Does a black-hole merger create time? A bound from GW250114

Rodion Kazennov[1]

[1] *Independent researcher*

(Dated: 3 November 2026)

Muller and Maguire proposed that new time is created together with new space, so that a black-hole merger would delay its own gravitational-wave signal. Their version ties the extra time to a volume and therefore needs a preferred frame. We test a version that does not: the extra time grows with the horizon area, $\Delta t = \beta\sqrt{\Delta A/4\pi}/c$, and switches on when the common horizon forms, a moment we read off numerical-relativity simulations. We search the open LIGO data of GW250114 for such a delay with two waveform families, after checking on real detector noise that the analysis returns no delay when none is present and finds one of the size $\beta = 1$ when it is. We find $\beta < 0.74$ (IMRPhenomXPHM) and $\beta < 0.79$ (NRSur7dq4) at 95% credibility; the natural value $\beta = 1$ is disfavoured at about 2.6σ with both families, and a fluctuating variant with $|\beta| = \sqrt{\pi}$ is excluded. A second event, GW250207, seen here by two detectors only, is consistent with both $\beta = 0$ and $\beta = 1$; combined with GW250114 (IMRPhenomXPHM) it moves the bound from $\beta < 0.74$ to $\beta < 0.79$. The original volume-based delay of about 1.5 ms for a GW150914-like merger would correspond to $\kappa \approx 4$ here and is far outside the posterior. This is a bound on one specific hypothesis, not evidence for it.

I. THE IDEA

Muller and Maguire [1] proposed that the flow of time is the creation of new time along with new space, and that a binary black-hole merger, which creates space, would delay its gravitational-wave signal, by about 1.5 ms for GW150914. Their version ties the new time to a volume. A volume depends on how spacetime is sliced into space and time, so the proposal needs a preferred frame, and tests of Lorentz invariance in the gravitational sector leave very little room for one [9].

We test a version without a preferred frame. The extra time is tied to the growth of the total horizon area. The areas of the initial (well separated) and final (settled) horizons are invariant, so their difference does not depend on a choice of slicing:

$$\Delta t = \beta\,\frac{1}{c}\sqrt{\frac{\Delta A}{4\pi}}, \qquad \Delta A = A_{\rm f} - (A_1 + A_2), \tag{1}$$

with one free number β. In units of the total mass, $\kappa \equiv \Delta t/(GM/c^3)$, with M the total mass in the detector frame, since the delay is measured on the detector's clock. For an equal-mass, non-spinning binary the SXS simulation SXS:BBH:1132 [2] gives $\Delta A = 0.562\,(A_1 + A_2)$ with $A_1 + A_2 = 8\pi M^2$ (in $G = c = 1$), so $\sqrt{\Delta A/4\pi} = 1.06\,M$ and $\kappa = 1.06\,\beta$. This factor is exact only for that configuration; a GW150914-like simulation (SXS:BBH:0305) gives about 1.09. Both events used here have mass ratios near one and small spins, and we use $\kappa = 1.06\beta$ throughout.

We consider two variants: a deterministic one, $\Delta t > 0$ with β of order one, so $\kappa \approx 1.06$ for $\beta = 1$; and a fluctuating one, by analogy with the "everpresent Λ" of causal sets, where Δt has a random sign and a size of order one; we take $|\beta| = \sqrt{\pi}$, $|\kappa| \approx 1.9$, a normalisation chosen by analogy, not predicted. The idea is phenomenological: neither β nor the shape of the onset is derived from a theory.

II. WHEN THE TIME WOULD BE CREATED

We take the moment from numerical relativity. In the SXS catalogue [2] the common horizon appears $7.3\,M$ (SXS:BBH:0305) and about $8\,M$ (SXS:BBH:1132; 7.9 and $8.1\,M$ at two resolutions) before the peak of the wave at infinity, and the area jumps at once by 73% and 63% of its total growth; 99% is reached by $+1.2\ldots1.4\,M$ after the peak. These times are in the simulation's coordinate time, which depends on the gauge. We therefore let the delay switch on as a smooth step,

$$\begin{aligned} h(t) &= h_{\rm GR}\big(t - \Delta t\,S(t)\big), \\ S(t) &= \tfrac{1}{2}\left[1 + \tanh\frac{t - t_{\rm peak} - \delta}{\tau}\right], \end{aligned} \tag{2}$$

with the centre $\delta \in [-10.7, -4.7]\,M$ and the width $\tau \in [1, 2]\,M$: the simulation values with $\pm 3\,M$ of room. Figure 1 shows the effect on the GW250114 waveform.

III. ANALYSIS

a. Data. GW250114 from the Gravitational Wave Open Science Center (GWOSC), Hanford (H1) and Livingston (L1), 8 s at 4096 Hz, analysed from 20 to 1792 Hz. The noise spectrum is a median Welch estimate from data away from the event. The second event, GW250207, is analysed with L1 and Virgo (V1); its H1 data are gated around the event and are not used, so its network signal-to-noise ratio here is about 47 at the best fit against 68.9 in the catalogue.

b. Sampler. bilby [3] with dynesty [4], 500 live points, distance and time marginalised. Detector calibration uncertainty is a cubic spline per detector with 5% in amplitude and 0.05 rad in phase, a generic choice rather than the collaboration's envelopes.

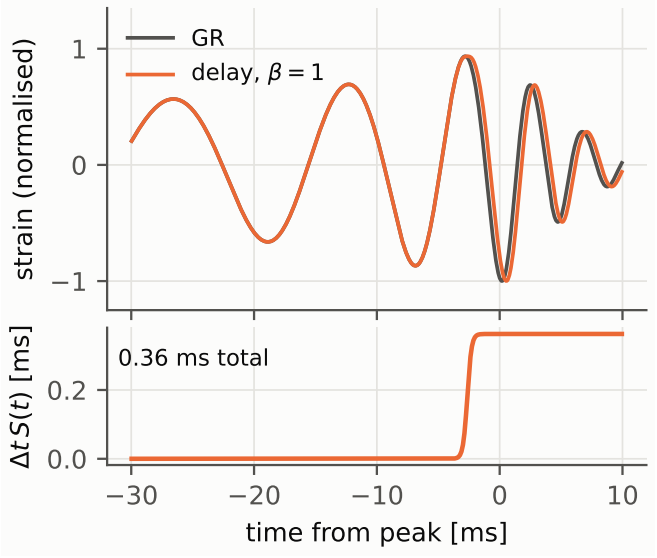


FIG. 1. The model of Eq. (2) on the best-fitting GR waveform of GW250114 (IMRPhenomXPHM, plus polarisation). Grey: GR. Orange: the same signal with the delay of $\beta = 1$ switched on $7.7\,M$ before the peak with width $1\,M$. Lower panel: the delay $\Delta t\, S(t)$ in milliseconds.

c. Priors. κ uniform in $[-15, 15]$; τ log-uniform in $[1, 2]\,M$; δ uniform in $[-10.7, -4.7]\,M$; standard priors on all other parameters.

d. Waveforms. IMRPhenomXPHM [5] and the numerical-relativity surrogate NRSur7dq4 [6]. The surrogate is trained for mass ratios up to 4 and spins up to 0.8; the whole GW250114 posterior lies well inside (mass ratio ≥ 0.74, spins ≤ 0.78). The event just clears the surrogate's length limit for the dominant (2,2) mode at a 20 Hz start ($\gtrsim 66\,M_\odot$ total in the detector frame; the posterior spans 68.6–72.7 $M_\odot$).

e. Reproducibility. Runs were moved between machines only after the log-likelihood at three fixed parameter points agreed with the first machine to better than 10^{-6} relative (2×10^{-8} measured).

IV. TESTS OF THE METHOD

Before the real data, each waveform family and each event had to pass a calibration test: a GR signal ($\kappa = 0$) injected into real detector noise 300 s before the event must return $\kappa = 0$ inside its 99% interval. A power test asks the opposite: a delay with $\beta = 1$ injected into the same noise must be found, with $\kappa = 0$ outside its 99% interval. Table I and Fig. 2 give the results.

All GR injections return $\kappa = 0$ inside their intervals. The $\beta = 1$ injection is found: no posterior sample of 3 971 lies at $\kappa \leq 0$. Its median is 0.27 (about 1.1 standard deviations) above the injected value; with one stretch of noise we cannot tell this apart from a systematic shift. The GW250207 injection used a randomly chosen signal with network SNR about 36, while the real event turned out louder (47 at the best fit). We therefore repeat the test with the injection placed at the real data's best-fitting parameters. It also returns $\kappa = 0$ inside its interval (median -0.29, 90% $[-1.42, +1.19]$), so GW250207 enters the joint bound.

TABLE I. Injections into real detector noise. Intervals are 90%. The best-fit row uses XPHM.

noise, signal, model	$\kappa_{\rm inj}$	median	90%
GW250114, GR, XPHM	0	$+0.08$	$[-0.48, +0.66]$
GW250114, GR, NRSur7dq4	0	-0.06	$[-0.67, +0.49]$
GW250114, $\beta = 1$, XPHM	1.06	$+1.33$	$[+0.90, +1.70]$
GW250207, GR, XPHM	0	$+0.48$	$[-0.83, +1.59]$
GW250207, GR (best fit)	0	-0.29	$[-1.42, +1.19]$

V. RESULTS

Table II and Fig. 3 give the results. For GW250114 the two waveform families agree: we fixed in advance that medians closer than one posterior standard deviation (0.28) count as agreement, and they differ by 0.085. The nuisance parameters τ and δ stay inside their priors; masses and spins are as expected. The tail beyond $\beta = 1$ holds only 14 and 18 of about 3 300 samples, so "2.6σ" should be read as 2.5–2.7σ given sampling noise, and the limits as 0.74 ± 0.01 and 0.79 ± 0.01 (bootstrap); the scatter between independent runs of 500 live points is not included. Across sub-ranges of the onset prior δ the 95% limit on κ lies between 0.72 and 0.85. The fluctuating variant ($|\kappa| \approx 1.9$) lies outside both posteriors: the largest $|\kappa|$ among about 3 300 samples is 1.33 and 1.30, so its posterior probability is below about 3×10^{-4}. A direct comparison of the delay model with GR on GW250114 gives $\ln \mathcal{B} = -2.7 \pm 0.4$ in favour of GR; this number depends on the prior range of κ, and we do not use it for the bound.

GW250207 alone does not distinguish $\beta = 0$ from $\beta = 1$. Its luminosity distance comes out at 577 Mpc against 187 Mpc in the catalogue; in our run distance and inclination are strongly correlated (coefficient 0.91), which we attribute to the missing H1 data, and κ correlates only weakly with distance (coefficient 0.27). Since both runs use the same uniform prior on κ, the posterior of a common κ is the product of the two posteriors (Gaussian kernel estimates). The combined bound is $\beta < 0.79$, with $\beta = 1$ at 2.5σ: the second event, centred higher, weakens the IMRPhenomXPHM bound from 0.74 to 0.79 rather than tightening it.

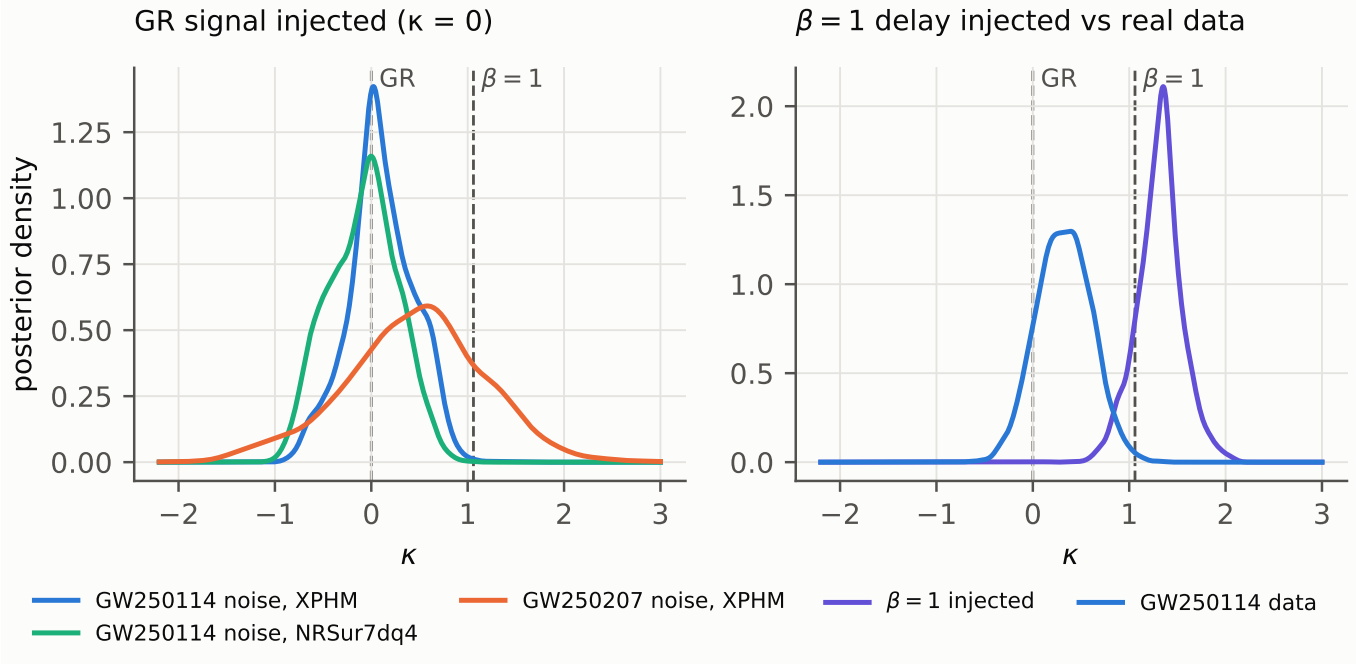


FIG. 2. Tests of the method. Left: κ posteriors of the three GR injections into real noise. Right: the $\beta = 1$ injection into GW250114 noise against the GW250114 data. Injected values are marked.

TABLE II. Real data. κ medians with 68% and 90% intervals, the 95% upper limit on β, and where $\beta = 1$ lies in the posterior (one-sided Gaussian equivalent).

	median	68%	90%	β_{95}	$\beta = 1$
GW250114, XPHM	+0.32	$[+0.04, +0.61]$	$[-0.14, +0.79]$	0.74	2.6σ
GW250114, NRSur	+0.41	$[+0.11, +0.69]$	$[-0.08, +0.84]$	0.79	2.55σ
GW250207, XPHM	+0.78	$[+0.17, +1.32]$	$[-0.26, +1.64]$	1.55	0.5σ
both, XPHM	+0.42	—	$[-0.02, +0.84]$	0.79	2.5σ

VI. HOW MUCH THE TEMPLATE ALONE CAN MOVE κ

Without noise, we injected a GR signal made with one family and profiled the likelihood in κ with the other, at two parameter points with four restarts each. Measured against each family analysing its own signal, IMRPhenomXPHM reads an NRSur7dq4 signal 0.27–0.29 lower in κ; the other way round the shift is 0.05–0.18 (Fig. 4). The real-data posteriors differ by 0.085 in the same direction, which moves the 95% bound on β from 0.74 to 0.79. The largest template-only shift, 0.29, is about one posterior standard deviation; if a shift of that size were present in the real data, $\beta = 1$ would lie at about 1.6σ instead of 2.6σ and the 95% bound would be near $\beta = 1.0$. We therefore quote the significance as statistical only, with the two families as the check on waveform systematics.

VII. WHAT THIS DOES AND DOES NOT SHOW

It is a bound, not a signal: $\kappa = 0$ lies at the 13th and 9th percentile of the two GW250114 posteriors. Each calibration test uses one stretch of noise and gives no false-alarm rate; several noise segments and merged independent runs would make the 2.6σ firmer. The onset is a tanh step, while in numerical relativity it is a jump followed by a few M of slower growth. The detector calibration model is generic.

The closest existing work is the bound of Grimaldi *et al.* [7] on a shift of the merger time for GW250114 in a parametrised effective-one-body model [8], $\delta\Delta t = 0.5^{+9.1}_{-5.8}\,M$ at 90%. That shift moves the point where the ringdown is attached and is a different deformation from ours. Their analysis frees further deviation parameters at the same time, so the widths of the two intervals ($0.93\,M$ for our κ against their $14.9\,M$) are not directly comparable. The inspiral–merger–ringdown consistency test fits the two halves of the signal separately, so it absorbs a constant delay; we have not shown whether it is blind to

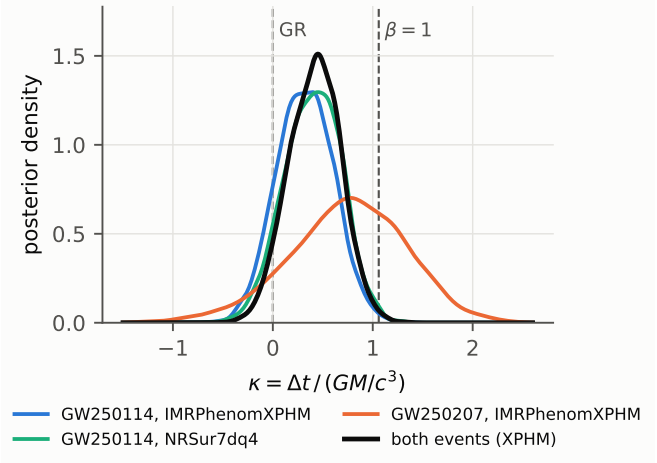


FIG. 3. κ posteriors. GW250114 with both waveform families, GW250207, and their product. Dashed: GR ($\kappa = 0$) and $\beta = 1$ ($\kappa = 1.06$).

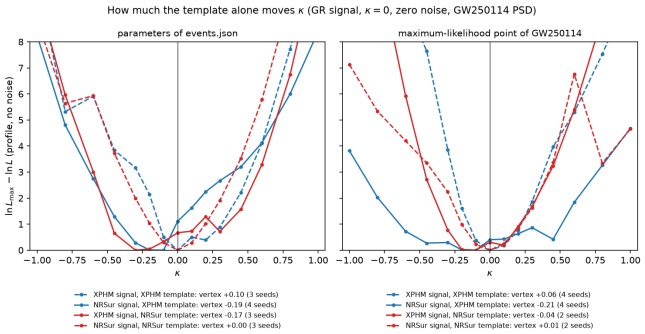


FIG. 4. Profile likelihood in κ for a noiseless GR signal of one family analysed with the other, upper envelope over restarts at two parameter points, with parabola vertices. The profiles are rough at the level of 1–3 in $\ln L$, so the vertices are good to about ± 0.1.

a delay that builds up through the merger.

We found no earlier work that proposes a delay tied to horizon-area growth or bounds one with data (INSPIRE searches). We do not claim the first bound on a merger time shift. Eleven more events in the open catalogue pass a cut of network SNR above 20 inside the surrogate's range; the analysis of the three loudest is under way.

VIII. CONCLUSION

If a binary black-hole merger created time in proportion to the square root of its horizon-area growth, GW250114 limits the constant to $\beta < 0.74$ (IMRPhenomXPHM) or $\beta < 0.79$ (NRSur7dq4) at 95%. The natural value $\beta = 1$ is disfavoured at about 2.6σ (statistical) with both families, and the fluctuating variant with $|\beta| = \sqrt{\pi}$ lies outside both posteriors. Adding GW250207 gives $\beta < 0.79$ with IMRPhenomXPHM.

ACKNOWLEDGMENTS

This research has made use of data obtained from the Gravitational Wave Open Science Center (gwosc.org), a service of the LIGO Scientific Collaboration, the Virgo Collaboration, and KAGRA. We thank the SXS collaboration for the public simulation catalogue.

[1] R. A. Muller and S. Maguire, arXiv:1606.07975 (2016).
[2] M. A. Scheel *et al.*, Class. Quantum Grav. **42**, 195017 (2025), arXiv:2505.13378.
[3] G. Ashton *et al.*, Astrophys. J. Suppl. **241**, 27 (2019), arXiv:1811.02042.
[4] J. S. Speagle, Mon. Not. R. Astron. Soc. **493**, 3132 (2020), arXiv:1904.02180.
[5] G. Pratten *et al.*, Phys. Rev. D **103**, 104056 (2021), arXiv:2004.06503.
[6] V. Varma *et al.*, Phys. Rev. Research **1**, 033015 (2019), arXiv:1905.09300.
[7] L. Grimaldi, E. Maggio, L. Pompili and A. Buonanno, Phys. Rev. D **113**, L061506 (2026), arXiv:2601.13173.
[8] E. Maggio, H. O. Silva, A. Buonanno and A. Ghosh, Phys. Rev. D **108**, 024043 (2023), arXiv:2212.09655.
[9] M. Mewes, Phys. Rev. D **99**, 104062 (2019), arXiv:1905.00409.