Mészáros, A., Balázs, L. G., et al. 2006, *A&A*, **453**, 797
 Barthelmy, S. D., Barbier, L. M., Cummings, J. R., et al. 2005, *SSRv*, **120**, 143
 Baumgardt, H., & Hilker, M. 2018, *MNRAS*, **478**, 1520
 Beloborodov, A. M., & Mészáros, P. 2017, *SSRv*, **207**, 87
 Bhat, N. P., Meegan, C. A., von Kienlin, A., et al. 2016, *ApJS*, **223**, 28
 Bhat, P. N., Briggs, M. S., Connaughton, V., et al. 2012, *ApJ*, **744**, 141
 Blackburn, L., Briggs, M. S., Camp, J., et al. 2015, *ApJS*, **217**, 8
 Casadio, C., Blinov, D., Readhead, A. C. S., et al. 2021, *MNRAS*, **507**, L6
 Chandra, P., & Frail, D. A. 2012, *ApJ*, **746**, 156
 Connaughton, V., Burns, E., Goldstein, A., et al. 2016, *ApJL*, **826**, L6
 Cucchiara, A., Levan, A. J., Fox, D. B., et al. 2011, *ApJ*, **736**, 7
 Davidson, R., Bhat, P. N., & Li, G. 2011, in AIP Conf. Ser. 1358, Gamma Ray Bursts 1358, ed. J. E. McEnery, J. L. Racusin, & N. Gehrels (New York: AIP), 17
 Fermi GBM Team 2021, GCN, **30633**, 1
 Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *PASP*, **125**, 306
 Gehrels, N., Chincarini, G., Giommi, P., et al. 2004, *ApJ*, **611**, 1005
 Gehrels, N., Norris, J. P., Barthelmy, S. D., et al. 2006, *Natur*, **444**, 1044
 Goldstein, A., Hamburg, R., Wood, J., et al. 2019, arXiv:1903.12597
 Goldstein, A., Veres, P., Burns, E., et al. 2017, *ApJL*, **848**, L14
 Greene, J. E., Strader, J., & Ho, L. C. 2020, *ARA&A*, **58**, 257
 Hakkila, J., Lesage, S., McAfee, S., et al. 2018, *ApJ*, **863**, 77
 Hui, C. & MoonBEAM Team 2021, BAAS, **53**, 315.02
 Hurley, K., Tsvetkova, A. E., Svinikin, D. S., et al. 2019, *ApJ*, **871**, 121
 Jakobsson, P., Levan, A., Fynbo, J. P. U., et al. 2006, *A&A*, **447**, 897
 Ji, L., Kovetz, E. D., & Kamionkowski, M. 2018, *PhRvD*, **98**, 123523
 Jordán, A., McLaughlin, D. E., Côté, P., et al. 2007, *ApJS*, **171**, 101
 Kalantari, Z., Ibrahim, A., Rahimi Tabar, M. R., & Rahvar, S. 2021, arXiv:2105.00585
 Kaneko, Y., Preece, R. D., Briggs, M. S., et al. 2006, *ApJS*, **166**, 298
 Kass, R. E., & Raftery, A. E. 1995, *J. Am. Stat. Assoc.*, **90**, 773
 Keeton, C. R., & Moustakas, L. A. 2009, *ApJ*, **699**, 1720
 Krauss, L. M., & Small, T. A. 1991, *ApJ*, **378**, 22
 Kumar, P., & Zhang, B. 2015, *PhR*, **561**, 1
 Lan, L., Lü, H.-J., Zhong, S.-Q., et al. 2018, *ApJ*, **862**, 155
 Mao, S. 1992, *ApJL*, **389**, L41
 Meegan, C., Lichti, G., Bhat, P. N., et al. 2009, *ApJ*, **702**, 791
 Mukherjee, O., & Nemiroff, R. J. 2021a, *RNAAS*, **5**, 103
 Mukherjee, O., & Nemiroff, R. J. 2021b, *RNAAS*, **5**, 183
 Mukherjee, O., & Nemiroff, R. J. 2021c, GCN, 30842
 Nemiroff, R. J. 1989, *ApJ*, **341**, 579
 Nemiroff, R. J., Marani, G. F., Norris, J. P., & Bonnell, J. T. 2001, *PhRvL*, **86**, 580
 Nemiroff, R. J., Wickramasinghe, W. A. D. T., Norris, J. P., et al. 1994, *ApJ*, **432**, 478
 Norris, J. P., Bonnell, J. T., Kazanas, D., et al. 2005, *ApJ*, **627**, 324
 Norris, J. P., Marani, G. F., & Bonnell, J. T. 2000, *ApJ*, **534**, 248
 Norris, J. P., Nemiroff, R. J., Bonnell, J. T., et al. 1996, *ApJ*, **459**, 393
 Oguri, M. 2019, *RPPh*, **82**, 126901
 Paynter, J., Webster, R., & Thrane, E. 2021, *NatAs*, **5**, 560
 Poolakkil, S., Preece, R., Fletcher, C., et al. 2021, *ApJ*, **913**, 60
 Refsdal, S. 1964, *MNRAS*, **128**, 307
 Rodney, S. A., Brammer, G. B., Pierel, J. D. R., et al. 2021, *NatAs*, **2021NatAs.tmp..164R**
 Rouco Escorial, A., Fong, W., Veres, P., et al. 2021, *ApJ*, **912**, 95
 Schneider, P., Ehlers, J., & Falco, E. E. 1992, Gravitational Lenses (Berlin: Springer)
 Speagle, J. S. 2020, *MNRAS*, **493**, 3132
 Thrane, E., & Talbot, C. 2019, *PASA*, **36**, e010
 Ukwatta, T. N., Stamatikos, M., Dhuga, K. S., et al. 2010, *ApJ*, **711**, 1073
 Venturi, T., Paragi, Z., Lindqvist, M., et al. 2020, arXiv:2007.02347
 Veres, P. & Fermi-GBM Team 2021, GCN, **30645**
 Veres, P., Bagoly, Z., Horvath, I., Meszaros, A., & Balazs, L. G. 2009, arXiv:0912.3928
 von Kienlin, A., Beckmann, V., Rau, A., et al. 2003, *A&A*, **411**, L299
 von Kienlin, A., Meegan, C. A., Paciesas, W. S., et al. 2020, *ApJ*, **893**, 46
 Wang, Y., Jiang, L.-Y., Li, C.-K., et al. 2021, *ApJL*, **918**, L34
 Winkler, C., Courvoisier, T. J. L., Di Cocco, G., et al. 2003, *A&A*, **411**, L1
 Woo, J.-H., Treu, T., Barth, A. J., et al. 2010, *ApJ*, **716**, 269
 Yang, X., Lü, H.-J., Yuan, H.-Y., et al. 2021, arXiv:2107.11050
 Zhang, B.-B., Burrows, D. N., Zhang, B., et al. 2012, *ApJ*, **748**, 132